



THE EXINDA VIRTUAL  
APPLIANCE GUIDE  
ADMINISTRATION GUIDE

Find out how to set up and configure Exinda Network Orchestrator in different environments and how to customize advanced features.



The information and content in this document is provided for informational purposes only and is provided "as is" with no warranties of any kind, either express or implied, including without limitation any warranties of merchantability, fitness for a particular purpose, and non-infringement. GFI Software disclaims and in no event shall be liable for any losses or damages of any kind, including any consequential or incidental damages in connection with the furnishing, performance or use of this document. The information is obtained from publicly available sources. Though reasonable effort has been made to ensure the accuracy of the data provided, GFI makes no warranty, promise or guarantee about the completeness, accuracy, recency or adequacy of information contained in this document and is not responsible for misprints, out-of-date information, or errors. GFI reserves the right to revise or update its products, software or documentation without notice. You must take full responsibility for your use and application of any GFI product or service. No part of this documentation may be reproduced in any form by any means without prior written authorization of GFI Software.

If you believe there are any factual errors in this document, please contact us and we will review your concerns as soon as practical.

GFI trademarks or registered trademarks of GFI Software or its affiliates in the US and other countries. Any other trademarks contained herein are the property of their respective owners.

Exinda Network Orchestrator Exinda SD-WAN is copyright of Exinda Inc. - 1999-2021 Exinda Inc. All rights reserved.

Document Version: 1.07.4.9

Last updated (month/day/year): 11/25/2021

# Contents

|                                                                          |     |
|--------------------------------------------------------------------------|-----|
| 0.1 How an Exinda implements WAN optimization .....                      | 9   |
| 0.2 What is network orchestration? .....                                 | 9   |
| 0.3 Exinda system components .....                                       | 9   |
| 0.3.1 Exinda Appliance .....                                             | 9   |
| 0.3.2 Exinda Web UI .....                                                | 10  |
| 0.3.3 Exinda Management Center .....                                     | 10  |
| 0.3.4 Exinda Solution Center .....                                       | 10  |
| 0.3.5 Exinda Service Delivery Point (SDP) .....                          | 10  |
| 0.4 The Exinda product line .....                                        | 10  |
| 0.4.1 Exinda physical appliances .....                                   | 10  |
| 0.4.2 Exinda virtual appliances .....                                    | 21  |
| 0.4.3 Exinda Network Expansion Module Comparison .....                   | 21  |
| 0.5 Product naming conventions .....                                     | 24  |
| 0.1 WUI Guided Tour .....                                                | 25  |
| 0.6 Deployment options .....                                             | 26  |
| 0.6.1 Key terms .....                                                    | 26  |
| 0.6.2 Basic characteristics and behaviors of Exinda Appliances .....     | 26  |
| 0.6.3 In-path topologies .....                                           | 27  |
| 0.6.4 Out-of-path topologies .....                                       | 36  |
| 0.6.5 Clustering topologies .....                                        | 65  |
| 0.7 Upgrading and downgrading .....                                      | 69  |
| 0.7.1 Upgrading to the latest firmware version .....                     | 69  |
| 0.7.2 Rolling back to the previously installed version of ExOS .....     | 70  |
| 0.8 Installing an Exinda Appliance .....                                 | 71  |
| 0.8.1 Gathering required information .....                               | 72  |
| 0.8.2 Connecting the appliance to the physical network .....             | 72  |
| 0.8.3 Installing an Exinda Virtual Appliance .....                       | 73  |
| 0.8.4 Scenario .....                                                     | 80  |
| 0.8.5 Use Case .....                                                     | 80  |
| 0.8.6 Scenario .....                                                     | 82  |
| 0.8.7 Use-case .....                                                     | 82  |
| 0.8.8 Use Cases .....                                                    | 83  |
| 0.8.9 Related Topics .....                                               | 97  |
| 0.8.10 Related Topics .....                                              | 97  |
| 0.8.11 Related Topics .....                                              | 107 |
| 0.8.12 Related Topics .....                                              | 113 |
| 0.8.13 Related Topics .....                                              | 114 |
| 0.8.14 Related Topics .....                                              | 121 |
| 0.8.15 Related Topics .....                                              | 126 |
| 0.8.16 Related Topics .....                                              | 127 |
| 0.8.17 Related Topics .....                                              | 129 |
| 0.8.18 Related Topics .....                                              | 131 |
| 0.8.19 Related Topics .....                                              | 131 |
| 0.8.20 Related Topics .....                                              | 139 |
| 0.8.21 Related Topics .....                                              | 140 |
| 0.8.22 Creating an initial configuration using the Basic Wizard .....    | 142 |
| 0.8.23 Licensing information .....                                       | 145 |
| 0.9 Managing multiple appliances with the Exinda Management Center ..... | 150 |
| 0.9.1 Getting started with EMC .....                                     | 151 |

|                                                                                        |     |
|----------------------------------------------------------------------------------------|-----|
| 0.9.2 Deploying Exinda Management Center as a Virtual Machine .....                    | 151 |
| 0.9.3 How EMC fits into the appliance feedback loop .....                              | 153 |
| 0.9.4 Exinda Management Center Concepts .....                                          | 153 |
| 0.9.5 Best Practices .....                                                             | 153 |
| 0.9.6 Configuring the EMC .....                                                        | 155 |
| 0.9.7 Importing appliance configuration .....                                          | 159 |
| 0.9.8 Configuring an appliance manually .....                                          | 164 |
| 0.9.9 Configuring a bridge .....                                                       | 167 |
| 0.9.10 Optimizer Policy Tree .....                                                     | 170 |
| 0.9.11 Sending configuration changes to the appliances .....                           | 176 |
| 0.10 Defining a network environment .....                                              | 178 |
| 0.10.1 Adding network objects .....                                                    | 178 |
| 0.10.2 Working with dynamically created network objects .....                          | 188 |
| 0.10.3 Working with users and groups as objects .....                                  | 190 |
| 0.10.4 Configuring VLAN objects .....                                                  | 192 |
| 0.10.5 Adding protocol objects .....                                                   | 196 |
| 0.10.6 Adding application objects .....                                                | 197 |
| 0.10.7 Adding and updating application group objects .....                             | 200 |
| 0.10.8 Configuring anonymous proxy detection and monitoring .....                      | 206 |
| 0.10.9 Configuring service level agreement objects .....                               | 208 |
| 0.10.10 Creating an HTML response object .....                                         | 213 |
| 0.10.11 Configuring schedule objects .....                                             | 215 |
| 0.10.12 Configuring adaptive response limits .....                                     | 217 |
| 0.10.13 Configuring application performance score objects .....                        | 223 |
| 0.10.14 Configuring an application performance metric object .....                     | 236 |
| 0.11 Monitoring your network .....                                                     | 238 |
| 0.11.1 Dashboards .....                                                                | 239 |
| 0.11.2 Monitoring network traffic in real time .....                                   | 244 |
| 0.11.3 Monitoring network interfaces .....                                             | 256 |
| 0.11.4 Monitoring network throughput .....                                             | 259 |
| 0.11.5 Monitoring service levels .....                                                 | 261 |
| 0.11.6 Monitoring applications .....                                                   | 272 |
| 0.11.7 Monitoring network users .....                                                  | 282 |
| 0.11.8 Monitoring hosts traffic volume .....                                           | 285 |
| 0.11.9 Monitoring network conversations .....                                          | 288 |
| 0.11.10 Monitoring subnets .....                                                       | 291 |
| 0.11.11 Monitoring virtual circuits .....                                              | 296 |
| 0.11.12 Monitoring the effects of controls .....                                       | 299 |
| 0.11.13 Monitoring optimization reports .....                                          | 307 |
| 0.11.14 Monitoring Exinda Appliance system performance .....                           | 312 |
| 0.11.15 Viewing monitoring statistics .....                                            | 321 |
| 0.12 Monitoring applications with the Exinda Solution Center .....                     | 330 |
| 0.12.1 How performance reports work .....                                              | 331 |
| 0.12.2 Using Application Performance reports .....                                     | 332 |
| 0.12.3 Bandwidth usage .....                                                           | 334 |
| 0.12.4 Using the Application Performance Monitor VoIP report .....                     | 335 |
| 0.12.5 Recreational Traffic .....                                                      | 337 |
| 0.12.6 Data Center Continuity .....                                                    | 337 |
| 0.12.7 RIAA Notice Prevention .....                                                    | 338 |
| 0.12.8 Using Network Governance reports .....                                          | 340 |
| 0.12.9 Answers to common questions about Solution Center Application Performance ..... | 341 |
| 0.12.10 Adding and deleting Solutions .....                                            | 342 |

|         |                                                                                    |     |
|---------|------------------------------------------------------------------------------------|-----|
| 0.12.11 | Setting a new baseline                                                             | 343 |
| 0.12.12 | Working with Application Performance charts                                        | 343 |
| 0.12.13 | Investigating a poor application performance score (APS)                           | 344 |
| 0.12.14 | Investigating unusual performance                                                  | 344 |
| 0.12.15 | Deleting an Application Performance report                                         | 344 |
| 0.13    | Managing network traffic                                                           | 345 |
| 0.13.1  | The Exinda policy tree                                                             | 346 |
| 0.13.2  | Circuits                                                                           | 349 |
| 0.13.3  | Virtual Circuits                                                                   | 356 |
| 0.13.4  | Policies overview                                                                  | 371 |
| 0.13.5  | The optimizer wizard                                                               | 389 |
| 0.13.6  | Calculating network performance metrics                                            | 393 |
| 0.14    | Configuring for common use cases and scenarios                                     | 398 |
| 0.14.1  | Monitoring and controlling traffic in a captive portal system                      | 398 |
| 0.14.2  | Backhauling Internet traffic                                                       | 403 |
| 0.14.3  | Setting and enforcing quotas                                                       | 405 |
| 0.14.4  | Creating Applications from DSCP-marked traffic (like Riverbed accelerated traffic) | 409 |
| 0.14.5  | Clustering and high availability                                                   | 410 |
| 0.14.6  | Controlling anonymous proxy traffic                                                | 418 |
| 0.15    | Managing Exinda Appliances with EMC                                                | 419 |
| 0.15.1  | Viewing appliances in the tenancy                                                  | 419 |
| 0.15.2  | Moving appliances within the tenancy                                               | 420 |
| 0.15.3  | Configuration Library                                                              | 422 |
| 0.15.4  | Configuring your Appliances through the CLI                                        | 459 |
| 0.15.5  | Reports in EMC                                                                     | 460 |
| 0.16    | Service Delivery Point (SDP)                                                       | 464 |
| 0.16.1  | SDP Web User Interface (WUI) and Features                                          | 465 |
| 0.16.2  | Getting started with SDP                                                           | 466 |
| 0.16.3  | Changing the default view settings in SDP WUI                                      | 474 |
| 0.16.4  | Tools                                                                              | 479 |
| 0.16.5  | Managing appliances in SDP                                                         | 481 |
| 0.16.6  | The SDP dashboard                                                                  | 486 |
| 0.16.7  | Viewing reports in SDP                                                             | 487 |
| 0.16.8  | Viewing the config log in SDP                                                      | 489 |
| 0.16.9  | Install SDP as a Virtual Appliance                                                 | 490 |
| 0.17    | Introduction                                                                       | 494 |
| 0.17.1  | Features                                                                           | 494 |
| 0.17.2  | Hardware series                                                                    | 494 |
| 0.17.3  | Exinda SD-WAN Management Interface                                                 | 497 |
| 0.17.4  | Maintenance and usage                                                              | 531 |
| 0.17.5  | Abbreviations and acronyms                                                         | 532 |
| 0.17.6  | Support information                                                                | 534 |
| 0.18    | Network settings                                                                   | 535 |
| 0.18.1  | NIC configuration                                                                  | 535 |
| 0.18.2  | IP address configuration                                                           | 537 |
| 0.18.3  | Routes configuration                                                               | 540 |
| 0.18.4  | DNS and domain names configuration                                                 | 541 |
| 0.18.5  | HTTP proxy configuration                                                           | 544 |
| 0.18.6  | Email configuration                                                                | 544 |
| 0.18.7  | SNMP configuration                                                                 | 547 |
| 0.18.8  | Integrate with Active Directory                                                    | 553 |
| 0.18.9  | IPMI Configuration                                                                 | 572 |

|                                                                 |     |
|-----------------------------------------------------------------|-----|
| 0.18.10 Overview of QoS by host .....                           | 576 |
| 0.19 System Setup .....                                         | 600 |
| 0.19.1 Date and Time Configuration .....                        | 600 |
| 0.19.2 UI Access Configuration .....                            | 603 |
| 0.19.3 SDP Configuration .....                                  | 605 |
| 0.19.4 Configure SQL Access .....                               | 605 |
| 0.19.5 Monitoring Configuration .....                           | 621 |
| 0.19.6 Netflow Configuration .....                              | 625 |
| 0.19.7 Create a Scheduled Job .....                             | 628 |
| 0.19.8 Alerts .....                                             | 630 |
| 0.19.9 Control Configuration .....                              | 633 |
| 0.19.10 Disk Storage Explained .....                            | 633 |
| 0.19.11 Virtualization .....                                    | 640 |
| 0.20 Certificates .....                                         | 642 |
| 0.20.1 Managing Certificates and CA Certificates .....          | 642 |
| 0.20.2 View all certificates and private keys .....             | 645 |
| 0.21 Optimization services .....                                | 646 |
| 0.21.1 How Appliance Discovery Works .....                      | 646 |
| 0.21.2 Configuring the Optimization Services .....              | 649 |
| 0.21.3 Universal Acceleration Service .....                     | 650 |
| 0.21.4 Protocol-specific Acceleration .....                     | 654 |
| 0.21.5 Data caching .....                                       | 680 |
| 0.22 Authentication .....                                       | 699 |
| 0.22.1 Display a List of Active Users .....                     | 699 |
| 0.22.2 Local User Accounts .....                                | 700 |
| 0.22.3 AAA .....                                                | 701 |
| 0.22.4 LDAP Authentication .....                                | 702 |
| 0.22.5 Radius Authentication .....                              | 702 |
| 0.22.6 TACACS+ authentication .....                             | 703 |
| 0.23 System Maintenance .....                                   | 703 |
| 0.23.1 Manage System Configuration .....                        | 703 |
| 0.23.2 Factory Defaults .....                                   | 706 |
| 0.23.3 Reboot/Shutdown .....                                    | 706 |
| 0.24 System Tools .....                                         | 708 |
| 0.24.1 Ping .....                                               | 708 |
| 0.24.2 Traceroute .....                                         | 709 |
| 0.24.3 DNS Lookup .....                                         | 709 |
| 0.24.4 Query a remote IPMI Exinda appliance .....               | 710 |
| 0.24.5 iPerf Client .....                                       | 711 |
| 0.24.6 iPerf Server .....                                       | 712 |
| 0.25 Diagnostics .....                                          | 715 |
| 0.25.1 Diagnostics Files .....                                  | 715 |
| 0.25.2 Acceleration Diagnostics .....                           | 716 |
| 0.25.3 Monitor .....                                            | 719 |
| 0.25.4 NIC Diagnostics .....                                    | 720 |
| 0.25.5 Optimizer Diagnostics .....                              | 721 |
| 0.25.6 RAID Diagnostics .....                                   | 722 |
| 0.25.7 TCP Dump .....                                           | 723 |
| 0.25.8 View the status of an alert .....                        | 725 |
| 0.25.9 View the status of the community .....                   | 726 |
| 0.25.10 Open a case with Exinda Networks Support Services ..... | 727 |
| 0.26 Log Files .....                                            | 727 |

|                                                                                                |     |
|------------------------------------------------------------------------------------------------|-----|
| 0.26.1 Viewing System Log Files .....                                                          | 728 |
| 0.26.2 Live Log .....                                                                          | 728 |
| 0.26.3 Tail Log .....                                                                          | 728 |
| 0.26.4 System Logging Configuration .....                                                      | 729 |
| 0.27 Troubleshoot problems with MAPI acceleration .....                                        | 730 |
| 0.27.1 Outlook cannot connect to the Exchange server .....                                     | 731 |
| 0.27.2 Outlook slow to send or receive emails .....                                            | 731 |
| 0.27.3 Decrease in acceleration of MAPI traffic .....                                          | 731 |
| 0.27.4 Reduction ratio for MAPI is different between Client-side and Server-side Exindas ..... | 732 |
| 0.28 Troubleshoot issues with TCP acceleration .....                                           | 732 |
| 0.29 Troubleshoot issues with SMB file acceleration .....                                      | 732 |
| 0.30 Troubleshoot issues with Active Directory configuration .....                             | 733 |
| 0.30.1 Exinda Appliance Reboots Every Night .....                                              | 733 |
| 0.30.2 WMI Service is not running .....                                                        | 734 |
| 0.30.3 System account showing in traffic reports .....                                         | 734 |
| 0.30.4 No Communication Between the Exinda AD Connector and the Exinda Appliance .....         | 734 |
| 0.30.5 Exinda AD Connector stops running .....                                                 | 734 |
| 0.30.6 Excluded Users Still Appear on the Exinda Appliance .....                               | 735 |
| 0.30.7 Changes to the Exinda Active Directory Controller have no effect .....                  | 735 |
| 0.30.8 The IP addresses are not being mapped to the AD users and groups .....                  | 736 |
| 0.31 Troubleshooting Edge Cache .....                                                          | 737 |
| 0.32 Topology troubleshooting .....                                                            | 739 |
| 0.33 Using the Command Line Interface .....                                                    | 740 |
| 0.33.1 Accessing the Command Line Interface .....                                              | 740 |
| 0.33.2 CLI Configuration Jumpstart .....                                                       | 741 |
| 0.33.3 Configure command line options .....                                                    | 742 |
| 0.34 Exinda End User License Agreement (EULA) .....                                            | 744 |
| 0.35 GNU General Public License (GPL) .....                                                    | 745 |
| 0.35.1 Preamble .....                                                                          | 745 |
| 0.35.2 TERMS AND CONDITIONS .....                                                              | 746 |
| 0.36 BSD 2.0 .....                                                                             | 752 |
| 0.37 EMC Notice .....                                                                          | 754 |
| 0.38 Compliances .....                                                                         | 754 |
| 0.38.1 CE .....                                                                                | 754 |
| 0.38.2 FCC Class A .....                                                                       | 754 |
| 0.39 Safety Guidelines .....                                                                   | 754 |
| 0.39.1 Lithium Battery Caution .....                                                           | 755 |
| 0.40 Predefined Applications and Supported L7 Signatures .....                                 | 756 |
| 0.41 Predefined Application Groups .....                                                       | 795 |

# Introduction

Every day critical business network traffic and recreational network traffic compete for bandwidth on strained networks. The Exinda Network Orchestrator inspects, monitors and manages network traffic, maximizing speed and data flow efficiency, giving priority to mission critical business applications across your LANs and WANs.

## 0.1 How an Exinda implements WAN optimization

As soon as you connect it to your network, an Exinda Appliance begins monitoring network traffic and gathering statistics to help you make informed optimization decisions. The Exinda provides a multitude of settings, parameters and tools you can use to tweak and squeeze every last byte of bandwidth from your network hardware.

- » Traffic shaping techniques, classifying and rationing bandwidth in alignment with your company goals and daily needs
- » Intelligent data caching for rapid access to frequently used files and data stores
- » Data deduplication to eliminate redundant data and free bandwidth
- » Network monitoring, analysis and management to identify and limit social network traffic, gaming traffic, streaming traffic and other non-essential traffic types

And getting started with Exinda Network Orchestrator is easy.

First you connect an Exinda Appliance to your network. Next, through the combination of an automatic, intelligent discovery process and manual definitions, the Exinda Appliance learns about your network. Then you specify policies to regulate traffic in your network. After that, you use Exinda's robust set of monitoring tools to gain total insight into the traffic on your network and adjust your policies as needed.

## 0.2 What is network orchestration?

Network orchestration is the idea that networks can be programmed to support applications, giving priority to one over another.

The Exinda Network Orchestrator provides the capability to detect and define data streams according to their origins, destinations and other characteristics. Then it gives you the capability to set up rules governing how much network resources a given data stream is allowed to consume.

An Exinda Network Orchestrator logically transforms your network from a group of disparate routers, hubs, switches, bridges, repeaters and blade servers, working independently, into a single, responsive, service-based asset.

## 0.3 Exinda system components

Exinda includes a number of required and optional components that can be installed in your organization's infrastructure.

### 0.3.1 Exinda Appliance

The Exinda product line includes a series of hardware and virtual network appliances designed to plug directly into your environment with minimal effort. Appliances come in a range of sizes to handle every networking scenario and size, from small offices with dozens of users to very large data centers that support hundreds of thousands.

For more information, refer to [The Exinda product line](#) (page 10).

## 0.3.2 Exinda Web UI

Exinda offers to user and administrator a Web User Interface that allows users to configure policies and monitor the appliances performances through a variety of dashboard and reports.

## 0.3.3 Exinda Management Center

The Exinda Management Center (EMC) provides complete management insight and configuration control of your Exinda Network Orchestrator appliances from one central console. All applications, devices, users, and activities across all network locations are managed from a central location giving IT Administrators the ability to manage network policies and manage appliance configuration across the entire organization.

For more information, refer to [Managing multiple appliances with the Exinda Management Center](#) (page 150).

## 0.3.4 Exinda Solution Center

The Exinda Solution Center provides a series of predefined monitors you can run to generate network performances reports for applications like FTP, SSH, Salesforce.com, Microsoft Office365, VoIP, and many more.

For more information, refer to [Monitoring applications with the Exinda Solution Center](#) (page 330).

## 0.3.5 Exinda Service Delivery Point (SDP)

The Exinda Service Delivery Point (SDP) is a high performance add-on designed for enterprise network environments looking to centrally manage multi-box Exinda deployments. It is available as both a hosted service and a virtual appliance.

SDP simplifies the tasks of installing, configuring, monitoring and reporting WAN optimization appliances. It is a key differentiator in the traffic shaping & WAN optimization space. A fundamental component of Exinda's Unified Performance Management solution, it rounds out the Exinda product line and makes it the most comprehensive and effective solution for achieving peak application performance.

With secure access via a Web browser, SDP subscribers gain full visibility into network usage and control over applications at any WAN site. SDP helps IT managers identify and control the underlying causes of poor network performance, whether it be unwanted recreational peer-to-peer traffic or a misconfigured server.

Custom reports provide a granular analysis of network usage, top applications and top URLs. This information is critical in setting an optimal network policy, throttling back applications and for future capacity planning.

For more information, refer to [SDP Web User Interface \(WUI\) and Features](#) (page 465).

## 0.4 The Exinda product line

The Exinda product line includes a series of hardware and virtual network appliances designed to plug directly into your environment with minimal effort. Appliances come in a range of sizes to handle every networking scenario and size, from small offices with dozens of users to very large data centers that support hundreds of thousands.

### 0.4.1 Exinda physical appliances

The tables below contain the technical specifications, hardware profiles and capacity guidelines for each Exinda Network Orchestrator model listed in order of capacity, starting with the smallest.

## Network Orchestrator 3062 Series



Screenshot 1: Front view of the Exinda 3062.



**A** Serial console (RJ-45 connector)

**B** Dedicated management port and IPMI interface

**C** LAN port of the first bridge (BR2)

**D** LAN port of the second bridge (BR4)

**E** DC power supply inlet

**F** USB ports (reserved)

**G** WAN port of the first bridge (BR2)

**H** WAN port of the second bridge (BR4)

**I** VGA port

Screenshot 2: Rear view of the Exinda 3062.

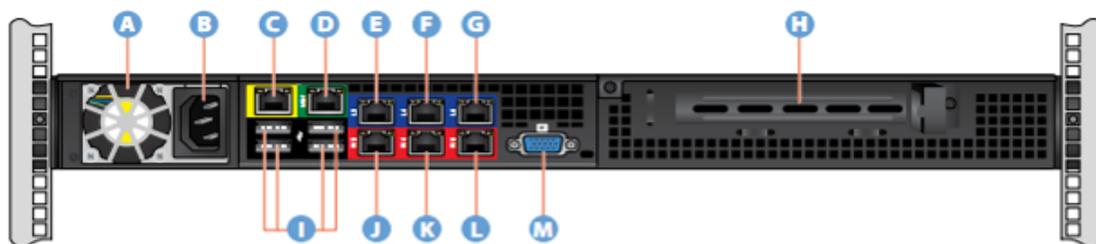
| Specification               | Details      |
|-----------------------------|--------------|
| Designed for                | Small Office |
| Supported Users             | Up to 1,600  |
| <b>Traffic Shaping</b>      |              |
| Shaping Throughput          | 150 Mbps     |
| Concurrent Flow             | 45,000       |
| New Connection Rate         | 4,000/s      |
| Packets Per Second          | 45,000/s     |
| Number of Traffic Policies  | 512          |
| <b>Traffic Acceleration</b> |              |
| Acceleration Throughput     | 20 Mbps      |
| Edge Cache Throughput       | 20 Mbps      |
| Optimized Connections       | 2,000        |
| <b>Network Diagnostics</b>  |              |

| Specification                  | Details                                            |
|--------------------------------|----------------------------------------------------|
| APS Objects                    | 100                                                |
| SLA Objects                    | 100                                                |
| PDF Reports                    | 20                                                 |
| <b>Hardware Specifications</b> |                                                    |
| Form Factor                    | Desktop or 1U rack mount                           |
| Data Store/Cache Size          | 500 GB                                             |
| NICs (Default)                 | 2 Bridge Pairs, or 1 Bridge Pair plus 1 Management |
| NICs (expandable to)           | -                                                  |
| Redundant Power                | No                                                 |

### Network Orchestrator 4062 Series



Screenshot 3: Front view of the Exinda4062.



- |                                                       |                                              |
|-------------------------------------------------------|----------------------------------------------|
| <b>A</b> Power supply module fan                      | <b>H</b> PCI-E card expansion slot           |
| <b>B</b> Power supply AC inlet                        | <b>I</b> USB keyboard ports                  |
| <b>C</b> Serial console (RJ-45 connector)             | <b>J</b> WAN port of the first bridge (BR2)  |
| <b>D</b> Dedicated management port and IPMI interface | <b>K</b> WAN port of the second bridge (BR4) |
| <b>E</b> LAN port of the first bridge (BR2)           | <b>L</b> WAN port of the third bridge (BR6)  |
| <b>F</b> LAN port of the second bridge (BR4)          | <b>M</b> VGA port                            |
| <b>G</b> LAN port of the third bridge (BR6)           |                                              |

Screenshot 4: Rear view of the Exinda 4062.

## Exinda Network Orchestrator 4062 v2

| Specification                  | Details                                                                         |
|--------------------------------|---------------------------------------------------------------------------------|
| Designed for                   | Medium Office                                                                   |
| Supported Users                | Up to 38,000                                                                    |
| <b>Traffic Shaping</b>         |                                                                                 |
| Shaping Throughput             | 1 Gbps                                                                          |
| Concurrent Flow                | 220,000                                                                         |
| New Connection Rate            | 10,000/s                                                                        |
| Packets Per Second             | 200,000/s                                                                       |
| Number of Traffic Policies     | 1024                                                                            |
| <b>Traffic Acceleration</b>    |                                                                                 |
| Acceleration Throughput        | 30 Mbps                                                                         |
| Edge Cache Throughput          | 50 Mbps                                                                         |
| Optimized Connections          | 6,000                                                                           |
| <b>Network Diagnostics</b>     |                                                                                 |
| APS Objects                    | 250                                                                             |
| SLA Objects                    | 250                                                                             |
| PDF Reports                    | 60                                                                              |
| <b>Hardware Specifications</b> |                                                                                 |
| Form Factor                    | Desktop or 1U rack mount                                                        |
| RAM                            | 16GB (2 x 8GB DIMMs)                                                            |
| Data Store/Cache Size          | SSD D3-S4510 960GB                                                              |
| NICs (Default)                 | 3 Bridge Pairs, 1 Management, 1 Cluster (10GbE and 1Gb Fiber options available) |
| NICs (expandable to)           | 5 Bypass Bridges                                                                |
| Redundant Power                | Yes                                                                             |

## Exinda Network Orchestrator 4062

| Specification          | Details       |
|------------------------|---------------|
| Designed for           | Medium Office |
| Supported Users        | Up to 38,000  |
| <b>Traffic Shaping</b> |               |
| Shaping Throughput     | 1 Gbps        |
| Concurrent Flow        | 220,000       |
| New Connection Rate    | 10,000/s      |
| Packets Per Second     | 200,000/s     |

| Specification                  | Details                                                                         |
|--------------------------------|---------------------------------------------------------------------------------|
| Number of Traffic Policies     | 1024                                                                            |
| <b>Traffic Acceleration</b>    |                                                                                 |
| Acceleration Throughput        | 30 Mbps                                                                         |
| Edge Cache Throughput          | 50 Mbps                                                                         |
| Optimized Connections          | 6,000                                                                           |
| <b>Network Diagnostics</b>     |                                                                                 |
| APS Objects                    | 250                                                                             |
| SLA Objects                    | 250                                                                             |
| PDF Reports                    | 60                                                                              |
| <b>Hardware Specifications</b> |                                                                                 |
| Form Factor                    | Desktop or 1U rack mount                                                        |
| Data Store/Cache Size          | 864GB                                                                           |
| NICs (Default)                 | 3 Bridge Pairs, 1 Management, 1 Cluster (10GbE and 1Gb Fiber options available) |
| NICs (expandable to)           | 5 Bypass Bridges                                                                |
| Redundant Power                | Yes                                                                             |

## Network Orchestrator 8063 Series



Screenshot 5: Front view of the Exinda 8062. The hardware of the Exinda 8062 and 8063 is the same.

Screenshot 6:

| Specification          | Details                     |
|------------------------|-----------------------------|
| Designed for           | Small to Medium Data Center |
| Supported Users        | Up to 250,000               |
| <b>Traffic Shaping</b> |                             |
| Shaping Throughput     | 5 Gbps                      |

| Specification                  | Details                               |
|--------------------------------|---------------------------------------|
| Concurrent Flow                | 500,000                               |
| New Connection Rate            | 20,000/s                              |
| Packets Per Second             | 650,000/s                             |
| Number of Traffic Policies     | 2048                                  |
| <b>Traffic Acceleration</b>    |                                       |
| Acceleration Throughput        | 150 Mbps                              |
| Edge Cache Throughput          | 175 Mbps                              |
| Optimized Connections          | 25,000                                |
| <b>Network Diagnostics</b>     |                                       |
| APS Objects                    | 300                                   |
| SLA Objects                    | 300                                   |
| PDF Reports                    | 100                                   |
| <b>Hardware Specifications</b> |                                       |
| Form Factor                    | Desktop or 1U rack mount              |
| Data Store/Cache Size          | 2 TB, RAID 10                         |
| Memory                         | 32 GB                                 |
| NICs (Default)                 | 1 management, 1 cluster, IPMI support |
| Interface NIC Slots            | 1 half height occupied, 1 full height |
| NICs (expandable to)           | 4 bypass bridges                      |
| Redundant Power                | Yes                                   |

## Network Orchestrator 8064 Series



Screenshot 7: Rear view of the Exinda 8064.

Screenshot 8:

| Specification          | Details                     |
|------------------------|-----------------------------|
| Designed for           | Small to Medium Data Center |
| Supported Users        | Up to 250,000               |
| <b>Traffic Shaping</b> |                             |
| Shaping Throughput     | 6 Gbps                      |
| Concurrent Flow        | 625,000                     |

| Specification                  | Details                               |
|--------------------------------|---------------------------------------|
| New Connection Rate            | 25,000/s                              |
| Packets Per Second             | 900,000/s                             |
| Number of Traffic Policies     | 2048                                  |
| <b>Traffic Acceleration</b>    |                                       |
| Acceleration Throughput        | 175 Mbps                              |
| Edge Cache Throughput          | 180 Mbps                              |
| Optimized Connections          | 25,000                                |
| <b>Network Diagnostics</b>     |                                       |
| APS Objects                    | 300                                   |
| SLA Objects                    | 300                                   |
| PDF Reports                    | 100                                   |
| <b>Hardware Specifications</b> |                                       |
| Form Factor                    | Desktop or 1U rack mount              |
| Data Store/Cache Size          | SSD 1.9 TB, RAID 1                    |
| Memory                         | 32 GB                                 |
| NICs (Default)                 | 1 management, 1 cluster, IPMI support |
| Interface NIC Slots            | 1 half height occupied, 1 full height |
| NICs (expandable to)           | 4 bypass bridges                      |
| Redundant Power                | Yes                                   |

## Network Orchestrator 10063 Series



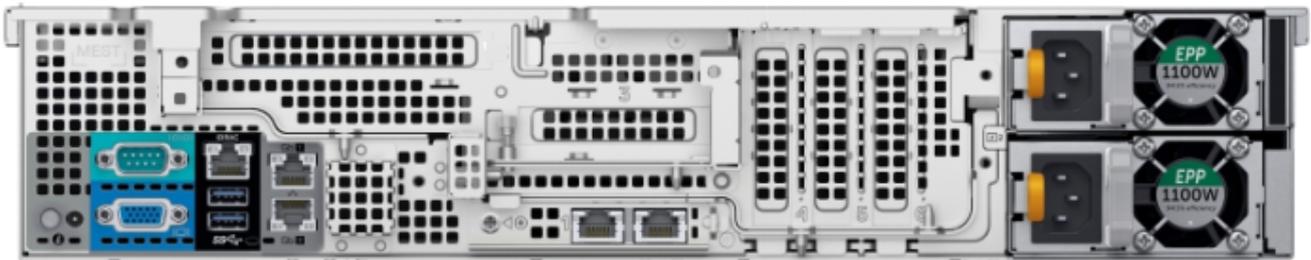
Screenshot 9: Front view of the Exinda 10062. The hardware of the Exinda 10062 and 10063 is the same.



Screenshot 10: Rear view of the Exinda 10062. The hardware of the Exinda 10062 and 10063 is same.

| Specification                  | Details                                           |
|--------------------------------|---------------------------------------------------|
| Designed for                   | Medium to Large Data Center                       |
| Supported Users                | Up to 400,000                                     |
| <b>Traffic Shaping</b>         |                                                   |
| Shaping Throughput             | 10 Gbps                                           |
| Concurrent Flow                | 1,200,000                                         |
| New Connection Rate            | 32,000/s                                          |
| Packets Per Second             | 1,400,000/s                                       |
| Number of Traffic Policies     | 4096                                              |
| <b>Traffic Acceleration</b>    |                                                   |
| Acceleration Throughput        | 500 Mbps                                          |
| Edge Cache Throughput          | 250 Mbps                                          |
| Optimized Connections          | 32,000                                            |
| <b>Network Diagnostics</b>     |                                                   |
| APS Objects                    | 300                                               |
| SLA Objects                    | 300                                               |
| PDF Reports                    | 100                                               |
| <b>Hardware Specifications</b> |                                                   |
| Form Factor                    | Desktop or 2U rack mount                          |
| Data Store/Cache Size          | 1.8 TB, RAID 10                                   |
| Memory                         | 64 GB                                             |
| NICs (Default)                 | 1 management, 3 extra on-board interfaces, 1 IPMI |
| Interface NIC Slots            | 3 half height, 2 full height                      |
| NICs (expandable to)           | 10 bypass bridges                                 |
| Redundant Power                | Yes                                               |

## Network Orchestrator 10064 Series



Screenshot 11: Rear view of the Exinda 10064.

| Specification                  | Details                                                                                                           |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Designed for                   | Medium to Large Data Center                                                                                       |
| Supported Users                | Up to 400,000                                                                                                     |
| <b>Traffic Shaping</b>         |                                                                                                                   |
| Shaping Throughput             | 10 Gbps                                                                                                           |
| Concurrent Flow                | 1,200,000                                                                                                         |
| New Connection Rate            | 32,000/s                                                                                                          |
| Packets Per Second             | 1,400,000/s                                                                                                       |
| Number of Traffic Policies     | 4096                                                                                                              |
| <b>Traffic Acceleration</b>    |                                                                                                                   |
| Acceleration Through-put       | 500 Mbps                                                                                                          |
| Edge Cache Through-put         | 250 Mbps                                                                                                          |
| Optimized Connections          | 32,000                                                                                                            |
| <b>Network Diagnostics</b>     |                                                                                                                   |
| APS Objects                    | 300                                                                                                               |
| SLA Objects                    | 300                                                                                                               |
| PDF Reports                    | 100                                                                                                               |
| <b>Hardware Specifications</b> |                                                                                                                   |
| Form Factor                    | Desktop or 2U rack mount                                                                                          |
| Data Store/Cache Size          | 1.8 TB, RAID 10                                                                                                   |
| Memory                         | 64 GB                                                                                                             |
| NICs (Default)                 | 1 management, 1 extra onboard interfaces that can be increased to 3 if the optional LOM card is installed, 1 IPMI |
| Interface NIC Slots            | 4 half height, 1 full height                                                                                      |

| Specification        | Details           |
|----------------------|-------------------|
| NICs (expandable to) | 10 bypass bridges |
| Redundant Power      | Yes               |

## Network Orchestrator 12063 Series

| Specification                  | Details                                           |
|--------------------------------|---------------------------------------------------|
| Designed for                   | Large Data Center                                 |
| Supported Users                | Up to 600,000                                     |
| <b>Traffic Shaping</b>         |                                                   |
| Shaping Throughput             | 15 Gbps                                           |
| Concurrent Flow                | 1,800,000                                         |
| New Connection Rate            | 38,000/s                                          |
| Packets Per Second             | 1,800,000/s                                       |
| Number of Traffic Policies     | 4096                                              |
| <b>Traffic Acceleration</b>    |                                                   |
| Acceleration Throughput        | 2 Gbps                                            |
| Edge Cache Throughput          | 500 Mbps                                          |
| Optimized Connections          | 49,000                                            |
| <b>Network Diagnostics</b>     |                                                   |
| APS Objects                    | 400                                               |
| SLA Objects                    | 400                                               |
| PDF Reports                    | 150                                               |
| <b>Hardware Specifications</b> |                                                   |
| Form Factor                    | Desktop or 2U rack mount                          |
| Data Store/Cache Size          | SSD 1.6 TB, RAID 2                                |
| Memory                         | 256 GB                                            |
| NICs (Default)                 | 1 management, 3 extra on-board interfaces, 1 IPMI |
| Interface NIC Slots            | 4 half height, 3 full height                      |
| NICs (expandable to)           | 18 bypass bridges                                 |
| Redundant Power                | Yes                                               |

## Network Orchestrator 12064 Series



Screenshot 12: Rear view of the Exinda 12064.

| Specification                  | Details                                          |
|--------------------------------|--------------------------------------------------|
| Designed for                   | Large Data Center                                |
| Supported Users                | Up to 600,000                                    |
| <b>Traffic Shaping</b>         |                                                  |
| Shaping Throughput             | 15 Gbps                                          |
| Concurrent Flow                | 1,800,000                                        |
| New Connection Rate            | 38,000/s                                         |
| Packets Per Second             | 1,800,000/s                                      |
| Number of Traffic Policies     | 4096                                             |
| <b>Traffic Acceleration</b>    |                                                  |
| Acceleration Throughput        | 2 Gbps                                           |
| Edge Cache Throughput          | 500 Mbps                                         |
| Optimized Connections          | 49,000                                           |
| <b>Network Diagnostics</b>     |                                                  |
| APS Objects                    | 400                                              |
| SLA Objects                    | 400                                              |
| PDF Reports                    | 150                                              |
| <b>Hardware Specifications</b> |                                                  |
| Form Factor                    | Desktop or 2U rack mount                         |
| Data Store/Cache Size          | SSD 1.6 TB, RAID 1                               |
| Memory                         | 256 GB                                           |
| NICs (Default)                 | 1 management, 3 extra onboard interfaces, 1 IPMI |
| Interface NIC Slots            | 1 half height, 7 full height                     |
| NICs (expandable to)           | 18 bypass bridges                                |
| Redundant Power                | Yes                                              |

## 0.4.2 Exinda virtual appliances

The virtual Exinda Network Orchestrator provides the same monitoring, reporting and control features as the Exinda hardware appliances. Capacity is determined by a combination of licensing and underlying hardware.

### Exinda Virtual

A virtual Exinda Network Orchestrator runs on a host machine under a hypervisor, using dedicated resources. The minimum dedicated hypervisor hardware requirements are listed in the table below:

| Model      | CPU (# x GHz)                                                   | Memory | Disk Space |
|------------|-----------------------------------------------------------------|--------|------------|
| EXNV-3062  | 4 x 2.0                                                         | 6 GB   | 250 GB     |
| EXNV-4062  | 4 x 2.4                                                         | 8 GB   | 250 GB     |
| EXNV-8063  | 8 x 2.4                                                         | 32 GB  | 500 GB     |
| EXNV-8064  | 8 x 2.4                                                         | 32 GB  | 500 GB     |
| EXNV-10063 | 12 x 2.4 (if license <= 300 M)<br>24 x 2.4 (if license > 300 M) | 64 GB  | 500 GB     |
| EXNV-10064 | 12 x 2.4 (if license <= 300 M)<br>24 x 2.4 (if license > 300 M) | 64 GB  | 1 TB       |
| EXNV-12063 | 12 x 2.4 (if license <= 300 M)<br>24 x 2.4 (if license > 300 M) | 128 GB | 1 TB       |
| EXNV-12064 | 12 x 2.4 (if license <= 300 M)<br>24 x 2.4 (if license > 300 M) | 128 GB | 1 TB       |

Requirements:

- » Intel Xeon class, 64-bit CPU with VT Enabled
- » Hard drive space on a single disk

#### NOTE

Disk extending techniques are not supported on virtual appliances. Adding additional storage requires a hard disk.

## 0.4.3 Exinda Network Expansion Module Comparison

Exinda Network Orchestrator provides a wide range of network modules that can be added to the standard setup, offering greater flexibility and configuration options.

Connector Type RJ-45

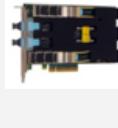
|                | EXN-RJ45-2P-FH 2-port<br>1 Gbps                                                     | EXN-RJ45-4P-HF 4-port<br>1 Gbps                                                     | EXN-RJ45-6P-F 6-port<br>1 Gbps                                                        | EXN-RJ45-2P-10G<br>2-port 10 Gbps                                                     |
|----------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
|                |  |  |  |  |
| Connector type | RJ-45                                                                               | RJ-45                                                                               | RJ-45                                                                                 | RJ-45                                                                                 |

|               | EXN-RJ45-2P-FH 2-port<br>1 Gbps                                                   | EXN-RJ45-4P-HF 4-port<br>1 Gbps                                                   | EXN-RJ45-6P-F 6-port<br>1 Gbps                                                      | EXN-RJ45-2P-10G<br>2-port 10 Gbps                                                   |
|---------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
|               |  |  |  |  |
| <b>4062</b>   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| slot 1        | Yes                                                                               | Yes                                                                               | N/A                                                                                 | N/A                                                                                 |
| <b>8064</b>   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| slot 1 (Half) | Yes*                                                                              | Yes*                                                                              | N/A                                                                                 | Yes*                                                                                |
| slot 2        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| <b>10064</b>  |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| slot 1        | Slot not usable                                                                   |                                                                                   |                                                                                     |                                                                                     |
| slot 2        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 3 (Half) | Yes*                                                                              | Yes*                                                                              | N/A                                                                                 | Yes*                                                                                |
| slot 4 (Half) | Yes*                                                                              | Yes*                                                                              | N/A                                                                                 | Yes*                                                                                |
| slot 5 (Half) | Yes*                                                                              | Yes*                                                                              | N/A                                                                                 | Yes*                                                                                |
| slot 6 (Half) | Slot not usable                                                                   |                                                                                   |                                                                                     |                                                                                     |
| <b>12064</b>  |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| slot 1        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 2        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 3        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 4        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 5        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 6 (Half) | Slot not usable                                                                   |                                                                                   |                                                                                     |                                                                                     |
| slot 7        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |
| slot 8        | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 |

\* The card fits in the half height slot when using the included half height bracket.

N/A - Not available in this slot for this hardware.

## Connector Type LC

|                       | EXN-FIB-2P-MS-FH 2-port 1 Gbps Short Range                                        | EXN-FIB-2P-SL-FH 2-port 1 Gbps Long Range                                         | EXN-FIB-4P-MS-F 4-port 1 Gbps Short Range                                         | EXN-FIB-4P-SL-FI 4-port 1 Gbps Long Range                                         | EXN-FIB10-2P-MS 2-port 10 Gbps Short Range                                        | EXN-FIB10-2P-SL 2-port 10 Gbps Long Range                                           | EXN-FIB40-2P-FH 2-port 40 Gbps Short Range                                          | EXN-FIB10-4P-MS-F 4-port 10 G Short Range                                           |
|-----------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
|                       |  |  |  |  |  |  |  |  |
| <b>Connector type</b> | LC                                                                                | LC                                                                                | LC                                                                                | LC                                                                                | LC                                                                                | LC                                                                                  | MTP                                                                                 | LC                                                                                  |
| <b>4062</b>           |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| slot 1                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | N/A                                                                               | N/A                                                                                 | N/A                                                                                 | N/A                                                                                 |
| <b>8064</b>           |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| slot 1 (Half)         | Yes*                                                                              | Yes*                                                                              | N/A                                                                               | N/A                                                                               | Yes*                                                                              | Yes*                                                                                | N/A                                                                                 | N/A                                                                                 |
| slot 2                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | N/A                                                                                 | Yes                                                                                 |
| <b>10064</b>          |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| slot 1                | Slot not usable                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| slot 2                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 3 (Half)         | Yes*                                                                              | Yes*                                                                              | N/A                                                                               | N/A                                                                               | Yes*                                                                              | Yes*                                                                                | N/A                                                                                 | N/A                                                                                 |
| slot 4 (Half)         | Yes*                                                                              | Yes*                                                                              | N/A                                                                               | N/A                                                                               | Yes*                                                                              | Yes*                                                                                | N/A                                                                                 | N/A                                                                                 |
| slot 5 (Half)         | Yes*                                                                              | Yes*                                                                              | N/A                                                                               | N/A                                                                               | Yes*                                                                              | Yes*                                                                                | N/A                                                                                 | N/A                                                                                 |
| slot 6 (Half)         | Slot not usable                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| <b>12064</b>          |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| slot 1                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 2                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 3                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 4                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 5                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 6 (Half)         | Slot not usable                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |                                                                                     |
| slot 7                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |
| slot 8                | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                               | Yes                                                                                 | Yes                                                                                 | Yes                                                                                 |

\* The card fits in the half height slot when using the included half height bracket.

N/A - Not available in this slot for this hardware.

Download the comparison document [here](#).

## 0.5 Product naming conventions

Encoded within the Exinda Network Orchestrator model numbers are the features and licensing of the appliance.

This is the model number syntax:

<hardware series> <software license> <hardware version>-<bandwidth parameters>

|                                                                       |                                                               |
|-----------------------------------------------------------------------|---------------------------------------------------------------|
| <b>hardware series</b>                                                | The hardware model number.                                    |
| <b>software license</b>                                               | The purchased license.                                        |
| <b>hardware version</b>                                               | The platform configuration version.                           |
| <b>bandwidth optimization   bandwidth acceleration / optimization</b> | The amount of bandwidth for acceleration, visibility and QoS. |

### Example:

Product model 8862-100/500 reflects the following information:

Series - 8000

Software license - x800 (acceleration, visibility and QoS control)

Hardware version - 6.2

Bandwidth - 100 Mbps (100 Mbps for acceleration, 500 Mbps for visibility and QoS)

# Getting started

This getting started guide steps through the basic process of installing, configuring and using your Exinda Network Orchestrator. Each step builds on the previous step and contains links to relevant help topics with detailed instructions to get your Exinda up and running in your network environment.

## 1 Install your Exinda Appliance

- » If you are using a Exinda Appliance hardware box, refer to [Installing the Hardware Appliance into the Network](#).
- » Installing a Exinda Virtual Appliance depends on the hypervisor it supports, refer to [Overview of the Virtual Appliances](#) for installation instructions on supported hypervisors.

## 2 Start monitoring the traffic

Start monitoring the traffic that passes through your Exinda Appliance. For more information, refer to [Monitoring your network](#) (page 238).

If you are looking to monitor particular traffic patterns or usage, you can configure objects to support this.

- » Create network objects to monitor the traffic usage of branches, departments, classes of devices etc. A network object can include one or more subnets and one or more IP addresses. For more information, refer to [Adding network objects](#) (page 178).
- » Configure active directory to identify traffic usage of users on the network. For more information, refer to [Integrate with Active Directory](#) (page 553).
- » Create an application object if a new or custom application is not identified by the system. For more information, refer to [Adding application objects](#) (page 197).

## 3 Configure Traffic Policy

Configuring traffic policy requires a bit of setup, however, the simplest solution is to run the [Optimizer Policy Wizard](#). By answering a few questions in the wizard, the system then sets up a traffic policy that effectively controls the general traffic scenarios. See [Optimizer Policy Tree](#) to understand how the policy configuration works.

You can also customize the traffic policy and have multiple policies in place to match your requirements. For more information, refer to [Policies overview](#) (page 371).

## 4 Create alerts and application performance monitors

Set alerts on various aspects of the traffic. You can monitor the user experience of particular applications and set an alert when the user experience becomes poor. You can monitor the availability of a site by pinging the IP address and define an alert when the latency exceeds your specified threshold or when the packet loss is severe. You can monitor for particular activity that may indicate an issue, such as asymmetric route detection, maximum accelerated connections exceeded, NIC collisions, or dropped packets, and so on.

- » To monitor the user experience of particular applications, create an [Application Performance Score](#) object.
- » To monitor a particular IP address for availability, create an [Site Service Level Agreements](#) object.
- » The alerts are sent by email assuming the appliance is configured to send e-mail. For more information, refer to [Email configuration](#) (page 544).

## 5 Find solutions and get notified

The Exinda appliance displays solutions and notifies you if certain undesirable thresholds are exceeded or if other notable traffic patterns are identified so that you can take action to tune your network.

- » For more information, refer to [Monitoring applications with the Exinda Solution Center](#) (page 330).
- » For more information, refer to [Alerts](#) (page 630).

## 6 Control and accelerate traffic accordingly

Go back and tune traffic policy to control and accelerate traffic to ensure business critical traffic is ensured the bandwidth that is needed, and that the traffic is prioritized and accelerated properly.

## 0.1 WUI Guided Tour

To view an interactive slide show of the UI, go to [Guided Tour](#)

## 0.6 Deployment options

An Exinda Appliance fits almost anywhere in your network environment. As a general rule, anywhere network packets move from one device to another, whether via physical cables or virtualization, you can plug in an Exinda Appliance.

In this section of the guide, you'll walk-through the most common Exinda Appliance deployments.

Many of the example topologies you'll see are used by Exinda customers from various industries around the globe. The list is by no means exhaustive, but you'll undoubtedly find something similar enough to your network environment for you to build on.

### 0.6.1 Key terms

Throughout the deployment options section, you'll see common networking terms you may already be familiar with. Depending on your background and experience, you may have seen the terms used slightly differently than the way they're used in this guide.

To make the deployment examples clear and concise, the terms and definitions are included here.

| Term                     | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| In-path                  | In-path describes a type of topology and refers to deploying an Exinda Appliance between network devices that send and receive data packets to each other, like a switch and a router. When an Exinda Appliance is in-path, it automatically inspects all packets traveling along its path.                                                                                                                                                             |
| Out-of-path              | An Exinda Appliance connected to only one network device is considered out-of-path. By default, that implies network packets don't naturally travel through the Exinda Appliance to get to their destinations. For example, connecting an Exinda Appliance to a switch or hub on the LAN. In this case, the Exinda Appliance behaves like any other network client and requires specific protocols or modes to enable packet monitoring and inspection. |
| Inline                   | In network terminology, an inline device receives packets and forwards them to their intended destination. Routers, firewalls and switches are examples of inline devices. The inline designation also alerts you the device is critical to network function. If the device goes down, network traffic is affected. In an in-path topology, an Exinda Appliance is deployed Exinda Appliancesinline in the network.                                     |
| Clustering               | Grouping Exinda Appliances together in systems to perform the same functions in each system or to mirror each other. Example scenarios include failover, load balancing and multipath networks.                                                                                                                                                                                                                                                         |
| Application Acceleration | Application Acceleration is an intelligent network enhancer based on a set of proprietary algorithms embedded in an Exinda Appliance. It reduces latency, increases network throughput, frees network capacity and a whole lot more. To get the benefits of Application Acceleration, you need at least two Exinda Appliances.                                                                                                                          |

### 0.6.2 Basic characteristics and behaviors of Exinda Appliances

All Exinda Appliance hardware models share some basic characteristics and behaviors. This list provides helpful information to keep in mind while planning and implementing a deployment.

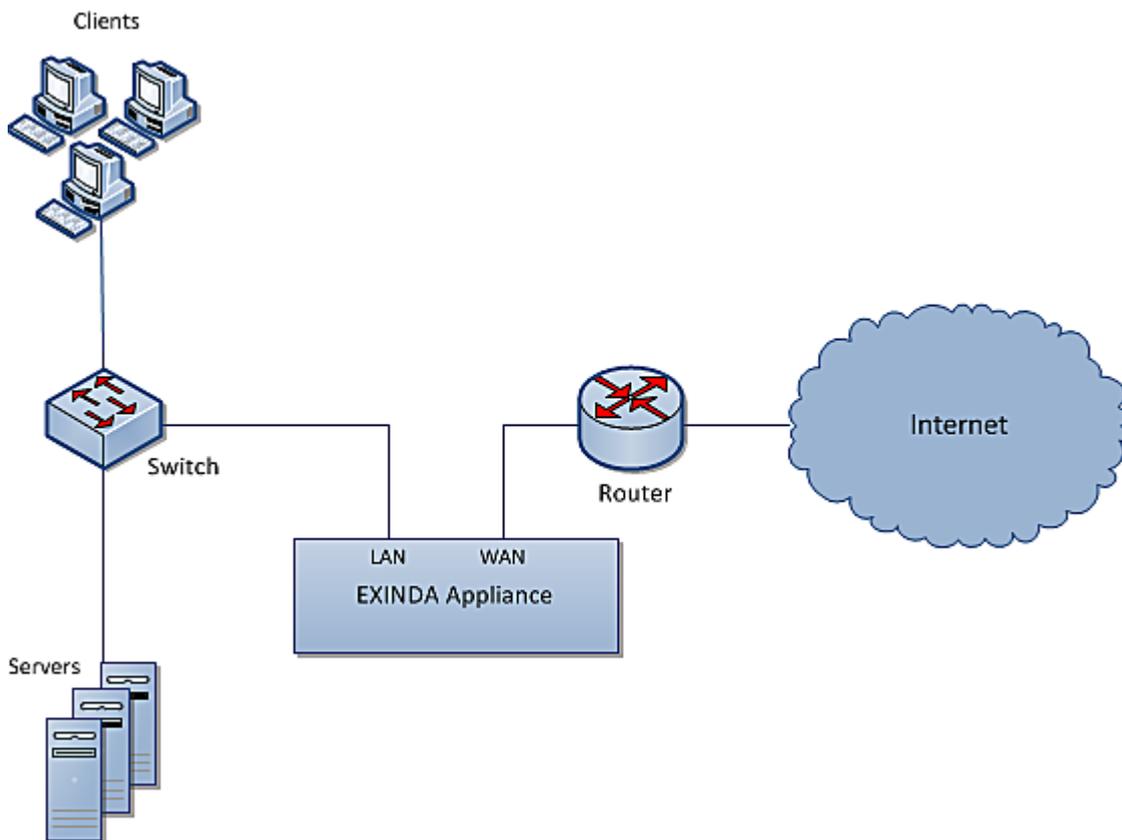
- » Every Exinda Appliance has at least one pair of hardware bypass ports marked LAN and WAN.
- » Exinda Appliance LAN and WAN ports failover to pass-through mode in the event of system failure or power loss.
- » It's best practice to deploy the Exinda Appliance powered off. This ensures the hardware bypass is working.

#### NOTE

There may be a short interruption to network connectivity while the Exinda appliance switches out of bypass mode during boot-up. Although switching in and out of bypass takes less than a millisecond, this may force neighboring equipment to renegotiate their layer 2 topology, which could take several seconds

### 0.6.3 In-path topologies

Exinda Appliances are often deployed between a core switch and a WAN/Internet router. In this type of deployment, the Exinda Appliance is referred to as inline, which is a way to describe a network device in a primary network path that receives packets and forwards them to their destinations. In this case, the Exinda Appliance receives packets from the core switch and sends them to the Internet/WAN router and vice versa.



Screenshot 13: Inline deployment of an Exinda Appliance

Usually, the WAN port on your Exinda appliance is cabled to the WAN/Internet router, using the crossover Ethernet cable. And the LAN port on your Exinda appliance is cabled to the core switch, using the straight Ethernet cable. If your appliance has a dedicated management port, it also needs to be cabled to an internal switch using an Ethernet cable. Both cables are shipped along with the appliance.

For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

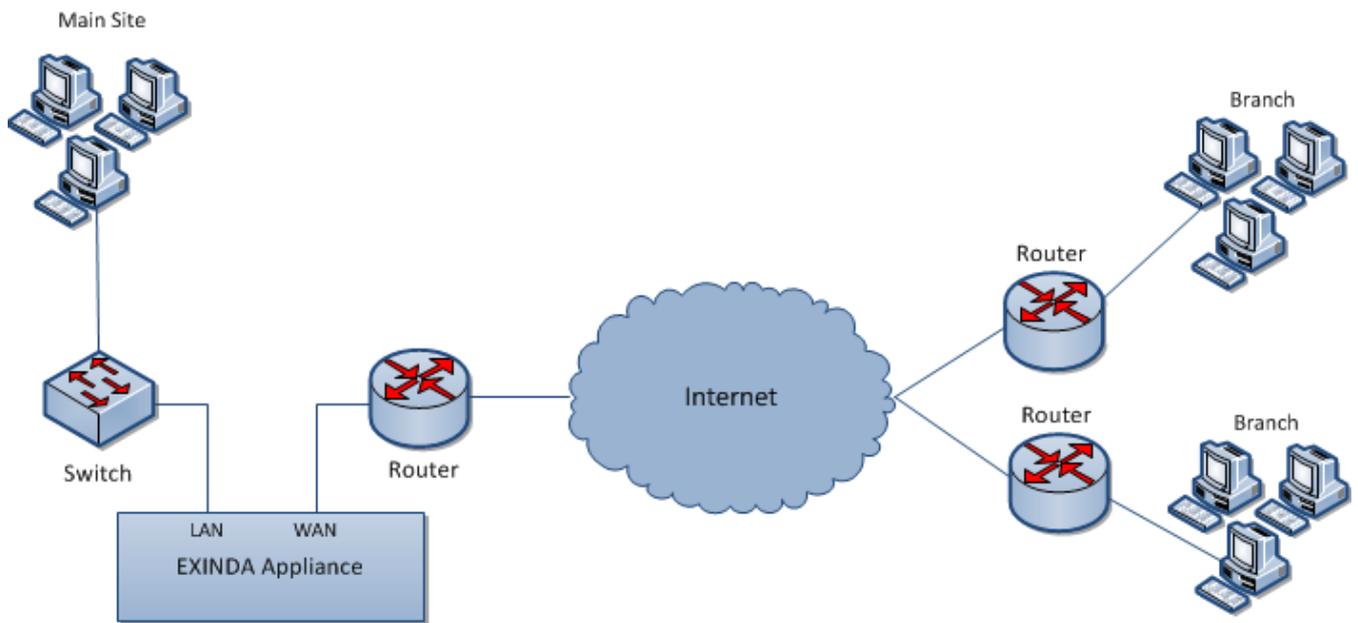
For specific information about your model, [download](#) its Quick Start Guide.

Once all Ethernet cables are in place, power the Exinda Appliance off and ensure the network connectivity. Then, power on the Exinda Appliance, let it fully boot and ensure network connectivity.

The following topics describe how to configure and operate your Exinda Appliance within various in-path topologies.

#### Main site Internet link topology: single site with one Exinda Appliance

A simple way to monitor network traffic between the Internet and your local network, is to plug-in your Exinda Appliance on the network path between your network users and the router, firewall or other gateway device that controls access from your LAN to the Internet.



Screenshot 14: Main site internet link deployment

The topology depicted in the diagram shows a basic Exinda Appliance network deployment with several sets of users from a Main Site and Branch sites linking over the Internet. In Exinda terminology, this general configuration is aptly named, Main Site Internet Link topology.

The left side of the diagram labeled "Main Site" represents a typical office environment, comprising network users, a switch, a router and an Exinda Appliance.

The middle of the diagram represents the Internet and the right side of the diagram represents all the people, machines and programs that want access to servers and applications hosted at the Main Site.

In this setup, the Exinda Appliance is connected to the switch and the router on the Main Site. So network traffic emanating from the Main Site to the Internet and traffic from the Internet to the Main Site must pass through the Exinda Appliance.

With the Exinda Appliance deployed between the switch and the router, you get visibility to all the traffic entering and leaving the Main Site network via the router.

### Installing the Exinda Appliance in a main site internet link topology

This install is straightforward and requires just a few steps.

The high level plan is to plug your Exinda Appliance inline between the switch and router.

In most network environments, the switch will already be physically connected to the router. We're going to temporarily disconnect the switch and router from each other, insert the Exinda Appliance between them and reconnect everything with the Exinda Appliance firmly in the middle.

#### NOTE

In network terminology, an "inline" device receives packets and forwards them to their intended destination. Routers, switches and firewalls are examples of inline devices. The inline designation also alerts you that the device is critical to network function. If the device goes down, network traffic is affected.

1. Connect the WAN port to your router/firewall using a crossover cable.
2. Connect the LAN port to the LAN switch.
3. Leave the Exinda Appliance powered off.

4. Check for Internet connectivity.

5. Turn on the Exinda Appliance and check for Internet connectivity again.

There are a few Exinda Appliance basics to keep in mind while planning a deployment. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

### Capabilities provided by a main site internet link topology

In a Main Site Internet Link topology, an Exinda Appliance:

- » Monitors all network traffic going to and returning from the Internet/WAN.
- » Monitors application specific network traffic going to and returning from the Internet.
- » Controls network traffic on the Main Site by allocating bandwidth to Main Site network users and setting priorities for Internet-based applications accessed by Main Site network users.
- » Monitors and controls network traffic from Branch users and the Internet to servers and applications hosted on the Main Site.

### Limitations of the main site internet link topology

- » Monitoring and controlling Branch Site traffic to the Internet is impossible because the Exinda Appliance is not inline between the branch offices and the Internet. If you want to monitor and control network traffic for a geographically dispersed user base that connects to the Internet through multiple routers outside your LAN, use a different deployment strategy.
- » One of Exinda's top features is Application Acceleration. With only one Exinda Appliance deployed, Application Acceleration isn't available because it requires at least two Exinda Appliances.

### Overcoming the limitations of the main site internet link topology

- » If you only have one Exinda Appliance but need to monitor and control network traffic to the Internet from users outside the Main Site, disable direct access to the Internet for branch offices and route all Internet-bound traffic through the Main Site.

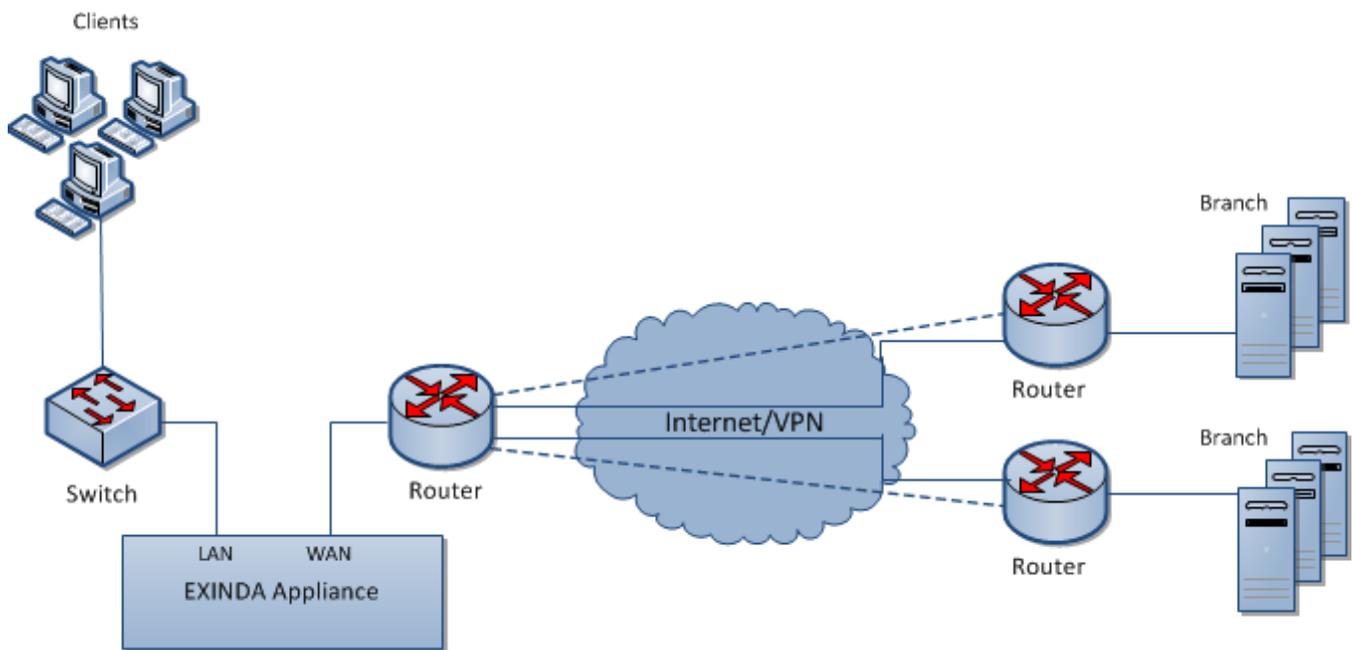
#### NOTE

Routing Internet-bound traffic from users outside your Main Site through the Main Site router introduces an extra leg of travel for network packets and could cause a spike in the Main Site's router utilization.

- » Install an Exinda Appliance at each branch. Essentially, you replicate the Main Site Internet Link topology at each branch. Installing multiple Exinda Appliances into your network environment allows you to implement Exinda's proprietary, performance-boosting Application Acceleration technology.

### Main site WAN link topology: multiple sites with one Exinda Appliance

This topology is used to monitor and control Internet and WAN traffic in the main site and WAN traffic from the branch offices. The Exinda Appliance monitors and controls Internet and WAN traffic on the main site and WAN traffic from branch offices. The Exinda guarantees traffic for the WAN and treats applications and users from different branch offices with different priorities.



Screenshot 15: Main Site WAN Link topology deployment

### Installing the Exinda Appliance in a main site WAN link topology

Plug your Exinda Appliance in line between the switch and router or firewall. If you have a VPN, refer to [Topologies with VPNs](#).

1. Connect the WAN port to your router/firewall using a crossover cable.
2. Connect the LAN port into the LAN switch.
3. Leave the Exinda Appliance powered off.
4. Check for Internet connectivity.
5. Turn on the Exinda Appliance and check for Internet connectivity again.

There are a few Exinda Appliance basics to keep in mind while planning a deployment. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

### Capabilities provided by the main site WAN link topology

In this topology, the Exinda appliance:

- » Monitors all traffic utilization and all applications to the Internet. You can distinguish between business relevant traffic and traffic used for recreational purposes.
- » Monitors usage of Internet and WAN traffic, e.g., how much of the link is being used by the Internet and each branch office?
- » Monitors and controls individual applications and users from each branch office.
- » Controls all traffic traversing the link. Allocate bandwidth to WAN and Internet applications.

### Limitations of the main site WAN link topology

- » Application Acceleration is not possible with a single appliance.
- » If a branch office connects to the Internet directly, the branch link cannot be monitored and controlled.

## Overcoming the limitations of the main site WAN link topology

» If you need to monitor and control network traffic to the Internet from users outside the Main Site, disable direct access to the Internet for branch offices and route all Internet-bound traffic through the Main Site.

### NOTE

Routing Internet-bound traffic from users outside your Main Site through the Main Site router introduces an extra leg of travel for network packets and could cause a spike in the Main Site's router utilization.

» Install an Exinda Appliance at each branch. Essentially, you replicate the Main Site Internet Link topology at each branch. Installing multiple Exinda Appliances into your network environment allows you to implement Exinda's proprietary, performance-boosting Application Acceleration technology.

## Distributed branch topology: multiple sites with multiple Exinda Appliances

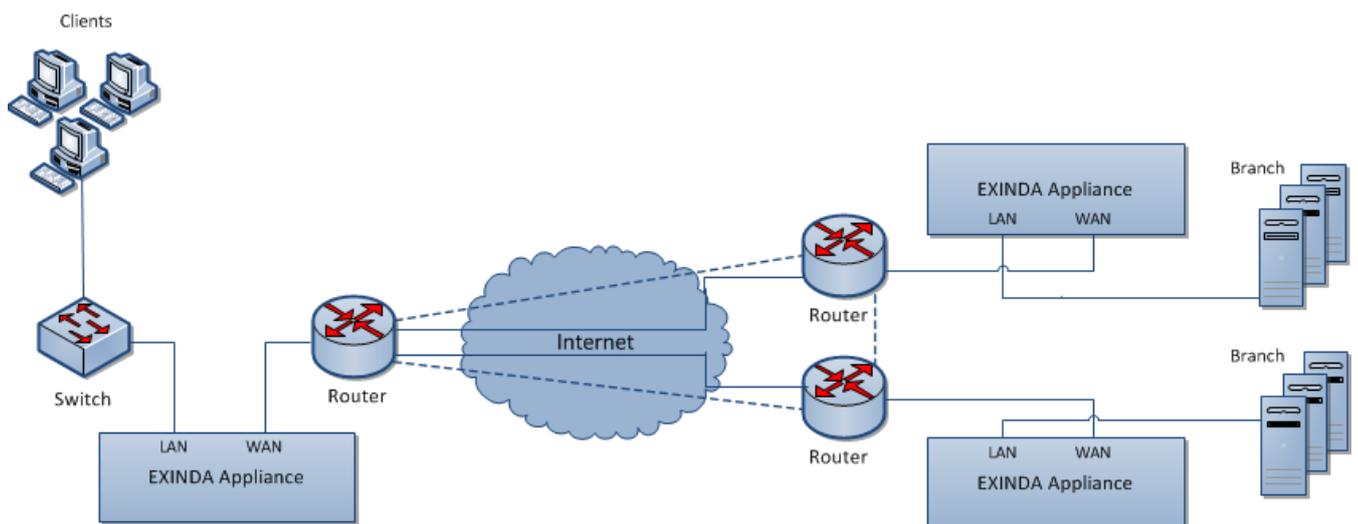
Deploying multiple Exinda Appliances provides a lot of flexibility for monitoring and shaping network traffic across multiple network sites. You also gain the ability to enable Exinda's Application Acceleration feature if your Exinda Appliance model supports it.

This topology requires at least two Exinda Appliances. In the network diagram below, there are three Exindas, once at each site. The sites are connected to the Internet and to each other via direct routes.

All the network traffic at each site passes through an Exinda Appliance on its way to and from the Internet and WAN. The Exindas may work together to enhance network performance and accelerate specific network traffic.

### NOTE

An acceleration license is required to take advantage of Application Acceleration. Check the [Exinda hardware appliances and technical specifications](#) to make sure your Exinda Appliance model supports Application Acceleration.



Screenshot 16: Distributed Branch topology deployment

## Installing the Exinda Appliance in a distributed branch topology

At each site:

1. Connect the WAN port to your router/firewall using a crossover cable.
2. Connect the LAN port to the LAN switch.
3. Leave the Exinda Appliance powered off.
4. Check for Internet connectivity.
5. Turn on the Exinda Appliance and check for Internet connectivity again.

There are a few Exinda Appliance basics to keep in mind while planning a deployment. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

### Capabilities provided by a distributed branch topology

In this topology, the Exinda appliances:

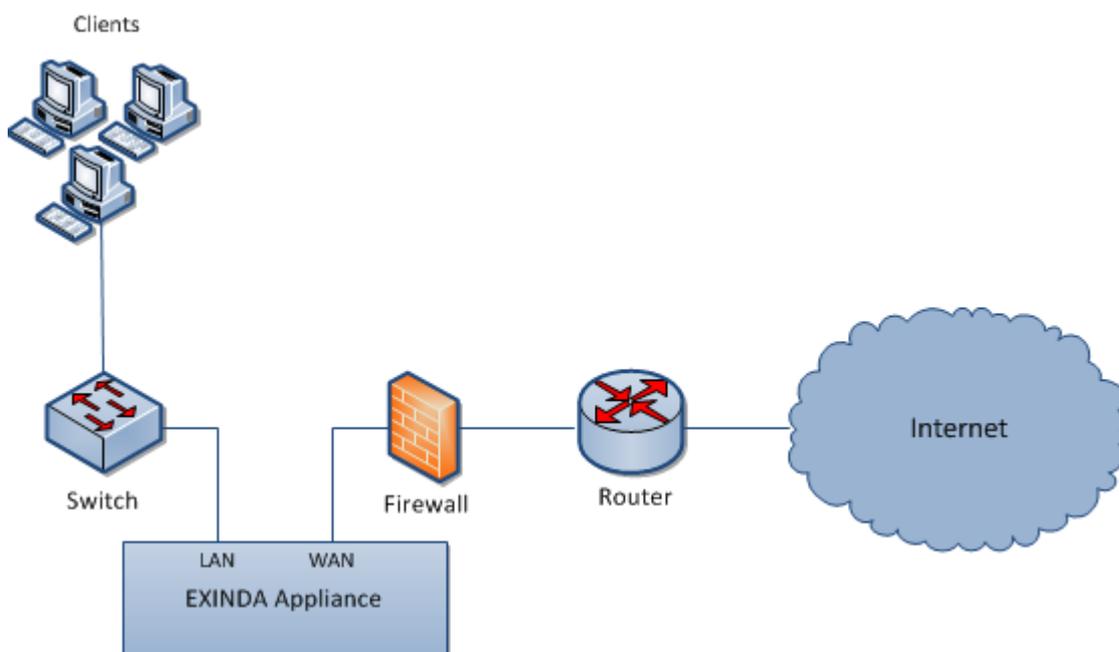
- » Monitor and control all traffic to and from the Internet and WAN.
- » May accelerate traffic between all WAN sites.
- » Monitor distribution of application traffic between all sites.
- » Prioritize and manage application performance in a fully meshed environment.
- » Control or block P2P and recreational applications site-wide.

### Limitation of the distributed branch topology

None. This is the most flexible and robust topology.

### Topologies with firewalls: Exinda Appliance in a network with a firewall

Firewall topologies vary significantly. Typically, an Exinda appliance is deployed between the switch and internal interface of the firewall. This ensures the Exinda Appliance sees all hosts on the LAN.



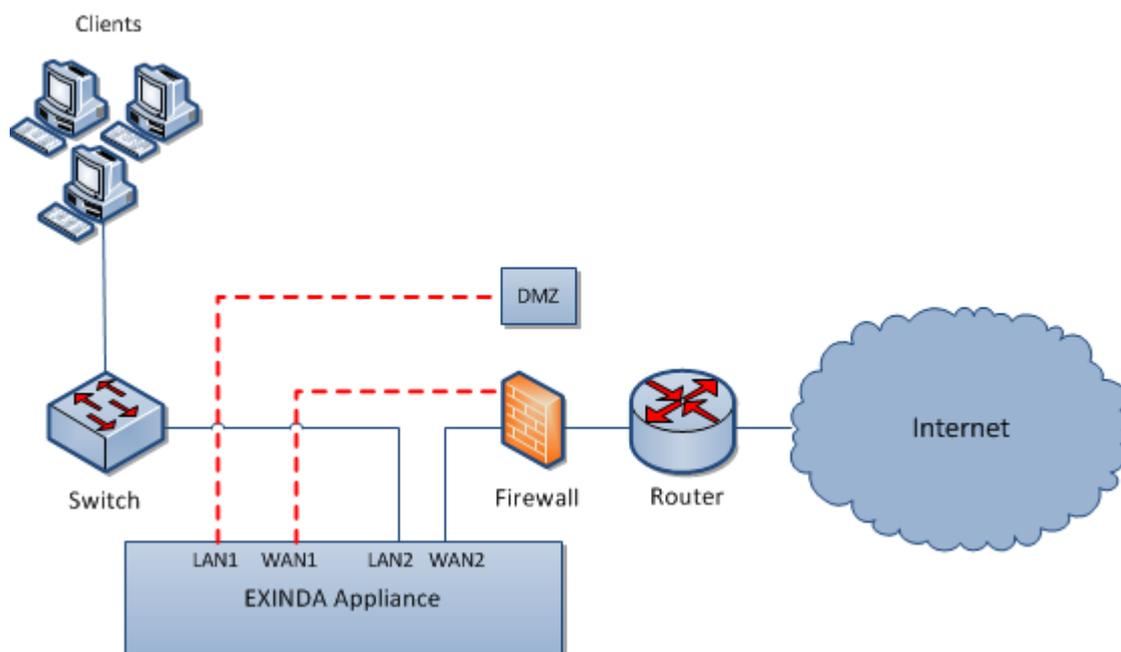
Screenshot 17: Topology with firewalls

## NOTE

Placing the Exinda Appliance between the router and external interface of the firewall will only monitor applications and IP addresses present on the external interface of the firewall. So, if your firewall performs Network Address Translation (NAT), the Exinda will only see the firewall's external IP address as the source address of the monitored flows.

## Deploying an Exinda Appliance in-path of a DMZ

The Exinda appliance can be deployed in-path of a DMZ, allowing monitoring, optimization and Application Acceleration of traffic to and from the DMZ.



## NOTE

Define a network object called DMZ and mark it as "Internal," so the Exinda appliance can ignore all traffic between the local LAN and the DMZ.

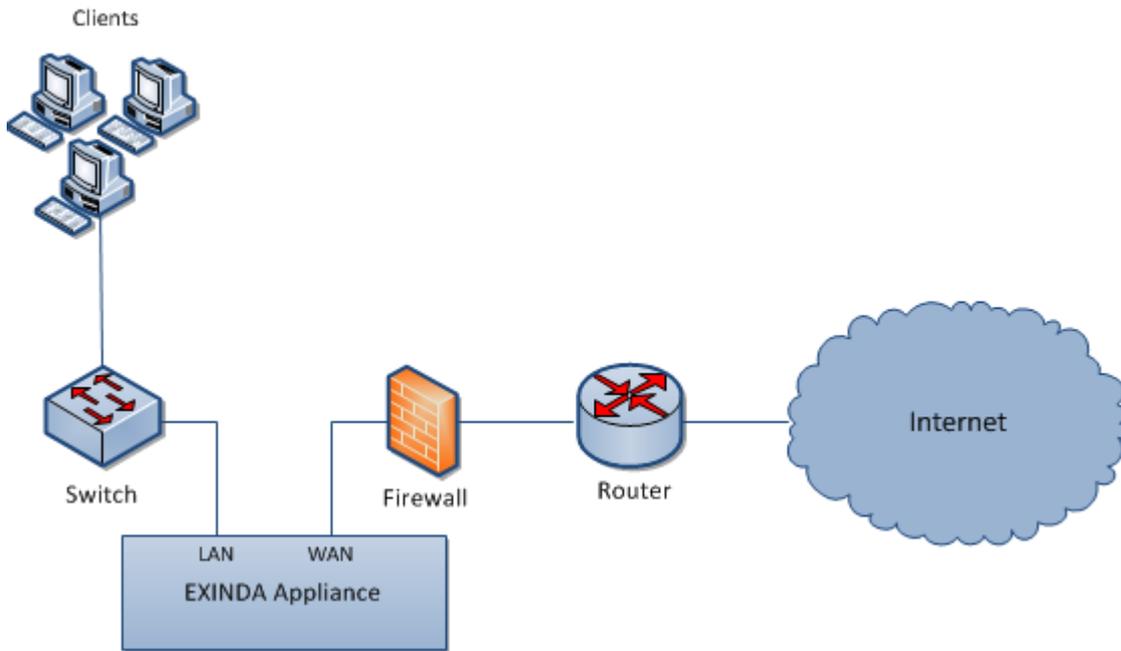
## Installing an Exinda Appliance in a network environment with a firewall

1. Enable the appropriate bridges on the IP Address configuration page.
2. Connect Exinda WAN2 into your router/firewall using a crossover cable.
3. Connect Exinda LAN2 into the LAN switch.
4. Connect Exinda LAN1 into the DMZ switch or host.
5. Connect Exinda WAN1 in the DMZ interface of the firewall using a crossover cable.

There are a few Exinda Appliance basics to keep in mind while planning a deployment. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

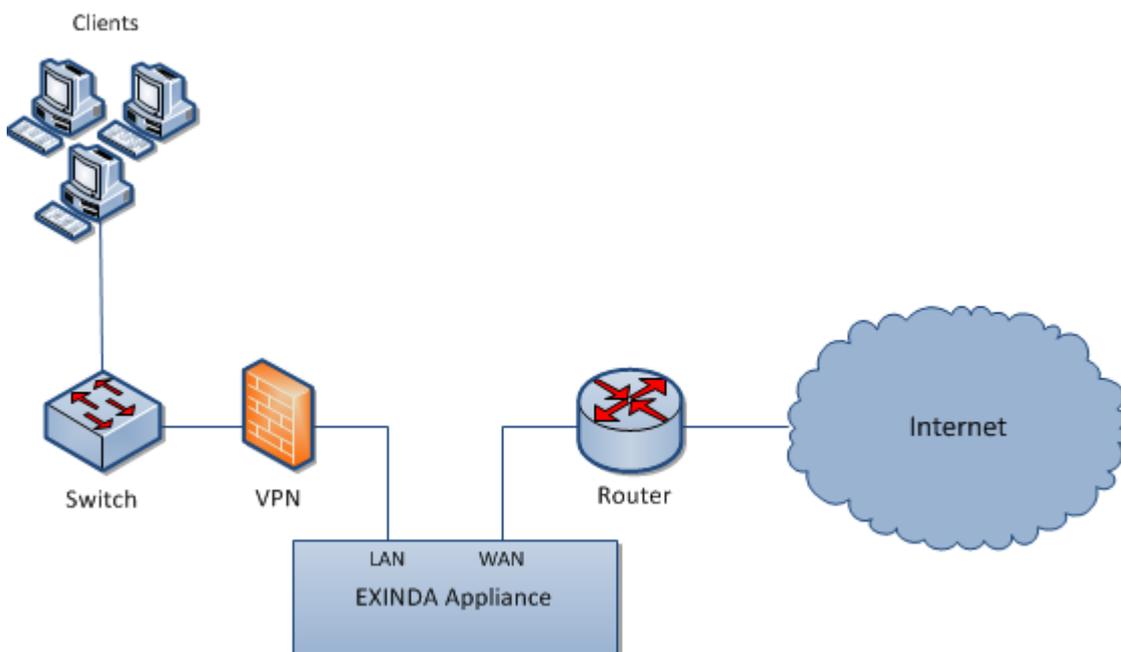
## Topologies with VPNs: Exinda Appliance in a network with a VPN

The typical deployment places an Exinda Appliance between an internal LAN switch and VPN terminator. This allows the Exinda to monitor and optimize traffic before it's encrypted and transported across the VPN tunnel.



Screenshot 18: Topologies with VPNs deployment - Scenario 1

In scenarios where the Exinda Appliance is placed between the VPN terminator and the router, the Exinda Appliance sees only encrypted tunnel traffic.



Screenshot 19: Topologies with VPNs deployment - Scenario 2

### Installing an Exinda Appliance in a network environment with a VPN

#### Scenario 1:

1. Connect the Exinda WAN port into the internal interface of the VPN terminator using a crossover cable.
2. Connect the Exinda LAN port into the LAN switch.

#### Scenario 2:

1. Connect the Exinda WAN port into the internal interface of the router.
2. Connect the Exinda LAN port into the external interface of the VPN terminator using a crossover cable.
3. Connect an Exinda unbridged interface (e.g., eth1 on a 4060) into the LAN switch and configure an address to manage the appliance.

There are a few Exinda Appliance basics to keep in mind while planning a deployment. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

### Capabilities of an Exinda Appliance in a network environment with a VPN

In VPN scenario 2, the Exinda Appliance:

- » Monitors and controls any unencrypted traffic to the WAN and Internet.
- » Monitors and prioritizes encrypted traffic between other VPN terminator sites. Only a single IP address will be visible per site.

### Limitations of an Exinda Appliance in a network environment with a VPN

In VPN scenario 2 the Exinda appliance cannot monitor and prioritize the encrypted traffic by application, internal hosts and servers.

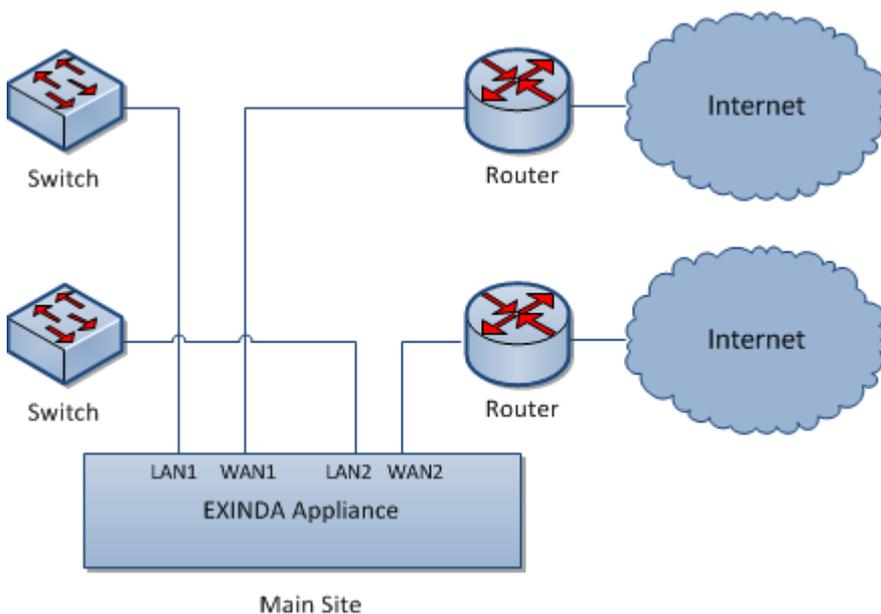
### Multiple link topology: Exinda Appliance in a network with multiple WAN connections

Some Exinda Appliances support multiple bridges, allowing multiple connection links through the appliance. This supports a single Exinda Appliance topology which allows monitoring and controlling Internet traffic to and from the main site as well as WAN traffic.

#### NOTE

We recommend using a bypass capable bridge. In the event of hardware failure, the Ethernet bypass allows your network to function even if the equipment doesn't. But keep in mind, for this strategy to be effective, once it's installed the bridge must be enabled on the [IP Address configuration page](#).

All Exinda Appliances support this topology, but some appliances only have a single bypass enabled bridge to provide Ethernet bypass in the event of hardware failure.



Screenshot 20: Multiple Link topology deployment

## Installing and Exinda Appliance in a multiple link topology

Plug your Exinda Appliance inline between the switch and router or firewall.

1. Connect the Exinda WAN1 port into your WAN router/firewall using a crossover cable.
2. Connect the Exinda LAN1 port into the LAN switch.
3. Connect the Exinda WAN2 port into your Internet router/firewall using a crossover cable.
4. Connect the Exinda LAN2 port into the LAN switch.

There are a few Exinda Appliance basics to keep in mind while planning a deployment. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

### 0.6.4 Out-of-path topologies

An Exinda Appliance connected to only one network device is considered out-of-path. By default, that implies network packets don't naturally travel through the Exinda Appliance to get to their destinations. For example, connecting an Exinda Appliance to a switch or hub on the LAN. In this case, the Exinda Appliance behaves like any other network client and requires specific protocols or modes to enable packet monitoring and inspection.

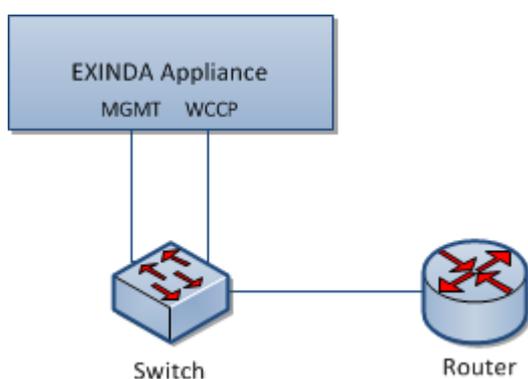
Learn how Exinda Appliances operate within an out-of-path topology.

---

### Accelerating traffic with Web Cache Communication Protocol (WCCP)

An Exinda appliance can accelerate traffic routed using Web Cache Communication Protocol (WCCP) v2. Use this topology when you want application acceleration, but do not wish to install the Exinda appliance inline.

To use WCCP v2 to route traffic to the Exinda appliance, configure the router and the Exinda appliance. To enable WCCP on the appliance, use the WCCP CLI command to assign an interface for WCCP. If you require authentication on the router, add the router's password on the appliance. Configured network objects determine the traffic direction.



Screenshot 21: Accelerate traffic with WCCP

WCCP out-of-path deployment includes a few limitations:

- » Only TCP applications can be routed to the Exinda.
- » The Router must support WCCP v2.
- » Additional load is placed on the router.

## Configuring the router to use WCCP

- » Consult the documentation of your router for instructions on configuring WCCP.
- » For Cisco routers running firmware release 12.0T, refer to [http://www.cisco.com/en/US/docs/ios/12\\_0t/12\\_0t3/feature/guide/wccp.html](http://www.cisco.com/en/US/docs/ios/12_0t/12_0t3/feature/guide/wccp.html)

## Enabling WCCP on the Exinda appliance

1. Access the CLI from the Web UI or via SSH, Telnet, or Serial Port in privileged (enable) mode and configure mode (configure terminal).
2. Assign an interface for WCCP.
  - a. For a unicast configuration, set the router IP address for each WCCP v2 service.

```
wccp interface <interface-name>
wccp service <service-group number> router <router-IP-address>
```

### EXAMPLE

Assign interface eth2 to WCCP v2 traffic with service class 10 from 192.168.0.1

```
wccp interface eth2
wccp service 10 router 192.168.0.1
```

### NOTE

Ensure you set the router to the highest IP address available on the router.

- b. For a multicast configuration, set a group-address for WCCP v2 traffic.

```
(config)# wccp interface <interface-name>
(config)# wccp service <service-group number> group-address <multicast-
address>
```

### EXAMPLE

Assign interface eth2 to WCCP v2 traffic with service class 10 from multicast address 192.168.0.1

```
wccp interface eth2
wccp service 10 group-address 224.1.1.1
```

3. If a password has been configured for a service on the router, add that password on the Exinda.

```
(config) # wccp service <service-group number> password <password>
```

## How network objects are used to determine traffic direction

For an Exinda appliance to determine traffic direction, you must define all internal subnets as internal network objects. Network objects can be edited from **Configuration > Objects > Network > Network Objects**.

After identifying the subnets as internal network objects, as traffic passes through the appliance, the appliance looks at the source IP and destination IP of the packet and matches them to the network objects to determine whether the source IP and destination IP should be considered internal or external, determining the packet direction.

Consider the following rules when comparing the location of an IP packet relative to an internal network object:

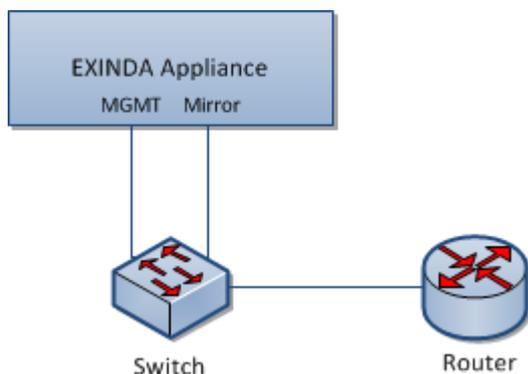
| Source IP | Destination IP | Result                                                                                                                                                                                                                                    |
|-----------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Internal  | External       | Packet is classified as outbound traffic.                                                                                                                                                                                                 |
| External  | Internal       | Packet is classified as inbound traffic.                                                                                                                                                                                                  |
| Internal  | Internal       | <ul style="list-style-type: none"> <li>» Traffic flowing from the lower IP to the higher IP is classified as outbound traffic.</li> <li>» Traffic flowing from the higher IP to the lower IP is classified as inbound traffic.</li> </ul> |
| External  | External       | <ul style="list-style-type: none"> <li>» Traffic flowing from the lower IP to the higher IP is classified as outbound traffic.</li> <li>» Traffic flowing from the higher IP to the lower IP is classified as inbound traffic.</li> </ul> |

To display the status of the WCCP service and verify that the connection between the appliance and the router is active

Use CLI to type `show wccp service <service-group number>`. The status of the service is displayed with the Router and Appliance IP addresses. If any error messages are displayed beside an IP address, resolve the issue with the configuration and re-verify the service.

### Overview of SPAN and mirror port monitoring

An Exinda Appliance can operate out-of-path or ON-LAN mode with any hub or switch that supports port mirroring or SPAN ports. Use this topology when you need to monitor traffic without installing the Exinda Appliance inline. The Exinda Appliance monitors and reports on all applications present on the SPAN/mirror port. It enables regular network audits and provides great flexibility in restricted and complex network environments.



Screenshot 22: Topology diagram showing how to cable MGMT and Mirror ports for Mirror/SPAN port monitoring.

To configure Mirror/SPAN port monitoring, perform the following tasks:

1. [Configuring Mirror Port Mode.](#)
2. [Enabling Mirror/SPAN traffic monitoring.](#)
3. [Configuring internal subnets as internal network objects.](#)

After enabling Mirror/SPAN monitoring and defining the appropriate Internal Network Objects, the Exinda appliance monitors traffic received on the Mirror/SPAN receiving port as if it were inline.

#### NOTE

Exinda interface reports will be blank because the Exinda Appliance has no insight into packet direction at the interface level.

## Configuring Mirror Port Mode

Before enabling Mirror/SPAN port monitoring, you must configure a switch port to mirror traffic to an unused port cabled to the Exinda Appliance. Alternatively, you can deploy a network hub in-path and directly cable the Exinda Appliance to the hub. A hub, by design, mirrors all traffic to all ports.

Any port not enslaved to a bridge or in use for another function, for example, cluster or WCCP, may be used to receive mirror port or SPAN port traffic.

## Enabling SPAN/Mirror traffic monitoring

Enable the SPAN/mirror port on an interface to monitor that type of traffic.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Configuration > System > Network > IP Address**.
5. To use an interface as a Mirror port, select the **Mirror** check box.



The screenshot shows the configuration page for interface 'eth2'. The 'Role' section has three checkboxes: 'Cluster' (unchecked), 'Mirror' (checked and highlighted with a red box), and 'WCCP' (unchecked). The 'Autoconf' section has two checkboxes: 'IPv4: DHCP' (unchecked) and 'IPv6: SLAAC' (unchecked). The 'Dynamic Addresses' field contains the value 'fe80::222:19ff:fed4:8dc5/64'. The 'Static Addresses' field is empty. The 'Comment' field is also empty.

6. Click **Apply Changes**. The selected interface now accepts Mirror/SPAN traffic.

The following commands can be executed from the CLI in order to enable or disable Mirror/SPAN port monitoring on an interface:

```
> en
# con t
(config) # mirror interface <inf>
(config) # no mirror interface <inf>
```

Once SPAN/Mirror monitoring is enabled and the appropriate Internal Network Objects have been defined, the Exinda Appliance will monitor traffic received on the Mirror/SPAN receiving port as if it were inline.

### NOTE

Exinda interface reports will be blank because the Exinda Appliance has no insight into packet direction at the interface level.

## Configuring internal subnets as internal network objects

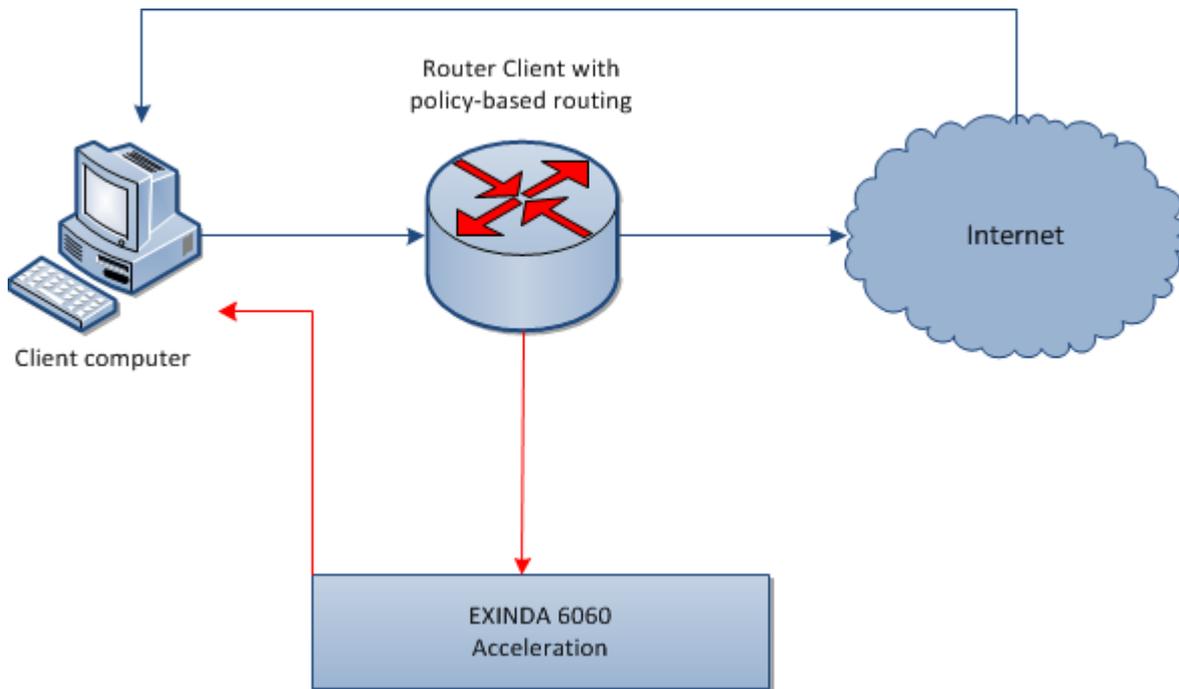
For the Exinda Appliance to determine traffic direction, all internal subnets should be defined as internal network objects. After identifying the subnets as internal network objects, as traffic passes through the appliance, the Exinda Appliance determines packet direction based on the following rules:

| Rule                                                                                                                                     | Result                                                                                                                                                                        |
|------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Packet's source IP matches an Internal Network Object<br>AND<br>Packet's destination IP DOES NOT match an Internal Network Object        | Packet is classified as outbound traffic.                                                                                                                                     |
| Packet's source IP DOES NOT match an Internal Network Object<br>AND<br>Packet's destination IP matches an Internal Network Object        | Packet is classified as inbound traffic.                                                                                                                                      |
| Packet's source IP matches an Internal Network Object<br>AND<br>Packet's destination IP matches an Internal Network Object               | Traffic flowing from the lower IP to the higher IP is classified as outbound traffic.<br>Traffic flowing from the higher IP to the lower IP is classified as inbound traffic. |
| Packet's source IP DOES NOT match an Internal Network Object<br>AND<br>Packet's destination IP DOES NOT match an Internal Network Object | Traffic flowing from the lower IP to the higher IP is classified as outbound traffic.<br>Traffic flowing from the higher IP to the lower IP is classified as inbound traffic. |

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Objects > Network**.
5. To change a network object to an internal object, click **Edit** and change the location to **Internal**.
6. Click **Apply Changes**.

### Directing traffic with policy-based routing

Using Policy Based Routing (PBR), you can deploy your Exinda appliance in the network out-of-path but retain in-path optimization capabilities. To achieve this, configure the router with the policy that determines whether traffic is sent to the requested destination or to the Exinda Appliance.

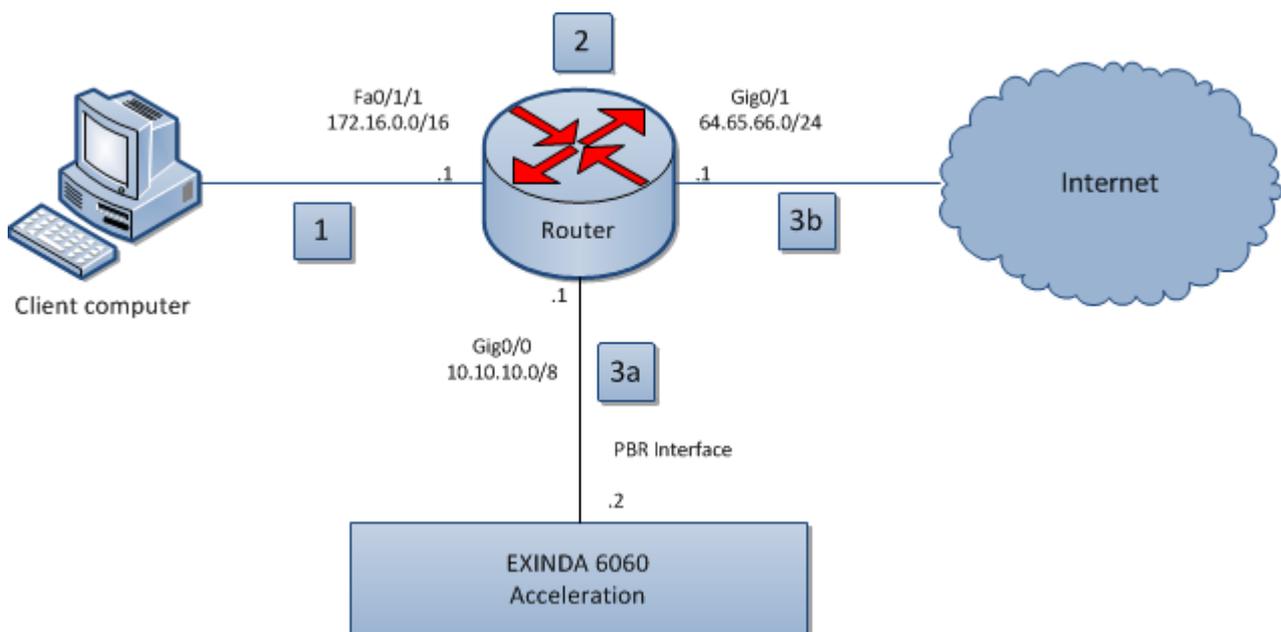


The Policy Based Routing feature offers the following benefits:

- » An Exinda Appliance can be deployed physically out-of-path, but logically in-path which makes it useful in virtual environments where an in-path deployment may undesirable or impossible.
- » Increased selectivity of traffic to be optimized. For example, redirect all web traffic to an Exinda Appliance.
- » Increased network topology configurability. Depending on the router's capability, the traffic is redirected based on access-lists which may use the port, source/destination address, etc.

### Overview of basic policy-based routing

A basic policy-based routing out-of-path deployment is depicted in the network diagram below. Deploying your Exinda Appliance in this topology requires three interfaces on the router.



Screenshot 23: Basic policy-based routing

## NOTE

All traffic between network components in this diagram is bi-directional.

- » Request to access a location on the Internet is made on a client computer.
- » The request is sent to the router where the source and destination are analyzed and compared to the policy configured on the router.
- » Based on the results of the analysis, the request is:
  - sent to the Exinda appliance for optimization, then back through the router to the requested destination.
  - sent directly to the requested destination.

For more information about basic policy-based routing, see these topics:

- » [Directing traffic with policy-based routing](#)
- » [Configuring the router for basic PBR topology](#)
- » [Configuring the out-of-path topology for policy-based routing](#)

### Configuring a router for a basic policy-based routing (PBR) topology

While deploying your Exinda Appliance in a basic PBR topology, you must specify how the router should handle traffic coming from the out-of-path Exinda Appliance, the WAN, and the LAN.

## CAUTION

Configure your router using IP addresses that correspond to your network. The IP addresses and netmasks used in this guide are for demonstration purposes only.

1. Launch the router command line interface.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

2. Configure the interface parameters for the Exinda Appliance installed out-of-path (Gig0/0).

- a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/0`
- b. Set the IP address of the out-of-path: Exinda Appliance. `hostname (config-if) # ip address 10.10.10.1 255.0.0.0`
- c. Set the duplex and speed parameters for the interface.

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

3. Configure the parameters for the WAN interface (Gig0/1).

- a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/1`

- b. Set the IP address and netmask of the WAN interface: `hostname (config-if)# ip address 64.65.66.1 255.255.255.0`
- c. Set the route map for policy routing to asymmetrical: `hostname (config-if)# ip policy route-map Asym`
- d. Set the duplex and speed parameters for the interface:

```
hostname (config-if)# duplex auto
hostname (config-if)# speed auto
```

#### 4. Configure the parameters for the LAN interface (Fa0/1/1).

- a. Specify the interface to configure: `hostname (config)# interface GigabitEthernet0/0`
- b. Set the IP address and netmask of the LAN interface: `hostname (config-if)# ip address 172.16.12.1 255.255.0.0`
- c. Set the route map for policy routing with the name `Asym`: `hostname (config-if)# ip policy route-map Asym`
- d. Set the duplex and speed parameters for the interface.

```
hostname (config-if)# duplex auto
hostname (config-if)# speed auto
```

#### 5. Create an access list named 120 that allows devices in the specified IP address range to access the network.

```
hostname (config)# access-list 120 permit ip 172.16.0.0 0.0.255.255 64.65.66.0
0.0.0.255
hostname (config)# access-list 120 permit ip 64.65.66.0 0.0.0.255 172.16.0.0
0.0.255.255
```

#### 6. Configure the route map to allow access to the routes specified in the access list (120), and route the traffic to the router.

```
route-map Asym permit 10
match ip address 120
set ip next-hop 10.10.10.2
```

### Configuring an out-of-path Exinda Appliance for policy based routing

To use policy-based routing, configure the interfaces on the Exinda Appliance with the appropriate settings, using either the Exinda Web UI or the CLI.

#### Using Exinda Web UI

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.

1. Click **Configuration > System > Network > IP Address**.
2. In the Interface Settings area, clear the BR10 checkbox. The bridge expands to display eth10 and eth11.

#### NOTE

If a virtual appliance is hosting the Exinda Appliance software, uncheck the BR2 option. The bridge expands to display eth2 and eth3.

3. Click **Configuration > System > Network > IP Address**.
4. In the eth11 area, select **PBR**.
5. In the **Static Addresses** field, type the IP address and netmask of the out-of-path Exinda Appliance.
6. In the **PBR Next-Hop Address** field, type the IP address of the router.
7. Click **Apply Changes**.
8. To save the changes to the configuration file, in the status bar click **Save**.
9. Click **Configuration > System > Tools > Console** to launch the Command Line Console.
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.
5. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
```

## Using CLI

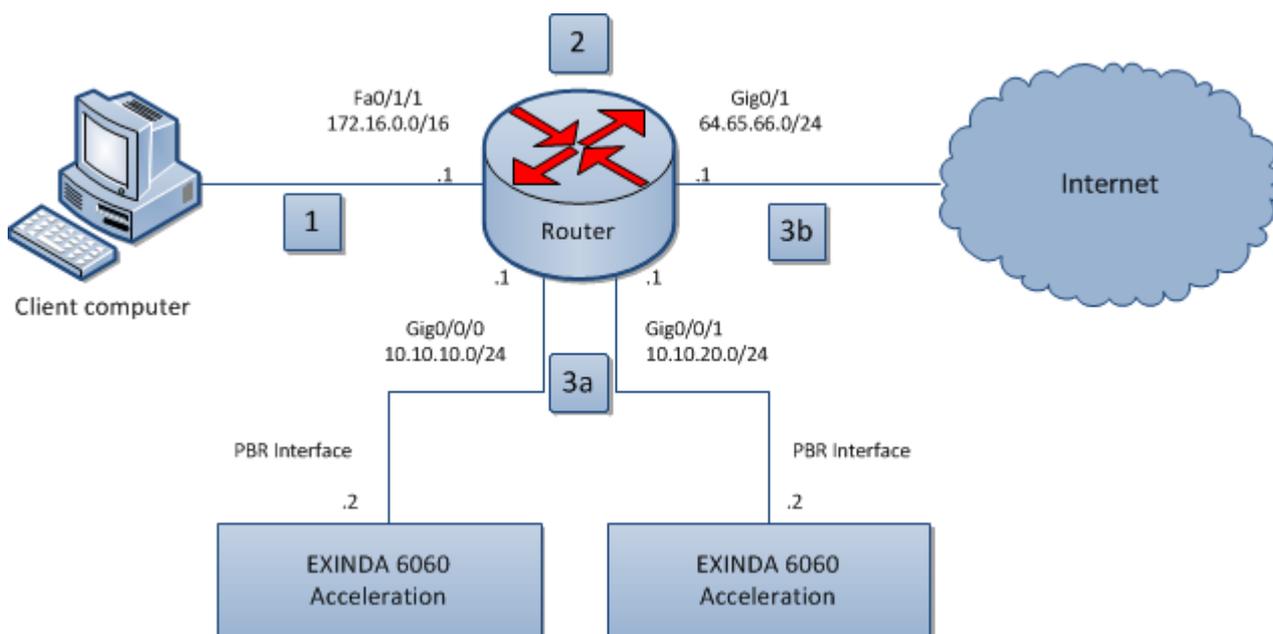
1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Configuration > System > Tools > Console**.
5. Type the appliance username and password at the prompts. Do one of the following:
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

3. Remove the BR10 bridge: `hostname (config)# no bridge BR10 enable`
4. Set the IP address and netmask of the out-of-path Exinda Appliance: `hostname (config)# interface eth11 ip address 10.10.10.2 /8`
5. Identify the interface to be used for policy-based routing: `hostname (config)# pbr interface eth11`
6. Set the IP address of the router: `hostname (config)# pbr interface ip next-hop 10.10.10.1`
7. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
```

### Overview of policy-based routing (PBR) with two subnets for failover and redundancy

You can deploy policy-based routing (PBR) on two subnets serviced by their respective Exinda Appliances. The configuration detailed below reroutes traffic from one Exinda Appliance to another in case of hardware failure.



Screenshot 24: Policy-based routing with two subnets

#### Note

All traffic between network components in this diagram is bidirectional.

- » Request to access a location on the Internet is made on a client computer.
- » The request is sent to the router, where the source destination is analyzed and compared to the policy configured on the router.
- » Based on the results of the analysis, the request is either:

- sent to the Exinda appliance for optimization, and then back through the router to the requested destination.
- or, sent directly to the requested destination.

For more information about policy-based routing with two subnets, see these topics:

- » [Configuring a router for PBR with two subnets](#)
- » [Configuring a n Exinda Appliance for PBR on two subnets](#)

### Configuring a router connected to two subnets and two Exinda Appliances for policy-based routing

To use policy-based routing (PBR), specify the behavior of the router interfaces connecting to the LAN, the WAN, and the out-of-path Exinda Appliances.

#### CAUTION

Use IP addresses, subnets and network masks that work in your environment. The IP addresses in these instructions are for demonstration purposes.

1. Launch the router command line interface.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

2. Configure the parameters for the LAN interface (Fa0/1/1).

a. Specify the interface to configure.

```
hostname (config)# interface GigabitEthernet0/0
```

b. Set the IP address and netmask of the LAN interface.

```
hostname (config-if)# ip address 172.16.12.1 255.255.0.0
```

c. Set a description for what the router interface is connecting to.

```
hostname (config-if)# description Connected to EX-IN
```

d. Set the route map for policy routing with the name `DivtEx1theEx2`.

```
hostname (config-if)# ip policy route-map DivtEx1theEx2
```

e. Set the duplex and speed parameters for the interface.

```
hostname (config-if)# duplex auto
hostname (config-if)# speed auto
```

3. Configure the interface parameters for the out-of-path Exinda Appliance #1 (Gig0/0/0).

- a. Specify the interface to configure.

```
hostname (config)# interface GigabitEthernet0/0/0
```

- b. Set a description for what the router interface is connecting to.

```
hostname (config-if)# description Connected to EX-OOP-1
```

- c. Set the IP address of the out-of-path Exinda Appliance.

```
hostname (config-if)# ip address 10.10.10.1 255.255.255.0
```

- d. Set the duplex and speed parameters for the interface.

```
hostname (config-if)# duplex auto  
hostname (config-if)# speed auto
```

#### 4. Configure the interface parameters for the out-of-path Exinda Appliance #2 (Gig0/0/1).

- a. Specify the interface to configure.

```
hostname (config-if)# ip policy route-map DivtExltheEx2
```

```
hostname (config)# interface GigabitEthernet0/0/1
```

- b. Set a description for what the router interface is connecting to.

```
hostname (config-if)# description Connected to EX-OOP-2
```

- c. Set the IP address of the out-of-path Exinda Appliance.

```
hostname (config-if)# ip address 10.10.20.1 255.255.255.0
```

- d. Set the duplex and speed parameters for the interface.

```
hostname (config-if)# duplex auto  
hostname (config-if)# speed auto
```

#### 5. Configure the parameters for the WAN interface (Gig0/1).

- a. Specify the interface to configure.

```
hostname (config)# interface GigabitEthernet0/1
```

- b. Set a description for what the router interface is connecting to.

```
hostname (config-if)# description Connected to WAN
```

- c. Set the IP address and netmask of the WAN interface

```
hostname (config-if)# ip address 64.65.66.1 255.255.255.0
```

d. Set the route map for policy routing to asymmetrical.

```
hostname (config-if)# ip policy route-map DivtEXOOP1thenEXOOP2
```

e. Set the duplex and speed parameters for the interface.

```
hostname (config-if)# duplex auto  
hostname (config-if)# speed auto
```

6. Create an access list named 120 that allows devices in the specified IP address range to access the network.

```
hostname (config)# access-list 120 permit ip 172.16.0.0 0.0.0.255 64.65.66.0  
0.255.255.255  
hostname (config)# access-list 120 permit ip 64.65.66.0 0.255.255.255  
172.16.0.0 0.0.255.255
```

7. Configure the route map to allow access to the routes specified in the access list (140), and route the traffic to the router.

```
route-map DivtEXOOP1thenEXOOP2 permit 10  
match ip address 140  
set ip next-hop 10.10.10.2 10.10.20.2
```

### Configuring an Exinda Appliance for policy-based routing

To use policy-based routing (PBR), configure the interfaces on the Exinda Appliance with the appropriate settings, using either the Exinda Web UI or the CLI.

#### Using Exinda Web UI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.
  - a. In the browser address field, type `https://Exinda_IP_address`.
  - b. Enter the appliance **User Name** and **Password**. Click **Login**. The Exinda Web UI is displayed.
  - c. Ensure you are in **Advanced** mode.
2. Click **Configuration > System > Network > IP Address**.
3. In the Interface Settings area, clear the BR10 checkbox. The bridge expands to display eth10 and eth11.

#### NOTE

If you're running an Exinda virtual appliance, uncheck the BR2 option. The bridge expands to display eth2 and eth3.

4. Click **Configuration > System > Network > IP Address**.
5. In the eth11 area, select **PBR**.
6. In the **Static Addresses** field, type the IP address and netmask of the out-of-path Exinda Appliance.

7. In the **PBR Next-Hop Address** field, type the IP address of the router.
8. Click **Apply Changes**.
9. To save the changes to the configuration file, in the status bar click **Save**.
10. Click **Configuration > System > Tools > Console** to launch the Command Line Console.
11. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
```

12. Repeat these steps on out-of-path Exinda Appliance #2.

## Using CLI

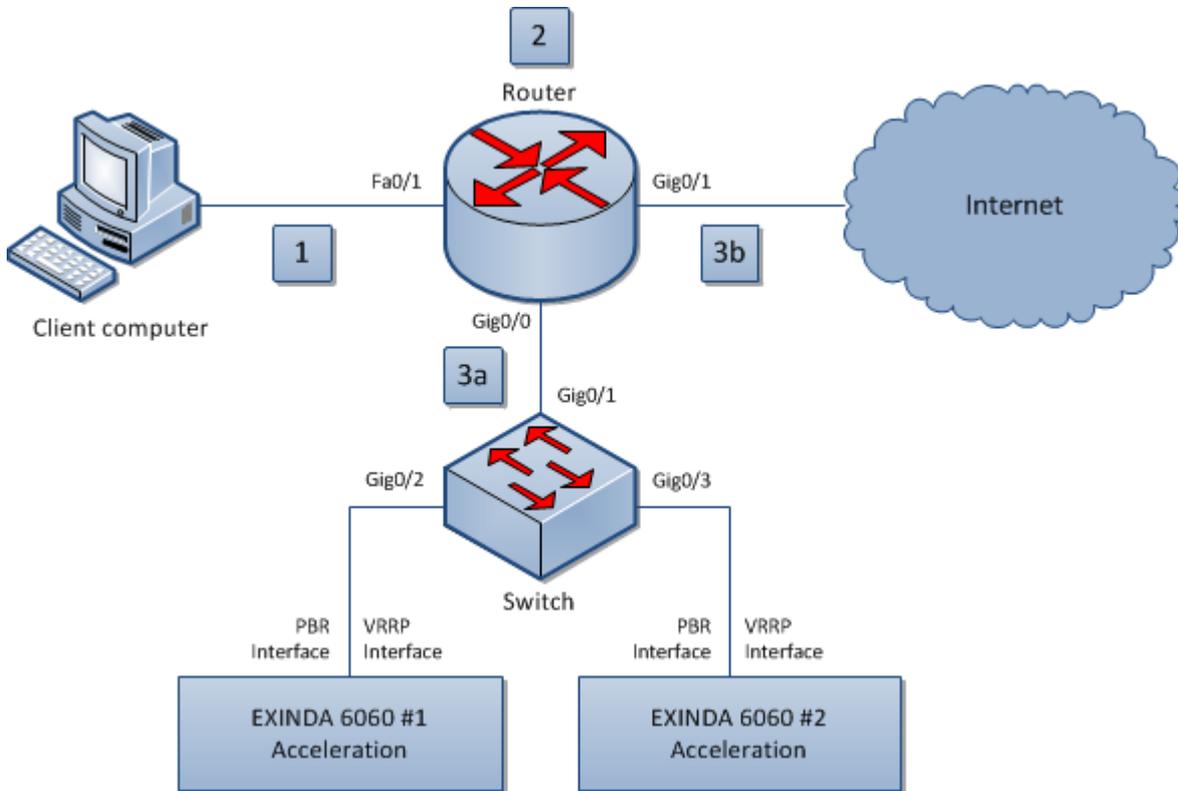
1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.
  - a. In the browser address field, type `https://Exinda_IP_address`.
  - b. Enter the appliance **User Name** and **Password**. Click **Login**. The Exinda Web UI is displayed.
  - c. Ensure you are in **Advanced** mode.
2. Click **Configuration > System > Tools > Console**.
3. Type the appliance username and password at the prompts.
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.
4. Remove the BR10 bridge: `hostname (config) # no bridge BR10 enable`
5. Set the IP address and netmask of the out-of-path Exinda Appliance: `hostname (config) # interface eth11 ip address 10.10.10.2 /24`
6. Identify the interface to be used for policy-based routing: `hostname (config) # pbr interface eth11`
7. Set the IP address of the router: `hostname (config) # pbr interface ip next-hop 10.10.20.1`
8. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
```

9. Repeat these steps on the out-of-path Exinda Appliance #2 using the appropriate IP address in Step 5.

### Overview of virtual router redundancy protocol (VRRP) with policy-based routing (PBR)

The VRRP provides automatic assignment of IP routers, using virtual routers as proxies for physical routers. An IP address acts as a gateway between the router and the Exinda Appliances. The VRRP receives traffic requests and distributes them to the appliances connected to it, providing greater reliability and steady traffic requests distribution.



Screenshot 25: Configuring VRRP with PBR

#### NOTE

All traffic between network components in this image is bi-directional.

- » The client computer requests access to a location on the Internet.
- » The request is sent to the router, where the source, destination is analyzed and compared to the policy configured on the router.
- » Based on the results of the analysis, the request is:
  - sent to the Exinda appliance for optimization, and then back through the router to the requested destination.
  - sent directly to the requested destination.

#### Configuring a router for virtual router redundancy protocol (VRRP) with policy-based routing (PBR)

To use VRRP with PBR, specify the behaviour of the router interfaces that connect to the LAN, WAN, and the out-of-path Exinda Appliances.

1. Launch the router command line interface.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.

- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`

The `hostname (config) #` prompt appears.

3. Configure the interface parameters for the switch installed between the router and the out-of-path Exinda Appliances (Gig0/0).

a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/0`

b. Set the IP address of the out-of-path Exinda Appliance: `hostname (config-if) # ip address 10.10.10.1 255.0.0.0`

c. Set the duplex and speed parameters for the interface:

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

3. Configure the parameters for the WAN interface (Gig0/1).

a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/1`

b. Set the IP address and netmask of the WAN interface: `hostname (config-if) # ip address 64.65.66.1 255.255.255.0`

c. Set the route map for policy routing to asymmetrical: `hostname (config-if) # ip policy route-map Asym`

d. Set the duplex and speed parameters for the interface:

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

4. Configure the parameters for the LAN interface (Fa0/1).

a. Specify the interface to configure: `hostname (config) # interface FastEthernet0/1`

b. Set the IP address and netmask of the LAN interface: `hostname (config-if) # ip address 172.16.12.1 255.255.0.0`

c. Set the route map for policy routing with the name `Asym`: `hostname (config-if) # ip policy route-map Asym`

d. Set the duplex and speed parameters for the interface:

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

5. Create an access list named 120 that allows devices in the specified IP address range to access the network:

```
hostname (config) # access-list 120 permit ip 172.16.0.0 0.0.255.255
64.65.66.0 0.0.0.255
hostname (config) # access-list 120 permit ip 64.65.66.0 0.0.0.255 172.16.0.0
0.0.255.255
```

6. Configure the route map to allow access to the routes specified in the access list (120), and route the traffic to the router:

```
route-map Asym permit 10
match ip address 120
set ip next-hop 10.10.10.100
```

### Configuring an Exinda Appliance for virtual router redundancy protocol (VRRP) and policy-based routing (PBR)

To use VRRP with PBR, the interfaces on the Exinda Appliance must be configured with the appropriate settings, using the Exinda Web UI or using CLI.

#### Using Exinda Web UI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.
  - a. In the browser address field, type `https://Exinda_IP_address`.
  - b. Enter the appliance **User Name** and **Password**. Click **Login**. The Exinda Web UI is displayed.
  - c. Ensure you are in **Advanced** mode.
2. Click **Configuration > System > Network > IP Address**.
3. In the Interface Settings area, clear the BR10 checkbox. The bridge expands to display eth10 and eth11.

#### NOTE

If a virtual appliance is hosting the Exinda Appliance software, uncheck BR2 option. The bridge expands to display eth2 and eth3.

4. Click **Configuration > System > Network > IP Address**.
5. In the eth11 area, select **PBR**.
6. In the **Static Addresses** field, type the IP address and netmask of the out-of-path Exinda Appliance.
7. In the **PBR Next-Hop Address** field, type the IP address of the router.
8. Select **Enable VRRP**.
9. In the **Virtual Router Address** field, type the group IP address.
10. Click **Apply Changes**.
11. To save the changes to the configuration file, in the status bar click **Save**.
12. Click **Configuration > System > Tools > Console** to launch the Command Line Console.
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.
3. To set the parameters of eth11, type the following commands:

```

hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
hostname (config)# vrrp interface eth11 delay 1
hostname (config)# vrrp interface eth11 priority 100
hostname (config)# vrrp interface eth11 vrid 50

```

4. Repeat these steps on eth10 on out-of-path Exinda Appliance #2.

## Using CLI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.

- a. In the browser address field, type `https://Exinda_IP_address`.
- b. Enter the appliance **User Name** and **Password**. Click **Login**. The Exinda Web UI is displayed.
- c. Ensure you are in **Advanced** mode.

2. Click **Configuration > System > Tools > Console**.

3. Type the appliance username and password at the prompts.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

3. Repeat these steps on the out-of-path Exinda Appliance #2 using the appropriate IP address in Step 5.

4. To set the VRRP parameters, type the following commands:

```

hostname (config)# vrrp interface eth11 delay 1
hostname (config)# vrrp interface eth11 enable
hostname (config)# vrrp interface eth11 priority 100
hostname (config)# vrrp interface eth11 vip 10.10.10.100
hostname (config)# vrrp interface eth11 vrid 50

```

5. To set the parameters of eth11, type the following commands:

```

hostname (config)# no interface eth11 shutdown
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# interface eth11 speed auto

```

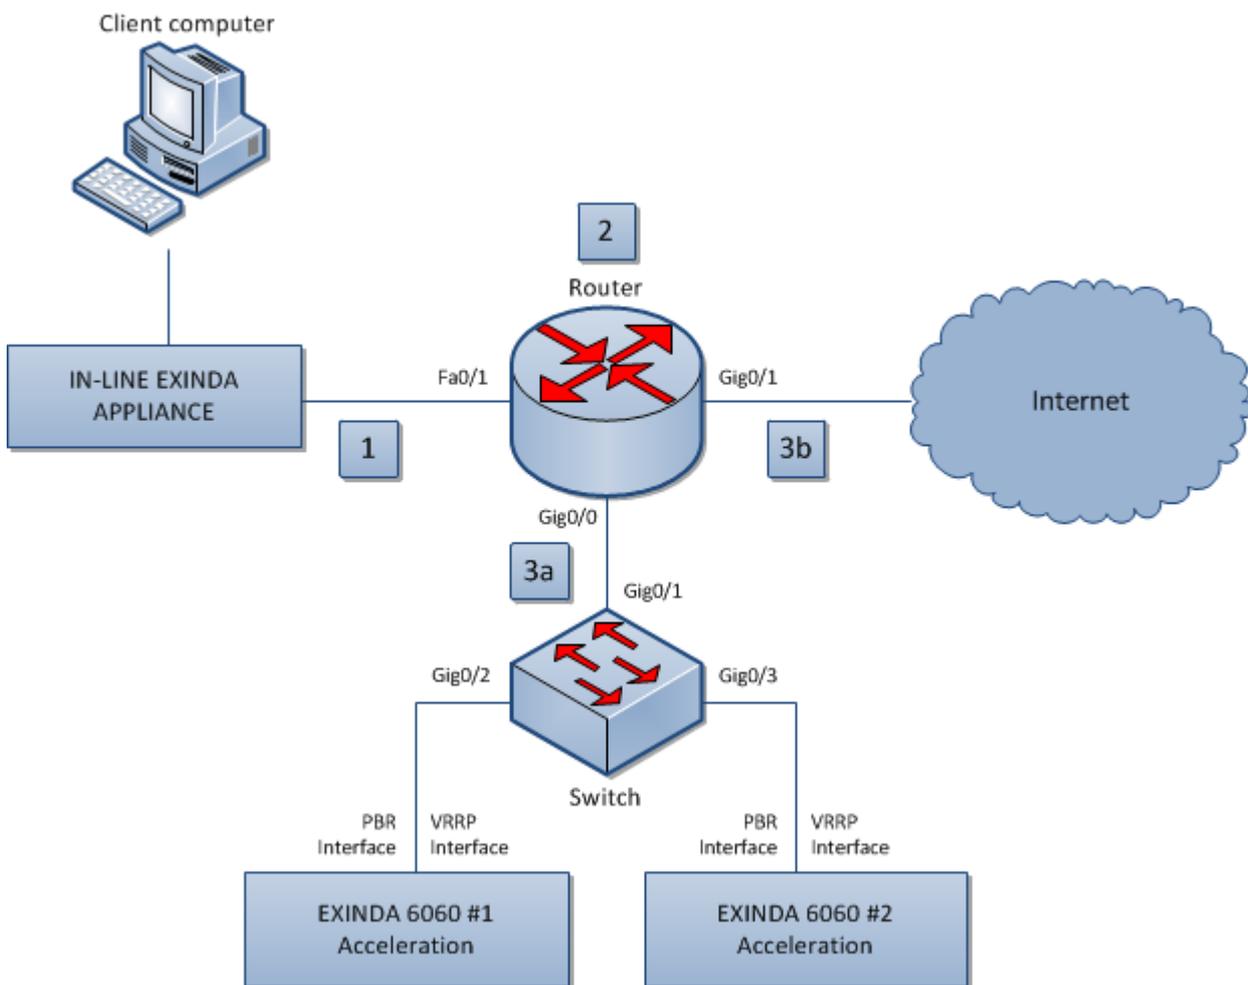
6. Set the IP address of the router: `hostname (config)# pbr interface ip next-hop 10.10.10.1`

7. Identify the interface to be used for policy-based routing: `hostname (config) # pbr interface eth11`
8. Set the IP address and netmask of the out-of-path Exinda Appliance #1: `hostname (config) # interface eth11 ip address 10.10.10.2 /8`
9. Remove the BR10 bridge: `hostname (config) # no bridge BR10 enable`

### Overview of virtual router redundancy protocol (VRRP) with policy-based routing (PBR) and VLANs

The VRRP provides automatic assignment of IP routers, using virtual routers as proxies for physical routers. A virtual router with a static IP address acts as a gateway between the router and the Exinda Appliances. The VRRP receives traffic requests and distributes them to the appliances connected to it, providing greater reliability and steady traffic requests distribution.

A switch can be configured to act as a VLAN to group Exinda Appliances connected to different physical switches. This out-of-path deployment decreases latency and increases performance for networks with VLANs.



Screenshot 26: VRRP with PBR and VLANs

#### NOTE

All traffic between the network components in this image is bi-directional.

- » The client computer requests access to a location on the Internet.
- » The request is sent to the router, where the source, destination is analyzed and compared to the policy configured on the router.
- » Based on the results of the analysis, the request is:
  - sent to the switch, which evenly distributes the traffic between the connected Exinda appliances, and then back through the router to the requested destination.
  - sent directly to the requested destination.

For more information about VRRP with PBR and VLANs, see these topics:

- » [Configuring a router for VRRP with PBR and VLANs](#)
- » [Configuring a switch for VRRP with PBR and VLANs](#)
- » [Configuring the for VRRP with PBR and VLANs](#)

### Configuring a router for virtual router redundancy protocol (VRRP) with policy-based routing (PBR) and VLANs

To use VRRP with PBR and VLANs, specify the behaviour of the router interfaces connecting to the switch, the LAN, and the WAN.

1. Launch the router command line interface.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

3. Configure the interface parameters for the switch installed between the router and the out-of-path Exinda Appliances (Gig0/0).

- a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/0`
- b. Set a description for what the router interface is connecting to: `hostname (config-if) # description Connected to Exinda Group`
- c. Specify that there is no IP address for the interface: `hostname (config-if) # no ip address`
- d. Set the duplex and speed parameters for the interface:

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

e. Configure the parameters for the VLAN 10 interface (Gig0/0.10).

- i. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/0.10`
- ii. Set the router interface to route between VLANs for the switch: `hostname (config-if) # encapsulation dot1Q 10`
- iii. Set the IP address and netmask of the VLAN 10 interface: `hostname (config-if) # ip address 10.10.10.1 255.255.255.0`

f. Configure the parameters for the VLAN 20 interface (Gig0/0.20).

- i. Specify the interface to configure: `hostname (config)# interface GigabitEthernet0/0.20`
- ii. Set the router interface to route between VLANs for the switch: `hostname (config-if)# encapsulation dot1Q 20`
- iii. Set the IP address and netmask of the VLAN 20 interface: `hostname (config-if)# ip address 10.10.20.1 255.255.255.0`

3. Configure the parameters for the WAN interface (Gig0/1).

- a. Specify the interface to configure: `hostname (config)# interface GigabitEthernet0/1`
- b. Set a description for what the router interface is connecting to: `hostname (config-if)# description Connected to WAN`
- c. Specify that there is no IP address for the interface: `hostname (config-if)# no ip address`
- d. Set the duplex and speed parameters for the interface:

```
hostname (config-if)# duplex auto
hostname (config-if)# speed auto
```

e. Configure the parameters for the VLAN 10 interface (Gig0/1.10).

- i. Specify the interface to configure: `hostname (config)# interface GigabitEthernet0/1.10`
- ii. Set the router interface to route between VLANs for the switch: `hostname (config-if)# encapsulation dot1Q 10`
- iii. Set the route map for policy routing with the name EXOOP1toEXIN: `hostname (config-if)# ip policy route-map EXOOP1toEXIN`

f. Configure the parameters for the VLAN 20 interface (Gig0/1.20).

- i. Specify the interface to configure: `hostname (config)# interface GigabitEthernet0/1.20`
- ii. Set the router interface to route between VLANs for the switch: `hostname (config-if)# encapsulation dot1Q 20`
- iii. Set the IP address and netmask of the LAN interface: `hostname (config-if)# ip address 10.10.40.1 255.255.255.0`
- iv. Set the route map for policy routing with the name LANToEXOOP2: `hostname (config-if)# ip policy route-map LANToEXOOP2`

4. Configure the parameters for the LAN interface (Fa0/1.10).

- a. Specify the interface to configure: `hostname (config)# interface FastEthernet0/1.10`
- b. Set a description for what the router interface is connecting to: `hostname (config-if)# description Connected to EX-IN`
- c. Set the router interface to route between VLANs for the switch: `hostname (config-if)# encapsulation dot1Q 10`

d. Set the IP address and netmask of the LAN interface: `hostname (config-if)# ip address 172.16.10.0 255.255.255.0`

e. Set the route map for policy routing with the name EXINToEXOOP1: `hostname (config-if)# ip policy route-map EXINToEXOOP1`

5. Configure the parameters for the LAN interface (Fa0/1.20).

a. Specify the interface to configure: `hostname (config)# interface FastEthernet0/1.20`

b. Set a description for what the router interface is connecting to: `hostname (config-if)# description Connected to EX-IN`

c. Set the router interface to route between VLANs for the switch: `hostname (config-if)# encapsulation dot1Q 20`

d. Set the IP address and netmask of the LAN interface: `hostname (config-if)# ip address 172.16.20.0 255.255.255.0`

e. Set the route map for policy routing with the name EXINToEXOOP2: `hostname (config-if)# ip policy route-map EXINToEXOOP2`

6. Create the following access lists to allow devices in the specified IP address range to access the network:

```
hostname (config)# access-list 100 permit ip 10.10.30.0 0.0.0.255 172.16.10.0 0.0.0.255
hostname (config)# access-list 101 permit ip 172.16.10.0 0.0.0.255 10.10.30.0 0.0.0.255
hostname (config)# access-list 102 permit ip 10.10.40.0 0.0.0.255 172.16.20.0 0.0.0.255
hostname (config)# access-list 103 permit ip 172.16.20.0 0.0.0.255 10.10.40.0 0.0.0.255
```

7. Configure the following route maps to allow access to the routes specified in the access lists, and route the traffic to the router.

```
route-map Asym permit 10
match ip address 100
set ip next-hop 10.10.10.100
!
route-map EXOOP1toEXIN permit 10
match ip address 100
set ip next-hop 10.10.10.100
!
route-map EXOOP2toEXIN permit 10
match ip address 102
set ip next-hop 10.10.20.100
!
route-map EXINToEXOOP1 permit 10
match ip address 101
set ip next-hop 10.10.10.100
!
route-map EXINToEXOOP2 permit 10
match ip address 103
set ip next-hop 10.10.20.100
```

## Configuring a switch for virtual router redundancy protocol (VRRP) with policy-based routing (PBR) and VLANs

To use VRRP with PBR and VLANs, specify the behaviour of the switch interfaces connecting to the router and the out-of-path Exinda Appliances.

1. Launch the command line interface for the switch.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`

The `hostname #` prompt appears.

- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`

The `hostname (config) #` prompt appears.

3. Configure the interface to the router.

```
interface GigabitEthernet0/1
switchport trunk encapsulation dot1q
switchport mode trunk
```

4. Configure the interface to Exinda Appliance #1.

```
interface GigabitEthernet0/2
switchport trunk encapsulation dot1q
switchport mode trunk
```

5. Configure the interface to Exinda Appliance #2.

```
interface GigabitEthernet0/3
switchport trunk encapsulation dot1q
switchport mode trunk
```

6. Set the IP addresses for the VLANs.

```
interface Vlan10
  no ip address
interface Vlan20
  no ip address
```

## Configuring an Exinda Appliance for virtual router redundancy protocol (VRRP) and policy-based routing (PBR) and VLANs

To use policy-based routing, configure the interfaces on the Exinda Appliance , using the the Exinda Web UI or CLI.

### Using Exinda Web UI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.

- a. In the browser address field, type `https://Exinda_IP_address`.The Exinda Web UI is displayed.
- b. Enter the appliance **User Name** and **Password**. Click **Login**.
- c. Ensure you are in **Advanced** mode.

2. Click **Configuration > System > Network > IP Address**.

3. In the Interface Settings area, clear the BR10 checkbox. The bridge expands to display eth10 and eth11.

#### NOTE

Clear the BR2 checkbox if you're using an Exinda virtual appliance. The bridge expands to display eth2 and eth3.

4. Click **Configuration > System > Network > IP Address**.

5. In the VLAN Settings area, select the **eth11** interface, and type the id **10**.

6. Click **Add VLAN**. The eth11.10 interface is created.

7. In the **eth11.10** area, select **PBR**.

8. In the **Static Addresses** field, type the IP address and netmask of the out-of-path Exinda Appliance.

9. In the **PBR Next-Hop Address** field, type the IP address of the router.

10. Select **Enable VRRP**.

11. In the **Virtual Router Address** field, type the group IP address.

12. Click **Apply Changes**.

13. To save the changes to the configuration file, in the status bar click **Save**.

14. Click **Configuration > System > Tools > Console** to launch the Command Line Console.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.

- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

3. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11.10 dhcp
hostname (config)# interface eth11.10 display
hostname (config)# interface eth11.10 duplex auto
hostname (config)# interface eth11.10 mtu 1500
hostname (config)# no interface eth11.10 shutdown
hostname (config)# interface eth11.10 speed auto
hostname (config)# vrrp interface eth11.10 delay 1
hostname (config)# vrrp interface eth11.10 priority 100
hostname (config)# vrrp interface eth11.10 vrid 50
```

4. Repeat these steps on eth10 of the out-of-path Exinda Appliance #2.

## Using CLI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.

a. In the browser address field, type `https://Exinda_IP_address`.

b. Enter the appliance **User Name** and **Password**. Click **Login**. The Exinda Web UI is displayed.

c. Ensure you are in **Advanced** mode.

3. Click **Configuration > System > Tools > Console**.

4. Type the appliance username and password at the prompts.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`

The `hostname #` prompt appears.

- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`

The `hostname (config) #` prompt appears.

3. Remove the BR10 bridge: `hostname (config) # no bridge BR10 enable`

4. Identify the VLAN interface: `hostname (config) # vlan vlan-id 10 interface eth11`

5. Set the IP address and netmask of `eth11.10` on the out-of-path Exinda Appliance #1: `hostname (config) # interface eth11.10 ip address 10.10.10.2 /8`

6. Identify the interface to be used for policy-based routing: `hostname (config) # pbr interface eth11.10`

7. Set the IP address of the router: `hostname (config) # pbr interface eth11.10 ip next-hop 10.10.10.1`

8. To set the parameters of `eth11.10`, type the following commands:

```
hostname (config) # no interface eth11.10 dhcp
hostname (config) # interface eth11.10 display
hostname (config) # interface eth11.10 duplex auto
hostname (config) # interface eth11.10 mtu 1500
hostname (config) # no interface eth11.10 shutdown
hostname (config) # interface eth11.10 speed auto
```

9. To set the VRRP parameters, type the following commands:

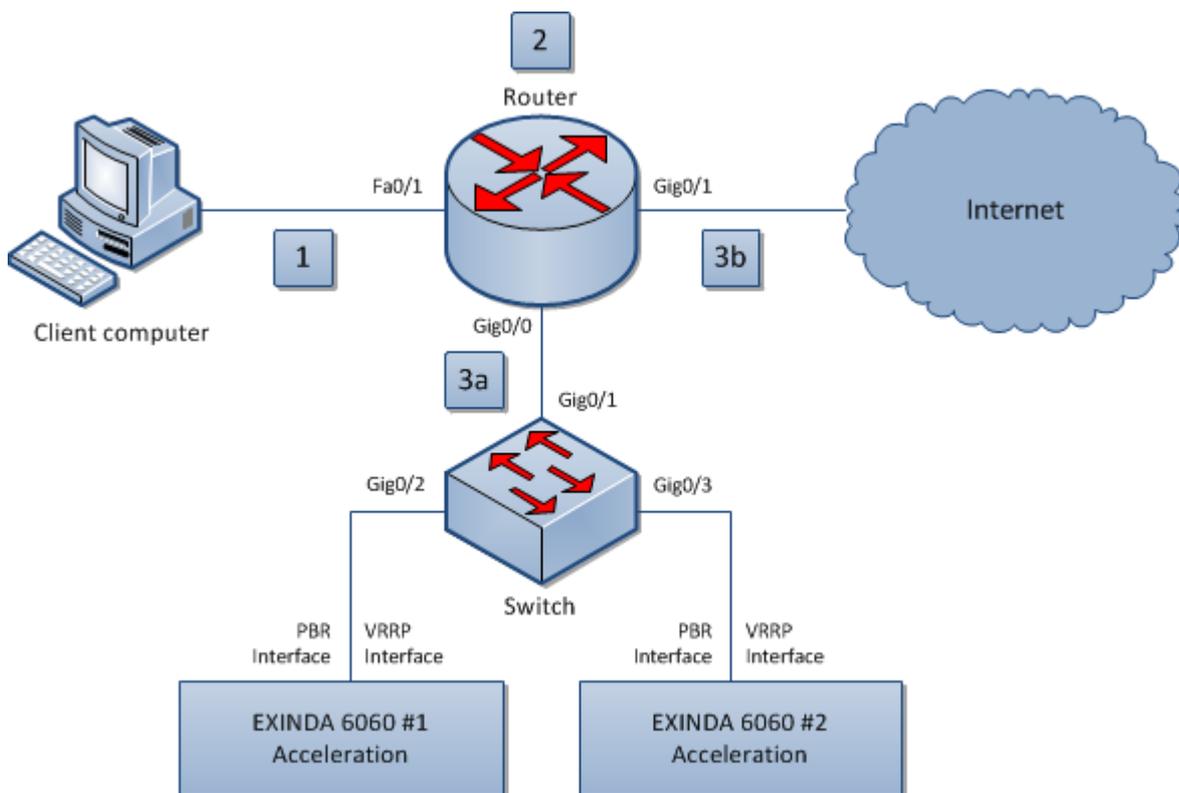
```
hostname (config) # vrrp interface eth11.10 delay 1
hostname (config) # vrrp interface eth11.10 enable
hostname (config) # vrrp interface eth11.10 priority 100
hostname (config) # vrrp interface eth11.10 vip 10.10.10.100
hostname (config) # vrrp interface eth
```

10. Repeat these steps on the out-of-path Exinda Appliance #2 using the appropriate IP address in Step 5.

## Overview of virtual router redundancy protocol (VRRP) with policy-based routing (PBR) and IP service levels of applications (SLA) tracking

The VRRP provides automatic assignment of IP routers, using virtual routers as proxies for physical routers. A virtual router with a static IP address acts as a gateway between the router and the Exinda Appliances. The VRRP receives traffic requests and distributes them to the appliances connected to it, providing greater reliability and steady traffic requests distribution.

To increase fault-tolerance, configure the router to monitor SLA tracking. IP SLA configurations act as a heartbeat mechanism between the router and applicable SLA hosts.



Screenshot 27: VRRP with PBR and IP SLA Tracking

#### NOTE

All traffic between network components in this image is bi-directional.

- » The client computer requests access to a location on the Internet.
- » The request is sent to the router, where the source, destination, and service levels for the requested application is analyzed and compared to the policy configured on the router.
- » Based on the results of the analysis, the request is: either
  - sent to the switch, which evenly distributes the traffic between the connected Exinda appliances, and then back through the router to the requested destination.
  - or, sent directly to the requested destination.

#### Limitations

This configuration is only supported on Cisco routers and only on the following releases:

Cisco IOS Release 12.3(11)T, 12.2(25)S, or Prior Releases

Configuration instructions can be found here: [https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_pi/configuration/12-4t/iri-12-4t-book/iri-pbr-mult-track.pdf](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/12-4t/iri-12-4t-book/iri-pbr-mult-track.pdf)

Cisco IOS Release 12.3(14)T, 12.2(33)SXH, and Later Releases

Configuration instructions can be found here: [https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_pi/configuration/12-2sx/iri-12-2sy-book/iri-pbr-mult-track.pdf](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/12-2sx/iri-12-2sy-book/iri-pbr-mult-track.pdf)

For more information about VRRP with PBR and VLANs with IP SLA tracking, see the topics listed below:

- » [Configuring a router for VRRP with PBR and IP SLA Tracking](#)
- » [Configuring an Exinda Appliance for VRRP with PBR and IP SLA Tracking](#)

### Configuring a router for virtual router redundancy protocol (VRRP) with policy-based routing (PBR) and IP service levels of applications (SLA) Tracking

To use VRRP with PBR and IP SLA Tracking, specify how the router should handle traffic coming from the out-of-path Exinda Appliance, the WAN, and the LAN.

1. Launch the router command line interface.

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

3. Configure the interface parameters for the Exinda Appliance installed out-of-path (Gig0/0).

- a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/0`
- b. Set the IP address of the out-of-path Exinda Appliance: `hostname (config-if) # ip address 10.10.10.1 255.0.0.0`
- c. Set the duplex and speed parameters for the interface:

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

3. Configure the parameters for the WAN interface (Gig0/1).

- a. Specify the interface to configure: `hostname (config) # interface GigabitEthernet0/1`
- b. Set the IP address and netmask of the WAN interface: `hostname (config-if) # ip address 64.65.66.1 255.255.255.0`
- c. Set the route map for policy routing to asymmetrical: `hostname (config-if) # ip policy route-map Asym`
- d. Set the duplex and speed parameters for the interface:

```
hostname (config-if) # duplex auto
hostname (config-if) # speed auto
```

4. Configure the parameters for the LAN interface (Fa0/1).

- a. Specify the interface to configure: `hostname (config) # interface FastEthernet0/1`
- b. Set the IP address and netmask of the LAN interface: `hostname (config-if) # ip address 172.16.12.1 255.255.0.0`
- c. Set the route map for policy routing with the name `Asym`: `hostname (config-if) # ip policy route-map Asym`
- d. Set the duplex and speed parameters for the interface:

```
hostname (config-if)# duplex auto
hostname (config-if)# speed auto
```

5. Create an access list named 120 that allows devices in the specified IP address range to access the network.

```
hostname (config)# access-list 120 permit ip 172.16.0.0 0.0.255.255
64.65.66.0 0.0.0.255
hostname (config)# access-list 120 permit ip 64.65.66.0 0.0.0.255 172.16.0.0
0.0.255.255
```

6. Configure the route map to allow access to the routes specified in the access list (120), and route the traffic to the router.

```
route-map Asym permit 10
match ip address 120
set ip next-hop verify-availability 10.10.10.100 1 track 1
```

7. Configure the IP address of the VRRP group and the schedule of the IP SLA operation.

```
hostname (config)# ip sla 5
hostname (config-ip-sla)#icmp-echo 10.10.10.100
hostname (config)#ip sla schedule 5 life forever start-time now
```

8. Set the tracking of the availability of the IP SLA operation.

```
track 1 ip sla 5
delay down 2 up 2
```

### Configure the Exinda Appliance for virtual router redundancy protocol (VRRP) with policy-based routing (PBR) and IP service levels for applications (SLA) tracking

To use VRRP with PBR and IP SLA Tracking, configure the interfaces on the Exinda Appliance, using the Exinda Web UI or CLI.

#### Using Exinda Web UI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.
  - a. In the browser address field, type `https://Exinda_IP_address`.
  - b. Enter the appliance **User Name** and **Password**. Click **Login**. The Exinda Web UI is displayed.
  - c. Ensure you are in **Advanced** mode.
2. Click **Configuration > System > Network > IP Address**.
3. In the Interface Settings area, clear the BR10 checkbox. The bridge expands to display eth10 and eth11.

#### NOTE

If a virtual appliance is hosting the Exinda Appliance software, clear the BR2 checkbox. The bridge expands to display eth2 and eth3.

4. Click **Configuration > System > Network > IP Address**.
  5. In the eth11 area, select **PBR**.
  6. In the **Static Addresses** field, type the IP address and netmask of the out-of-path Exinda Appliance.
  7. In the **PBR Next-Hop Address** field, type the IP address of the router.
  8. Select **Enable VRRP**.
  9. In the **Virtual Router Address** field, type the group IP address.
  10. Click **Apply Changes**.
  11. To save the changes to the configuration file, in the status bar click **Save**.
  12. Click **Configuration > System > Tools > Console**, launch the Command Line Console.
    - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
    - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.
3. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
hostname (config)# vrrp interface eth11 delay 1
hostname (config)# vrrp interface eth11 priority 100
hostname (config)# vrrp interface eth11 vrid 50
```

4. Repeat these steps on eth10 of the out-of-path Exinda Appliance #2.

## Using CLI

1. On the out-of-path Exinda Appliance # 1, launch the Exinda Web UI.
2. In the browser address field, type `https://Exinda_IP_address`. Enter the appliance **User Name** and **Password**. Click **Login** to show the Exinda Web UI.
3. Ensure you are in **Advanced** mode.
4. Click **Configuration > System > Tools > Console**.
5. Type the appliance username and password at the prompts.
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.
6. Remove the BR10 bridge: `hostname (config) # no bridge BR10 enable`

7. Set the IP address and netmask of the out-of-path Exinda Appliance #1: `hostname (config)# interface eth11 ip address 10.10.10.2 /8`
8. Identify the interface to be used for policy-based routing: `hostname (config)# pbr interface eth11`
9. Set the IP address of the router: `hostname (config)# pbr interface ip next-hop 10.10.10.1`
10. To set the parameters of eth11, type the following commands:

```
hostname (config)# no interface eth11 dhcp
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto
```

11. To set the VRRP parameters, type the following commands:

```
hostname (config)# vrrp interface eth11 delay 1
hostname (config)# vrrp interface eth11 enable
hostname (config)# vrrp interface eth11 priority 100
hostname (config)# vrrp interface eth11 vip 10.10.10.100
hostname (config)# vrrp interface eth11 vrid 50
```

12. Repeat these steps on the out-of-path Exinda Appliance #2 using the appropriate appliance information.

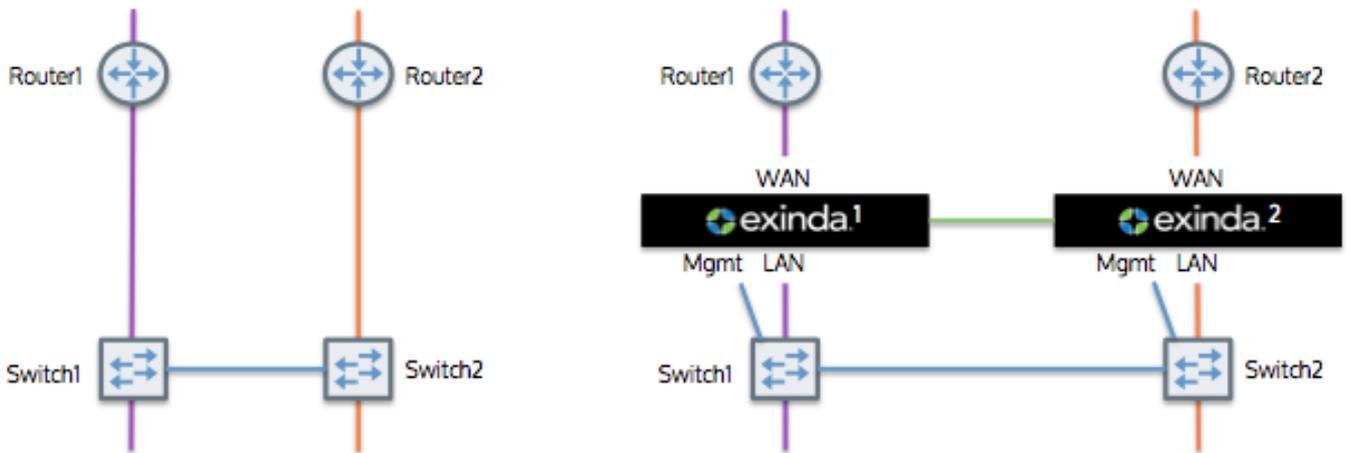
## 0.6.5 Clustering topologies

The three most common network environments that require Exinda clustering are:

- » The environment is critical and failover is necessary. Typically, one network link is active and the other is passive. The Exinda Appliance on the passive link clusters with the active Exinda Appliance so it has all the required data and configuration settings in case the active link fails.
- » The environment has dual paths with active links into one or more ISPs. In this case, the Exinda Appliances need to be clustered so each appliance exchanges all traffic with the others.
- » The environment is load balanced with traffic distributed across several servers. In this case, an Exinda Appliance is installed in front of each load balanced server and all other Exinda Appliances cluster so they each have the same configuration and network traffic can be reported on in aggregate.

### Configuring redundancy using multiple Exinda Appliances

In a redundant topology, multiple Exinda Appliances are deployed in a redundant architecture. The Exinda appliances are configured as a high availability cluster, act as a single appliance and maintain alternate paths for data to travel in case a cable breaks, a component loses power or piece of hardware fails.



Screenshot 28: Redundant topology

Each appliance processes traffic between its switch and router. The two appliances are directly connected and each forwards its traffic to the other. Forwarded traffic received from the other appliance does not get forwarded onto the LAN. Both appliances report on the traffic in the same way.

In this configuration, Exinda Appliances automatically synchronize their configuration settings.

### Capabilities of redundant topology

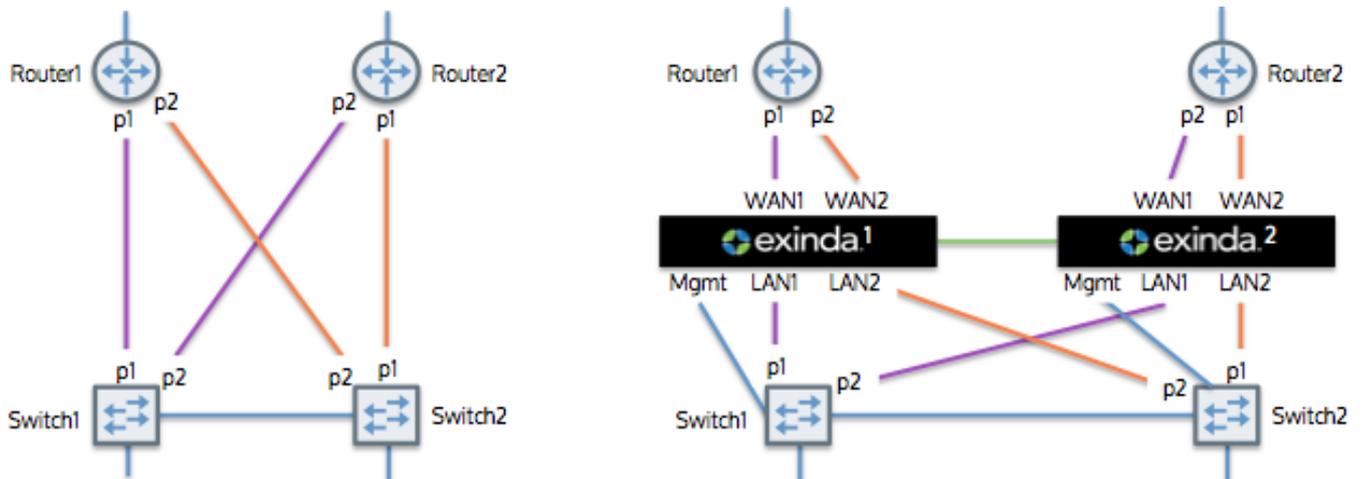
- » Monitoring of both links.
- » Optimization of both links.
- » Transparent failover.

### Installing a redundant topology

1. On each Exinda Appliance, assign an interface for cluster internal use and, an interface to manage the appliance.
2. Connect the cluster interfaces on each Exinda with a crossover cable.
3. Power up Exinda 1. After one minute, power up Exinda 2.
4. Connect Exinda 1 LAN into switch 1/port 1.
5. Connect Exinda 1 WAN into router 1.
6. Connect Exinda 2 LAN into switch 2/port 1.
7. Connect Exinda 2 WAN into router 2.
8. Connect Exinda 1 management interface into switch 1/port 2.
9. Connect Exinda 2 management interface into switch 2/port 2.

### Configuring load balancing and failover with multiple Exinda Appliances

In a load balancing topology, Exinda Appliances are linked and handle failover and load balancing for each other.



Screenshot 29: Load balancing and fail-over with multiple Exinda Appliances

In this topology, two Exinda Appliances are connected to two routers. The two appliances are also directly connected to each other and each forwards its traffic to the other appliance.

Note that the traffic received from the other Exinda appliance is not forwarded onto the LAN. Both appliances report on the data in the same way.

Exinda Appliances automatically synchronize configuration settings between the two appliances.

The Exinda appliances that support this topology include the 4060\*, 4061\*, 5000, 6010, 6060\*, 7000 and 10060\*.

#### NOTE

\* indicates that network expansion modules are required.

### Capabilities

- » Monitoring of both links.
- » Optimization of both links.
- » Redundancy of Exinda appliances.

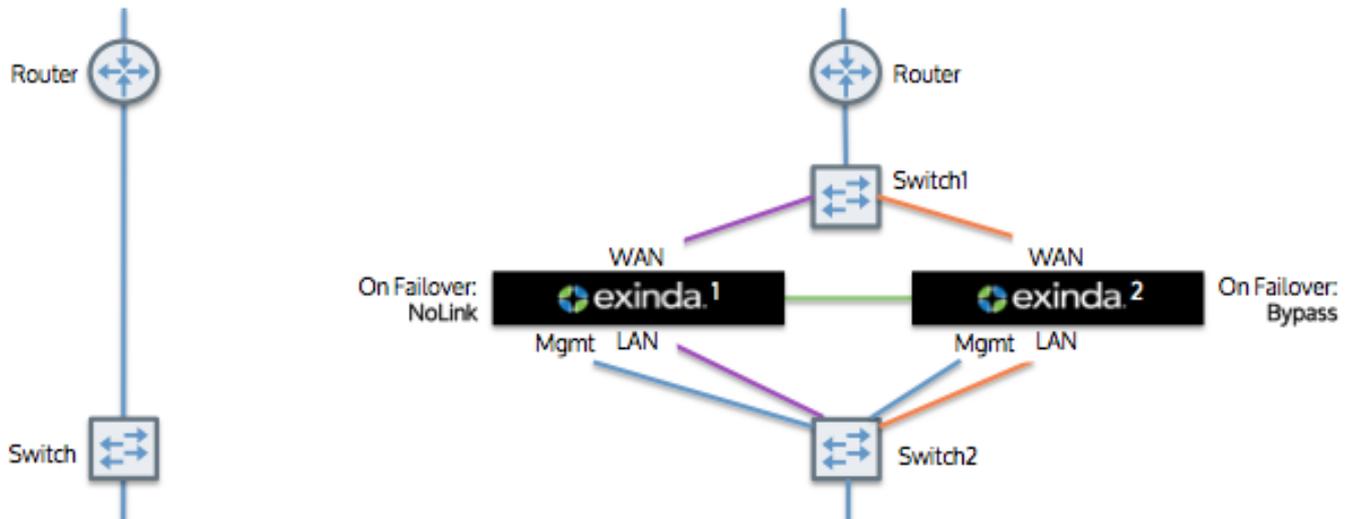
### Installation

1. On each Exinda, assign an interface for cluster internal use and, an interface for managing the appliance.
2. Connect the cluster interfaces on each Exinda with a crossover cable.
3. Power up Exinda 1. After one minute, power up Exinda 2.
4. Connect Exinda 1 LAN2 into switch 2/port 2.
5. Connect Exinda 1 WAN2 into router 1/port 2.
6. Connect Exinda 1 LAN1 into switch 1/port 1.
7. Connect Exinda 1 WAN1 into router 1/port 1.
8. Connect Exinda 2 LAN2 into switch 2/port 1.
9. Connect Exinda 2 WAN2 into router 2/port 1.
10. Connect Exinda 2 LAN1 into switch 1/port 2.
11. Connect Exinda 2 WAN1 into router 2/port 2.

12. Connect Exinda 1 MGMT into switch 1.
13. Connect Exinda 2 MGMT into switch 2.

### Configuring high availability mode

To achieve high availability or to support a greater number accelerated connections, deploy your Exinda Appliances in cluster mode. This configuration allows the your Exinda Appliances to act as a single appliance.



Screenshot 30: Topology without router redundancy

In this topology, both Exinda Appliances connect to the router via a WAN switch. The two appliances are directly connected to each other and each forwards its traffic to the other. Forwarded traffic between Exinda Appliances is not forwarded onto the LAN. Both appliances report on the data in the same way.

The Exinda Appliances automatically synchronize configuration settings between the two appliances.

#### NOTE

To configure for high availability so that there is an active path and a standby path:

Your WAN switch (Switch 1) and LAN switch (Switch 2) must support the Spanning Tree Protocol (STP). Configure the STP in Switch 2, so that port 1 (connected to Exinda 1) has higher priority than port 2 (connected to Exinda 2). If the link at switch 2/port 1 goes down (e.g., Exinda 1 loses power), then the switch will enable switch 2/port 2. Exinda1 should configure NoLink as the bridge failover option, Exinda 2 should configure Bypass.

- » Active Path: switch 1/port 1 to switch 2/port 1
- » Standby Path: switch 1/port 2 to switch 2/port 2

### Capabilities

- » Monitoring data available on both Exinda appliances.
- » Optimization available via Exinda 1 or Exinda 2.
- » Double the acceleration connections capability.
- » Redundancy of Exinda appliances.

## Installation

1. On each Exinda, assign an interface for cluster internal use and an interface for managing the appliance.
2. Connect the cluster interfaces on each Exinda with a crossover cable.
3. Power up Exinda 1. After 1 minute power up Exinda 2.
4. Connect Exinda 1 LAN into switch 2/port 1.
5. Connect Exinda 1 WAN into switch 1/port 1.
6. Connect Exinda 2 LAN into switch 2/port 2.
7. Connect Exinda 2 WAN into switch 1/port 2.
8. Connect the management interface of Exinda 1 into switch 2.
9. Connect the management interface of Exinda 2 into switch 2.
10. On Exinda 1, select "NoLink" for the LAN/WAN bridge failover mode.
11. On Exinda 2, select "Bypass" for the LAN/WAN bridge failover mode.

### NOTE

Switch2/port1 should have the highest STP priority.

## 0.7 Upgrading and downgrading

Exinda Network Orchestrator is upgraded regularly with new product features and improvements.

Learn how to upgrade and downgrade the firmware of your Exinda Appliance:

---

|                                                                      |    |
|----------------------------------------------------------------------|----|
| 0.7.1 Upgrading to the latest firmware version .....                 | 69 |
| 0.7.2 Rolling back to the previously installed version of ExOS ..... | 70 |

---

### 0.7.1 Upgrading to the latest firmware version

#### Important notes when upgrading

- » To avoid network impacts, schedule upgrades when network traffic is low.
- » Ensure you have an active internet connection before upgrading.
- » A valid Software Subscription (SS) is required to upgrade. Go to **Configuration > System > Maintenance > Firmware Update** to view your SS expiry date at the top of the page.
- » Keep a copy of your current configuration safe for disaster recovery. For more information, refer to [How to backup your Appliance settings](#) (page 704).
- » An Exinda appliance contains two partitions. One of these partitions stores and runs the current running version, which means that while retaining the running version and without stopping any of its services or operations, you can install a newer version on the other unused partition. This also makes downgrading easy.
- » If you are running an Exinda Virtual Appliance, check that the machine running the hypervisor meets the [system and hardware requirements](#) of the newer version.

- » When Exinda is upgraded from a previous version, in most cases not enough space is allocated to the new log service and the Size column displays zero. The administrator is required to manually free up space on the device and [resize](#) the log service at least to the minimum required value before volume can be encrypted or formatted.
- » If you are managing your Exinda Appliances centrally using the Exinda Management Center (EMC). You can upgrade the firmware version of all the appliances from EMC at the same time.

## Upgrade procedure

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Configuration > System > Maintenance > Firmware Update**.
5. Click **Check for Latest Update** to determine if a new firmware version is available.

Alternatively:

- If you already know the URL of the image file, select **Install from URL** and key-in the URL. The URL is usually published in the Release Notes.
  - If you had previously downloaded the image file onto the appliance, select **Install from downloaded file** and select the image from the list.
  - If the image file has been downloaded and stored on your computer, select **Install from local file** and navigate to the location of the file and select it.
6. You can also schedule the download and installation of the new update at a later date or time by checking **Schedule Installation** and specifying **Date** and **Time**. By default, the image gets downloaded straight away and only the installation gets scheduled. To schedule the download of the image to happen at the scheduled time, select **Schedule Image Download**.
  7. To restart the appliance after the scheduled installation, check **Reboot After Installation**. By default, the Exinda appliance will not reboot following a scheduled installation.
  8. Before installing or scheduling a new update, you must accept the **End User License Agreement (EULA)**.
  9. Click **Install**. This process may take a few minutes to complete.

**NOTE:** If the network connection fails while retrieving the latest file for the upgrade, you must manually restart the download. When the download restarts, any previously downloaded data is retained and only the remaining data is downloaded.

10. Reboot the appliance once the installation finishes. For more information, refer to [Reboot the Exinda Appliance](#) (page 706).

## 0.7.2 Rolling back to the previously installed version of ExOS

If you have updated the version of ExOS running on your Exinda Appliance, you can rollback to the previously installed version. Exinda Appliances have two partitions for installing ExOS updates. The current (running) ExOS version is installed on one partition and the previously installed version is on the other partition. This allows you to revert to the older ExOS version.

## IMPORTANT

When rolling back to a previous ExOS version, the system configuration changes to the state it was in the last time the older version ran. If you made changes to the system configuration since upgrading from the older version, they will be lost when the Exinda appliance is rolled back.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Maintenance** and switch to the **Firmware Update** tab.
6. In the Current Installed Images area, click **Switch Boot Partition**.
7. To finalize the ExOS install, reboot the appliance. The Exinda Appliance reboots, running the previous version of ExOS. For more information, refer to [Reboot/Shutdown](#) (page 706).

## 0.8 Installing an Exinda Appliance

Learn how to install, set up and configure your Exinda Appliance.

### NOTE

This information covers installing physical appliances. For information about installing a virtual appliance, refer to [Exinda Virtual Appliances](#).

See the following topics for more information:

---

|                                                              |     |
|--------------------------------------------------------------|-----|
| 0.8.1 Gathering required information .....                   | 72  |
| 0.8.2 Connecting the appliance to the physical network ..... | 72  |
| 0.8.3 Installing an Exinda Virtual Appliance .....           | 73  |
| 0.8.4 Scenario .....                                         | 80  |
| 0.8.5 Use Case .....                                         | 80  |
| 0.8.6 Scenario .....                                         | 82  |
| 0.8.7 Use-case .....                                         | 82  |
| 0.8.8 Use Cases .....                                        | 83  |
| 0.8.9 Related Topics .....                                   | 97  |
| 0.8.10 Related Topics .....                                  | 97  |
| 0.8.11 Related Topics .....                                  | 107 |
| 0.8.12 Related Topics .....                                  | 113 |
| 0.8.13 Related Topics .....                                  | 114 |
| 0.8.14 Related Topics .....                                  | 121 |
| 0.8.15 Related Topics .....                                  | 126 |

---

---

|                                                                       |     |
|-----------------------------------------------------------------------|-----|
| 0.8.16 Related Topics .....                                           | 127 |
| 0.8.17 Related Topics .....                                           | 129 |
| 0.8.18 Related Topics .....                                           | 131 |
| 0.8.19 Related Topics .....                                           | 131 |
| 0.8.20 Related Topics .....                                           | 139 |
| 0.8.21 Related Topics .....                                           | 140 |
| 0.8.22 Creating an initial configuration using the Basic Wizard ..... | 142 |
| 0.8.23 Licensing information .....                                    | 145 |

---

## 0.8.1 Gathering required information

Use this list prior to installing your Exinda Appliance to check that you have the requisite physical hardware and information to ensure a smooth installation.

1. Inspect the package contents. The package contents vary slightly depending on the model. In general, the package includes the following items:

- 1 × Exinda Appliance
- 2 × Power Leads
- 1 × UTP Straight Cable (usually blue)
- 1 × UTP Crossover Cable (usually red)
- 1 × Serial Cable
- 1 × Quick Start Guide

2. Determine the location of the appliance in your network. For more information, refer to [Deployment options](#) (page 26).

3. Collect all the information needed for system configuration. The initial configuration wizard requires the following:

- **Host Name** – A name you give the Exinda Appliance to refer to and remember it more easily.
- **Ethernet Negotiation Settings** – Any Ethernet negotiation settings with any particular equipment that connects the Exinda appliance, such as hard-coded Ethernet speed and duplex settings.
- **IP Address and Netmask** – An available IP address and netmask.
- **Default Gateway** – The default gateway IP address.
- **DNS Server(s)** – At least one DNS server is required so the Exinda Appliance can resolve host names.
- **SMTP Server** – An SMTP server, if you wish to receive e-mail notifications from the Exinda Appliance.
- **Time Zone** – The time zone the Exinda Appliance will be in.

## 0.8.2 Connecting the appliance to the physical network

There are basic characteristics and behaviors of Exinda Appliances you should know prior to installation. For more information, refer to [Basic characteristics and behaviors of Exinda Appliances](#) (page 26).

To physically connect an Exinda Appliance to a network, do the following:

1. With the appliance powered OFF, install a crossover cable (red) between the Exinda's WAN port (eth3) and the router/firewall.
2. Install a straight-through cable (blue) between the Exinda LAN port (eth2) and the LAN switch.
3. With the Exinda Appliance off, ensure connectivity through the appliance to ensure the Ethernet bypass is working.
4. Install a straight-through cable (blue) between the Exinda management port 1 and the LAN switch.
5. Plug in the power cord.
6. Switch on the appliance and ensure that the NIC lights flash and traffic is passing through.

After installing the appliance, run the initial configuration wizard. For more information, refer to [Creating an initial configuration using the Basic Wizard](#) (page 142).

### 0.8.3 Installing an Exinda Virtual Appliance

The Exinda range of unified performance management (UPM) appliances is available as fully featured virtual appliances. There is no difference between the software that runs on bare metal hardware and the software that runs on the virtual appliance. If the Exinda software detects it is running on a hypervisor, it automatically enables certain optimizations to ensure maximum performance.

Exinda provides support for Virtual Appliances running on the following hypervisors:

- » [VMware vSphere](#) (ESX and ESXi) (5.5 and 6.0)
- » [Citrix XenServer](#) (6.2 and 6.5)
- » [Microsoft Hyper-V](#) (Windows Server 2008 R2 and later)
- » [Linux KVM](#) (Intel VT or AMD-V)

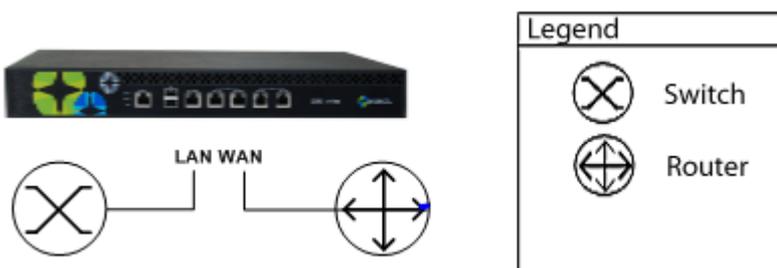
Exinda Appliances deploy the same ways as hardware appliances. For more information, refer to [Deployment options](#) (page 26).

The two typical deployment topologies are in-path and out-of-path.

#### Overview of in-path deployment

In-path deployments involve one or more LAN/WAN port pairs bridged together at layer 2. To be monitored and optimized by an Exinda Virtual Appliance, traffic must go through the bridge.

In an in-path topology, an Exinda Virtual Appliance is deployed inline. Using an image of a hardware appliance to represent a virtual appliance, an inline Exinda Virtual Appliance plugs into the network via the LAN and WAN ports like this:



Screenshot 31: Inline deployment

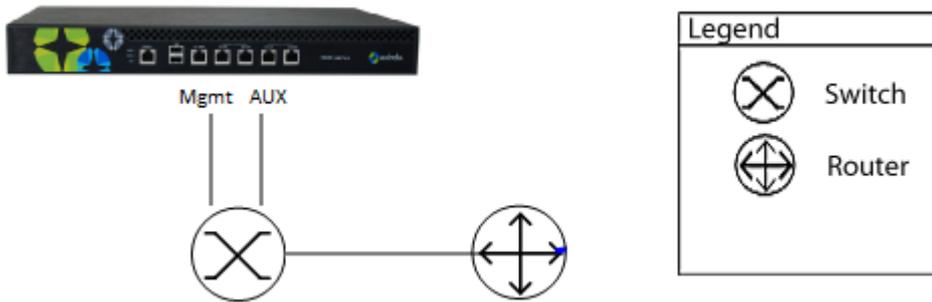
The challenge in a virtual environment is deciding how to pass traffic through the bridge. The method you choose depends on your virtual environment and hypervisor.

There are several options:

- » Dedicate two physical NICs on the host to be LAN and WAN ports.
- » Used shared NICs on the host to be LAN and WAN ports, but logically separate the traffic with VLAN tags.
- » Create a virtual network and direct traffic through virtual LAN/WAN ports.

## Overview of out-of-path deployment

Out-of-path deployments are typically used in SPAN port mirroring, WCCP, high availability (HA) and clustering scenarios. Using an image of a hardware appliance to represent a virtual appliance, an Exinda Virtual Appliance plugs into the network like this:



Screenshot 32: Out-of-path deployment

The setup is straightforward, requiring one NIC for management and another for auxiliary.

See the following topic for more information about installing and running an Exinda Virtual Appliance.

## Sizing and resource requirements

Exinda supports:

- » 50Mbps to 10Gbps for Diagnostics
- » 20Mbps to 10Gbps for Diagnostics and Shaping
- » 1Mbps to 1Gbps for Diagnostics, Shaping, and Acceleration

### NOTE

The recommended CPU is the minimum number of CPUs to achieve the performance numbers. The CPU must be VT Enabled & 64-bit.

The Exinda virtual appliance has a Flexible Storage option, with which you can adjust the size of the storage for Edge Cache, SMB1 cache and WAN Memory cache. By increasing the virtual file sizes for each of these caches, you can greatly improve the performance of your Exinda virtual appliance.

For Edge Cache and SMB1 Cache there is no limit to the size of the file created on the external storage. Use common sense when creating the file sizes. Exinda recommends that 80% of the actual file size be allocated to Edge Cache and SMB1 Cache. For example, if your SMB1 cache is 1TB then the recommendation is 800MB.

For WAN Memory, size should be based on the following:

- » For systems with 2GB RAM – Max WAN Memory Cache is 300GB
- » For all other systems – Max WAN Memory Cache is 1TB

Depending on the WAN bandwidth, use the following sizing specifications to estimate the host resources required for each of the Exinda virtual appliances.

| Virtual Appliance    | Diagnostics    | Diagnostics and Shaping | Diagnostics, Shaping, and Acceleration | Virtual CPUs (Qty × GHz) | Minimum RAM | Minimum Disk Storage |
|----------------------|----------------|-------------------------|----------------------------------------|--------------------------|-------------|----------------------|
| VM Small             | Up to 150 Mbps | Up to 50 Mbps           | Up to 10 Mbps                          | 4 × 2.0GHz               | 6GB         | 250GB                |
| VM Medium            | Up to 1 Gbps   | Up to 500 Mbps          | Up to 20 Mbps                          | 4 × 2.4GHz               | 8GB         | 250GB                |
| VM Large             | Up to 5 Gbps   | Up to 5 Gbps            | Up to 200 Mbps                         | 8 × 2.4GHz               | 32GB        | 500GB                |
| VM Extra Large       | Up to 10 Gbps  | Up to 10 Gbps           | Up to 1 Gbps                           | 12–24 × 2.4GHz           | 64GB        | 1TB                  |
| VM Extra Extra Large | Up to 15 Gbps  | Up to 15 Gbps           | Up to 1 Gbps                           | 12–24 × 2.4GHz           | 128GB       | 1TB                  |

**NOTE**

On the Exinda Virtual Appliance 10063, the number of CPUs depends on the licensed bandwidth for Diagnostics, Shaping, and Acceleration.

**NOTE**

In the case of Minimum Disk Storage, storage at higher throughput will require higher disk I/O bandwidth, so the underlying storage should be RAID-based, ideally RAID 10.

**IMPORTANT**

There are several factors that may mean more or fewer resources are required in individual environments.

Factors that may have an effect include:

- » Quality, speed, performance of the host CPUs.
- » Quality and performance of host NICs.
- » Host disk I/O bandwidth.

Use these topics for more guidance on sizing, specifications and resource requirements:

- » [VM Small](#)
- » [VM Medium](#)
- » [VM Large](#)
- » [VM Extra Large](#)
- » [VM Extra Extra Large](#)

## Exinda Virtual Appliance model VM Small specifications

| Requirement                                                         | Details                                                                                                                                                                                            |     |  |
|---------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|--|
| Diagnostics Licensed Bandwidth (in mbps)                            | 150M                                                                                                                                                                                               |     |  |
| Diagnostics and Shaping Licensed Bandwidth (in mbps)                | 50M                                                                                                                                                                                                |     |  |
| Diagnostics, Shaping, and Acceleration Licensed Bandwidth (in mbps) | 5M                                                                                                                                                                                                 | 10M |  |
| Max Concurrent Flows                                                | 150,000                                                                                                                                                                                            |     |  |
| Max L7 New Connection Rate                                          | 4,000                                                                                                                                                                                              |     |  |
| Maximum Accelerated Connections                                     | 2,000                                                                                                                                                                                              |     |  |
| Reports                                                             | 20                                                                                                                                                                                                 |     |  |
| SLAs                                                                | 100                                                                                                                                                                                                |     |  |
| APS Objects                                                         | 100                                                                                                                                                                                                |     |  |
| Policies                                                            | 512                                                                                                                                                                                                |     |  |
| Edge Cache Max Throughput (in mbps)                                 | 20                                                                                                                                                                                                 |     |  |
| Default Bridge (Interface)                                          | br1 (eth1, eth2)/bypass<br>br3 (eth3, eth4)/bypass                                                                                                                                                 |     |  |
| CPU (Qty × GHz)                                                     | 2 × 2.0GHz                                                                                                                                                                                         |     |  |
|                                                                     | <p><b>NOTE</b><br/>All virtual machine models must be run on hosts with Intel® Xeon® class CPUs. CPU ratings requirements are as quoted in the table. These CPUs must be VT Enabled and 64-Bit</p> |     |  |
| Minimum Storage (GB)                                                | 250GB                                                                                                                                                                                              |     |  |
| Minimum Memory (GB)                                                 | 6GB                                                                                                                                                                                                |     |  |
| EC-IOPS                                                             | 50                                                                                                                                                                                                 |     |  |
| Monitoring-IOPS                                                     | 150                                                                                                                                                                                                |     |  |

## Exinda Virtual Appliance model VM Medium specifications

| Specification                                                       | Details |      |      |
|---------------------------------------------------------------------|---------|------|------|
| Diagnostics Licensed Bandwidth <sup>1</sup> (in gbps)               | 1G      |      |      |
| Diagnostics and Shaping Licensed Bandwidth (in mbps)                | 100M    | 250M | 500M |
| Diagnostics, Shaping, and Acceleration Licensed Bandwidth (in mbps) | 20M     |      |      |
| Max Concurrent Flows                                                | 220,000 |      |      |

| Specification                    | Details                                                                                                                                                                                            |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max L7 New Connection Rate       | 10,000                                                                                                                                                                                             |
| Maximum Accelerated Connections  | 4,500                                                                                                                                                                                              |
| Reports                          | 60                                                                                                                                                                                                 |
| SLAs                             | 250                                                                                                                                                                                                |
| APS Objects                      | 250                                                                                                                                                                                                |
| Policies                         | 1,024                                                                                                                                                                                              |
| Edge Cache Max Throughput (Mbps) | 20                                                                                                                                                                                                 |
| Default Bridge (Interface)       | br10 (eth10, eth11)/bypass                                                                                                                                                                         |
| CPU (Qty × GHz)                  | 4 × 2.4Ghz                                                                                                                                                                                         |
|                                  | <p><b>NOTE</b><br/>All virtual machine models must be run on hosts with Intel® Xeon® class CPUs. CPU ratings requirements are as quoted in the table. These CPUs must be VT Enabled and 64-Bit</p> |
| Minimum Storage (GB)             | 250GB                                                                                                                                                                                              |
| Minimum Memory (GB)              | 8GB                                                                                                                                                                                                |
| EC-IOPS                          | 50                                                                                                                                                                                                 |
| Monitoring-IOPS                  | 150                                                                                                                                                                                                |

#### Exinda Virtual Appliance model 8064 specifications

| Specification                                                    | Details    |      |    |
|------------------------------------------------------------------|------------|------|----|
| Diagnostics Licensed Bandwidth (Gbps)                            | 5G         |      |    |
| Diagnostics and Shaping Licensed Bandwidth (Gbps)                | 3G         | 4G   | 5G |
| Diagnostics, Shaping, and Acceleration Licensed Bandwidth (Mbps) | 100M       | 200M |    |
| Max Concurrent Flows                                             | 1,200,000  |      |    |
| Max L7 New Connection Rate                                       | 25,000     |      |    |
| Maximum Accelerated Connections                                  | 20,000     |      |    |
| Reports                                                          | 100        |      |    |
| SLAs                                                             | 300        |      |    |
| APS Objects                                                      | 300        |      |    |
| Policies                                                         | 4,096      |      |    |
| Edge Cache Max Throughput (Mbps)                                 | 175        |      |    |
| CPU (Qty × GHz) <sup>2</sup>                                     | 8 × 2.4GHz |      |    |
| Minimum Storage (GB)                                             | 500GB      |      |    |
| Minimum Memory (GB)                                              | 32GB       |      |    |
| EC-IOPS                                                          | 80         |      |    |

| Specification              | Details                    |
|----------------------------|----------------------------|
| Monitoring-IOPS            | 150                        |
| Default Bridge (interface) | br10 (eth10, eth11)/bypass |

### Exinda Virtual Appliance model 10064 specifications

| Specification                                             | Details                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Diagnostics Licensed Bandwidth (Gbps)                     | 10G                                                                                                                                                                                                                                                                                                                                           |
| Diagnostics and Shaping Licensed Bandwidth                | 6G                      7G                      8G                      10G                                                                                                                                                                                                                                                                   |
| Diagnostics, Shaping, and Acceleration Licensed Bandwidth | 300                      400                      500                      1G                                                                                                                                                                                                                                                                 |
| Max Concurrent Flows                                      | 1,800,000                                                                                                                                                                                                                                                                                                                                     |
| Max L7 New Connection Rate                                | 30,000                                                                                                                                                                                                                                                                                                                                        |
| Maximum Accelerated Connections                           | 30,000                                                                                                                                                                                                                                                                                                                                        |
| Reports                                                   | 100                                                                                                                                                                                                                                                                                                                                           |
| SLAs                                                      | 300                                                                                                                                                                                                                                                                                                                                           |
| APS Objects                                               | 300                                                                                                                                                                                                                                                                                                                                           |
| Policies                                                  | 4,096                                                                                                                                                                                                                                                                                                                                         |
| Edge Cache Max Throughput (Mbps)                          | 250                                                                                                                                                                                                                                                                                                                                           |
| Default Bridge (Interface)                                | br10 (eth10 eth11)/bypass                                                                                                                                                                                                                                                                                                                     |
| CPU (Qty × GHz)                                           | 12 × 2.4 (<=300M3)<br>24 × 2.4 (>300M3)                                                                                                                                                                                                                                                                                                       |
|                                                           | <p><b>NOTE</b><br/>All virtual machine models must be run on hosts with Intel® Xeon® class CPUs. CPU ratings requirements are as quoted in the table. These CPUs must be VT Enabled and 64-Bit</p> <p><b>NOTE</b><br/>On the EXNV-10063, the number of CPUs depends on the licensed bandwidth for Diagnostics, Shaping, and Acceleration.</p> |
| Minimum Storage (TB)                                      | 2TB                                                                                                                                                                                                                                                                                                                                           |
| Minimum Memory (GB)                                       | 64GB                                                                                                                                                                                                                                                                                                                                          |
| EC-IOPS                                                   | 80                                                                                                                                                                                                                                                                                                                                            |
| Monitoring-IOPS                                           | 150                                                                                                                                                                                                                                                                                                                                           |

## Exinda Virtual Appliance model VM Extra Extra Large specifications

| Specification                                             | Details                                                                                                                                                                                            |     |     |     |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|-----|
| Diagnostics Licensed Bandwidth (Gbps)                     | 15                                                                                                                                                                                                 |     |     |     |
| Diagnostics and Shaping Licensed Bandwidth                | 6G                                                                                                                                                                                                 | 7G  | 8G  | 10G |
| Diagnostics, Shaping, and Acceleration Licensed Bandwidth | 300                                                                                                                                                                                                | 400 | 500 | 1G  |
| Max Concurrent Flows                                      | 1,800,000                                                                                                                                                                                          |     |     |     |
| Max L7 New Connection Rate                                | 30,000                                                                                                                                                                                             |     |     |     |
| Maximum Accelerated Connections                           | 30,000                                                                                                                                                                                             |     |     |     |
| Reports                                                   | 100                                                                                                                                                                                                |     |     |     |
| SLAs                                                      | 300                                                                                                                                                                                                |     |     |     |
| APS Objects                                               | 300                                                                                                                                                                                                |     |     |     |
| Policies                                                  | 4,096                                                                                                                                                                                              |     |     |     |
| Default Bridge (Interface)                                | br10 (eth10, eth11)/bypass                                                                                                                                                                         |     |     |     |
| Edge Cache Max Throughput (Mbps)                          | 250                                                                                                                                                                                                |     |     |     |
| CPU (Qty × GHz)                                           | 12 × 2.4 (<=300M3)<br>24 × 2.4 (>300M3)                                                                                                                                                            |     |     |     |
|                                                           | <p><b>NOTE</b><br/>All virtual machine models must be run on hosts with Intel® Xeon® class CPUs. CPU ratings requirements are as quoted in the table. These CPUs must be VT Enabled and 64-Bit</p> |     |     |     |
|                                                           | <p><b>NOTE</b><br/>On the EXNV-12063, the number of CPUs depends on the licensed bandwidth for Diagnostics, Shaping, and Acceleration.</p>                                                         |     |     |     |
| Minimum Storage (TB)                                      | 2TB                                                                                                                                                                                                |     |     |     |
| Minimum Memory (GB)                                       | 128GB                                                                                                                                                                                              |     |     |     |
| EC-IOPS                                                   | 80                                                                                                                                                                                                 |     |     |     |
| Monitoring-IOPS                                           | 150                                                                                                                                                                                                |     |     |     |

## Exinda Virtual Appliance use cases

Learn different ways of deploying the Exinda virtual appliance. These use cases cover both inline and out-of-path deployments. All of the hypervisors referenced in this documentation support out-of-path deployments. Inline deployments are supported on VMware vSphere and Citrix XenServer, but NOT on Microsoft Hyper-V.

## Inline deployment with externally attached LAN

In this use case, the Exinda Virtual Appliance is set up for inline mode deployment with an externally attached LAN. There are primarily two scenarios for inline deployment of the virtual appliance:

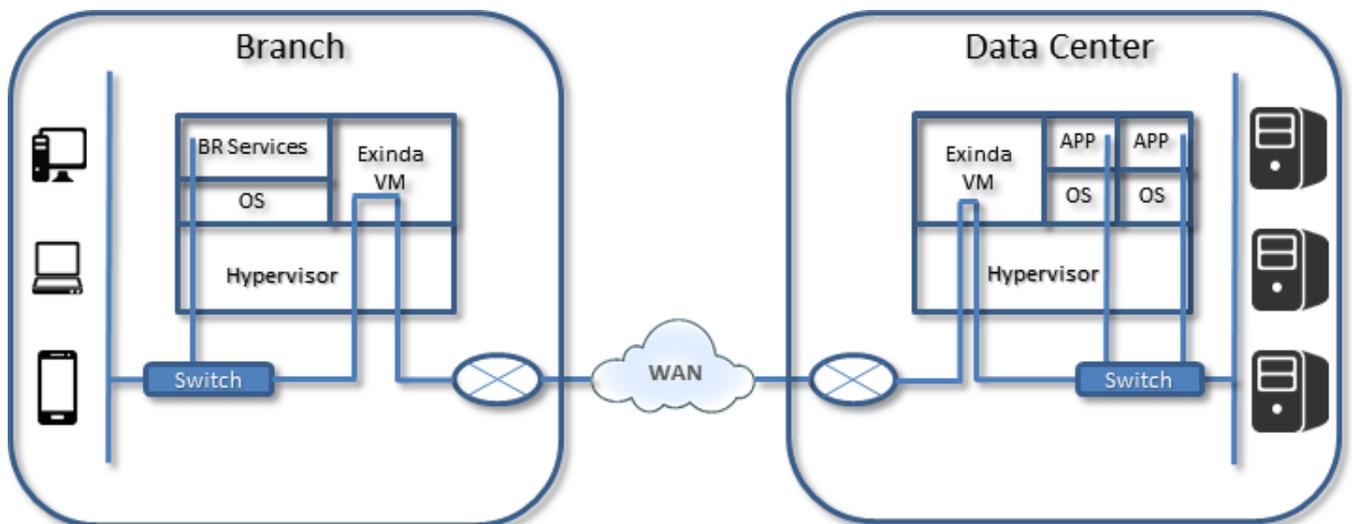
1. The LAN-side users and applications connect to the Exinda Virtual Appliance through a physical NIC interface.
2. The applications are virtualized and isolated on the same host as the Exinda, on the LAN side interface of the Exinda Virtual Appliance.

This use case discusses the first scenario.

### NOTE

This scenario is only possible using VMware vSphere (v5.5 and v6.0) and Citrix XenServer (v6.2 and v6.5) hosts. Microsoft Hyper-V does not support inline deployments.

## 0.8.4 Scenario

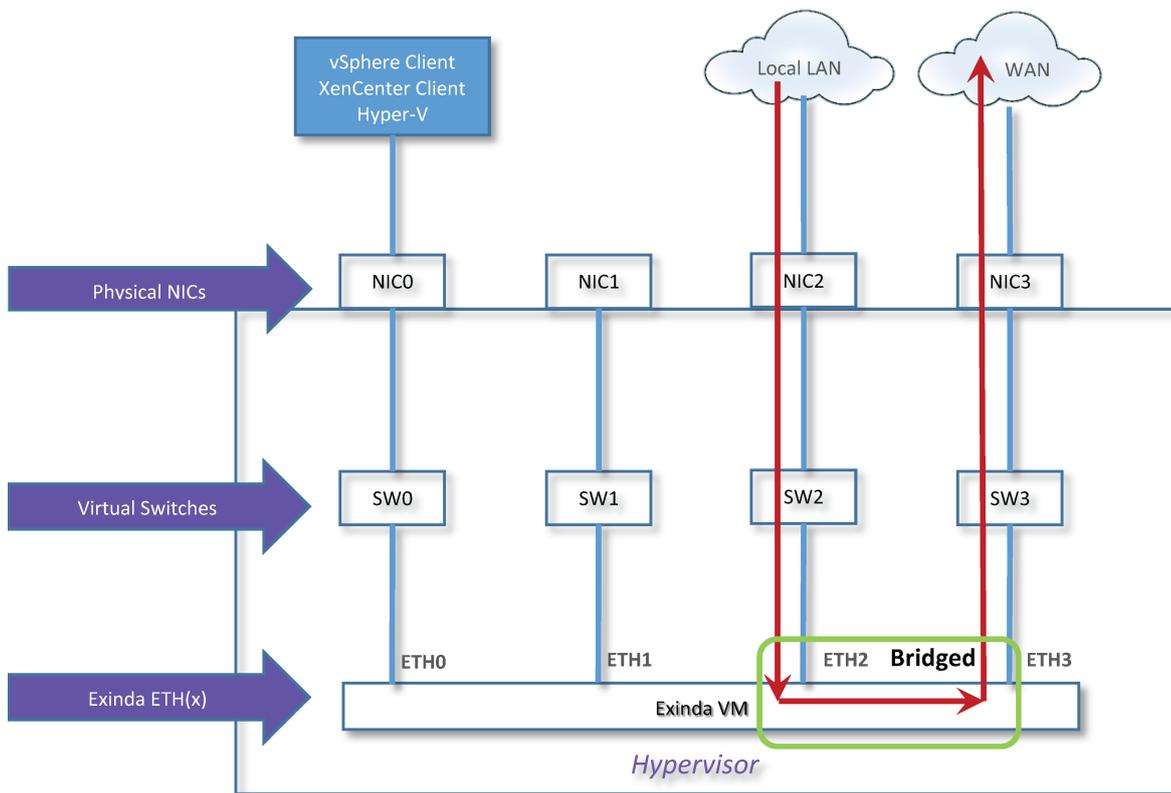


Screenshot 33: Scenario details

- » The branch users access virtualized local services (such as print, DHCP, and Active Directory).
- » At least two physical NIC interfaces are bridged together in the virtual Exinda Virtual Appliance.
- » User connections from the branch office to the Data Center applications are in line through the Exinda Virtual Appliance on both ends of the connections, and through external NIC interfaces.
- » The Exinda provides diagnostic, shaping, and acceleration for all traffic in this configuration.
- » Optionally, install the Exinda Virtual Appliance in a separate host in an inline mode configuration and connect through an external switch.

## 0.8.5 Use Case

In this scenario the hypervisor configuration is done at either a branch office with virtual infrastructure, or in the data center where the applications are running on the host VM. So, the local users accessing the WAN leverage the Exinda for diagnostics, shaping, and acceleration.



Screenshot 34: Use-case for Inline deployment with externally attached LAN

In this use-case:

- » The local LAN is sitting behind the hypervisor.
- » The host has four NIC interfaces. NIC 0 is dedicated for management of the system and NIC 1 is idle or used for other purposes.
- » NIC2 and NIC3 are mapped to SW2 and SW3.
- » SW2 and SW3 are mapped to Exinda Virtual Appliance ETH2 and ETH3.
- » ETH2 and ETH3 are mapped to NIC 2 and NIC 3, and are configured and bridged together by the Exinda virtual appliance.
- » The data path from a client on the local LAN goes through the Exinda virtual appliance in Inline Mode and out to the WAN.
- » Exinda Exinda Virtual Appliance software version is 7.4.2.
- » This deployment works with either local or external storage.

### Inline deployment with an isolated virtual LAN and virtual applications

In this use case, the Exinda Virtual Appliance is set up for an inline mode deployment with an isolated virtual LAN and virtual applications. There are primarily two scenarios for inline deployment of the virtual appliance:

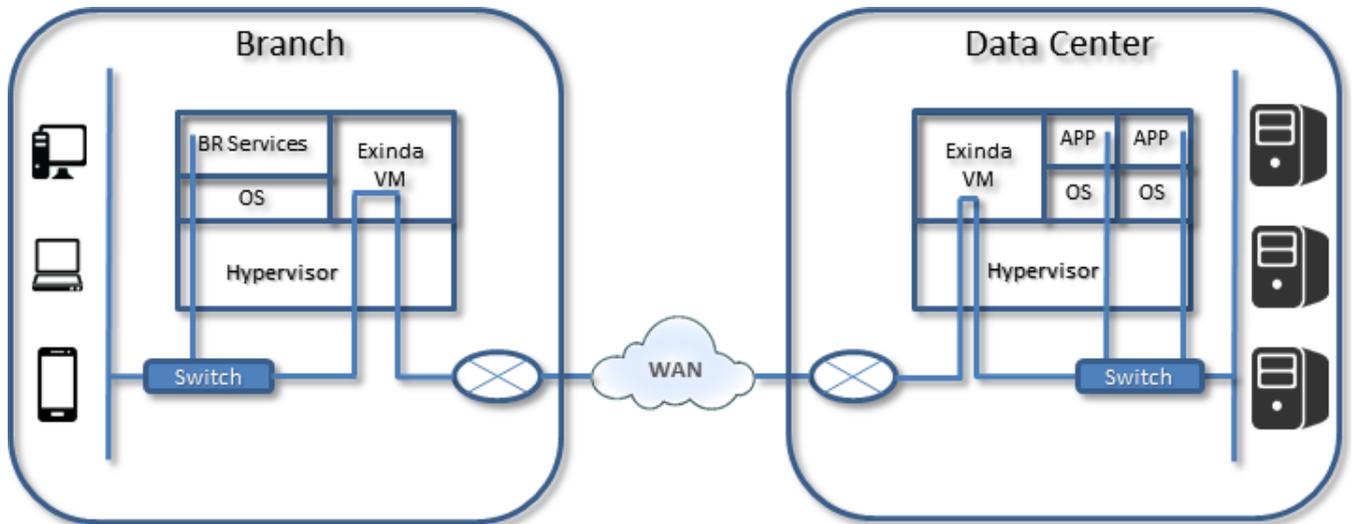
1. The LAN side users and applications connect to the Exinda Virtual Appliance through a physical NIC interface.
2. The applications are virtualized and isolated on the same host as the Exinda on the LAN side interface of the Exinda Virtual Appliance.

This use case discusses the second scenario.

## NOTE

This scenario is only possible using VMware vSphere (v5.5 and v6.0) and Citrix XenServer (v6.2 and v6.5) hosts. Microsoft Hyper-V does not support inline deployments.

### 0.8.6 Scenario

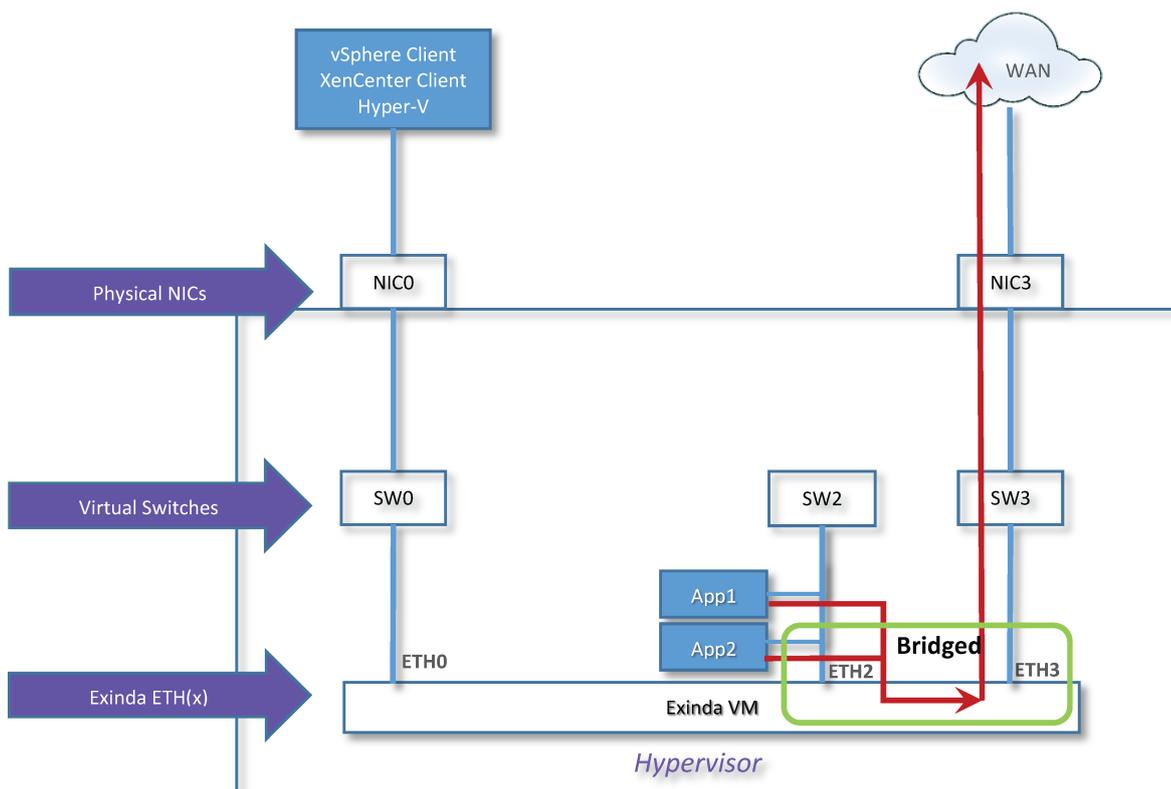


Screenshot 35: Scenario details

- » Branch users access virtualized local services (such as print, DHCP, and Active Directory).
- » One physical NIC interface is configured to the WAN side link.
- » Two virtual switches are bridged together in the virtual Exinda Virtual Appliance. The first switch is on the LAN side; the second one is for the WAN side. This results in isolating the applications behind the Exinda Virtual Appliance.
- » Users connecting from the branch office to the Data Center applications are in line through the Exinda virtual machine(s) on both ends of the connections and through a single external NIC interface to the WAN.
- » The Exinda appliance provides Diagnostics, Shaping, and Acceleration for all traffic in this configuration.

### 0.8.7 Use-case

This configuration will be in either a branch office with virtual infrastructure, or in the data center where the applications are running on the host VM and local users accessing the WAN leverage the Exinda for Diagnostics, Shaping, and Acceleration.



Screenshot 36: Use-case for Inline deployment with an isolated virtual LAN and virtual applications

- » In the diagram above, a virtual LAN is isolated and sits behind the Exinda Virtual Appliance all running on the same host hypervisor.
- » The host has two NICs; NIC 0 is dedicated for management of the system and NIC 1 is idle or used for other purposes.
- » All virtual application workloads are configured in the hypervisor to SW2.
- » SW2 is configured to map to Exinda Virtual Appliance ETH2
- » ETH2 is configured as part of a bridged connection defined as BR2.
- » BR2 bridges NIC 2 and NIC 3 together in the Exinda virtual appliance. The data path for any application connected to the SW2 virtual switch goes through the Exinda Virtual Appliance in inline mode through the ETH2/ETH3 bridged configuration and out the NIC3 interface to the WAN.
- » OPTIONAL: If this is a branch office with local users, configure local users to connect through the NIC2 physical interface and SW2/ETH2 Exinda Virtual Appliance interface and out to the WAN. This require mapping a third NIC interface.

### Out-of-band (WCCP) mode

You can set up your Exinda Virtual Appliance in out-of-band mode, using WCCP protocol for deployment. This deployment is typical for customers who have chosen to redirect a percentage of their traffic for acceleration and traffic shaping through the Exinda Virtual Appliance. In the event the Exinda Virtual Appliance fails, all traffic previously redirected to the Exinda Virtual Appliance goes through un-optimized and un-accelerated.

## 0.8.8 Use Cases

- » Out-of-band (WCCP) mode (VMware ESXi)
- » Out-of-band and High Availability (PBR/VRRP) Mode

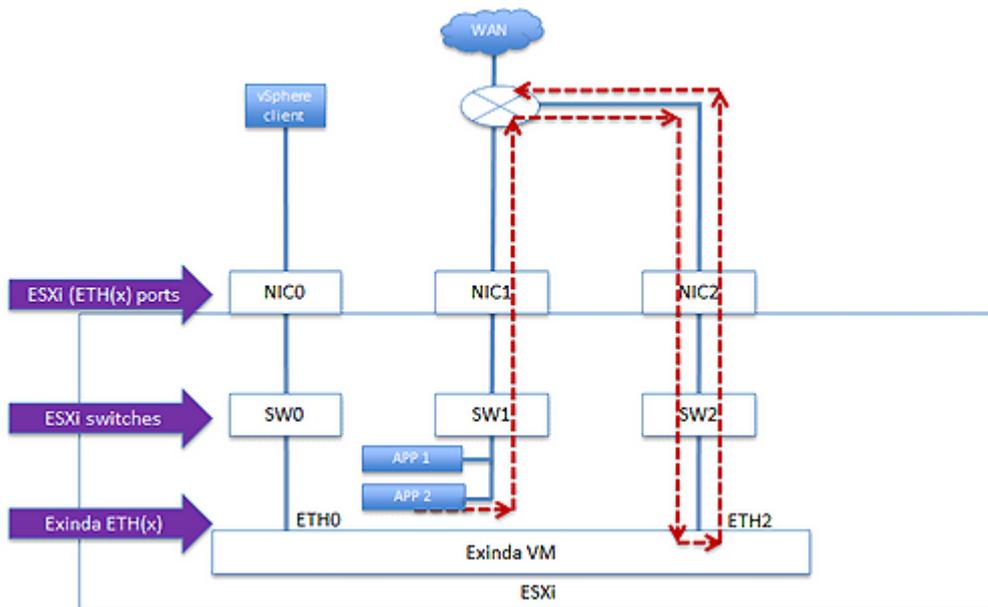
## Out-of-band (WCCP) mode (VMware ESXi)

This configuration can be done either at a branch office with virtual infrastructure, or in the data center where the application and local user traffic accessing the WAN is redirected to the Exinda Virtual Appliance for diagnostics, shaping, and acceleration.

### Scenario

- » Branch users access local services (print, DHCP, Active Directory) that are virtualized.
- » One physical NIC interface is dedicated for management of the virtual machines and hypervisor.
- » One physical NIC interface is configured on the hypervisor for LAN access, and has a routed connection to the WAN router.
- » One physical NIC interface is configured and mapped to the AUX port on the Exinda Virtual Appliance.
- » OPTIONAL: It is possible to configure and map all data traffic through a single NIC interface on the hypervisor; however, for performance reasons it is recommended to segment the un-optimized traffic from the optimized traffic.
- » User access from the branch office to the Data Center applications has two paths:
  - Path one is directly to the WAN router, with no traffic shaping or acceleration.
  - Path two is through the re-directed path invoked by the router (using WCCP) to the Exinda virtual appliance. The traffic is optimized and accelerated.
- » Traffic on the Data Center side has the same path as the branch side. Traffic that is selected to be optimized and accelerated is redirected to the Exinda Virtual Appliance through WCCP on the WAN router.
- » The Exinda provides Diagnostics, Shaping, and Acceleration for only redirected traffic in this configuration

### Use Case



Screenshot 37: Use-case for Out-of-band (WCCP) mode (VMware ESXi)

In this use case:

- » The virtual LAN with application servers (APP1 W2003 or 8 and APP2) are configured in the ESXi hypervisor on SW1 and mapped to NIC1. They have a direct path the WAN router.

» The physical server is a re-purposed Dell Server and has the following physical/logical interface mappings for illustration purposes only (you are encouraged to select your own server):

|                       |      |      |      |
|-----------------------|------|------|------|
| ESXi Ethernet Ports   | NIC0 | NIC1 | NIC2 |
| Exinda Ethernet Ports | ETH0 | ETH1 | ETH2 |
| ESXi Virtual Switches | SW0  | SW1  | SW2  |

- » The host has three NICs; NIC 0 is dedicated for management of the system.
- » NIC 1 is dedicated to all virtual application workloads hosted on the ESXi.
- » All virtual application workloads are configured in the ESXi to SW1.
- » The Exinda Virtual Appliance is configured on SW2 virtual switch and is mapped to the NIC2 interface.
- » The NIC2 interface has a direct connection to the WAN router, and is configured for WCCP GRE layer 3 mode between the router and the Exinda Virtual Appliance.
- » VMware version = 5.5 and 6.0
- » Exinda Virtual Appliance firmware version = 7.4.9
- » The data path for virtualized applications configured on SW1 takes two paths:
  - Path one – un-optimized and un-accelerated traffic is forwarded directly to the WAN router through NIC1.
  - Path two – traffic to be optimized (traffic shaped) and accelerated traffic is forwarded to the router for redirection through WCCP to the Exinda Virtual Appliance through NIC2/SW2 on the AUX port of the Exinda Virtual Appliance.
- » (Optional) It is possible to configure and map all data traffic on SW1 and SW2 to the NIC1 interface; however, for performance reasons it is recommended to segment the optimized traffic on its own NIC and virtual switch for performance reasons, and in the event of failure of the Exinda Virtual Appliance.

### Out-of-band and High Availability (PBR/VRRP) Mode

In this use case, there are multiple Exinda Virtual Appliances set up for out-of-band mode and configured for High Availability. PBR is the protocol used for redirection while the appliances act as VRRP nodes. One Exinda acts as the Master node, receiving all redirected traffic from the WAN router, and all remaining Exinda Virtual Appliances act as Backup nodes in case the Master fails.

This deployment would be typical for customers who are choosing to redirect a percentage of their traffic for acceleration and traffic shaping through the Exinda virtual appliances.

To configure PBR/VRRP on the Exinda Virtual Appliance, see [VRRP with PBR](#).

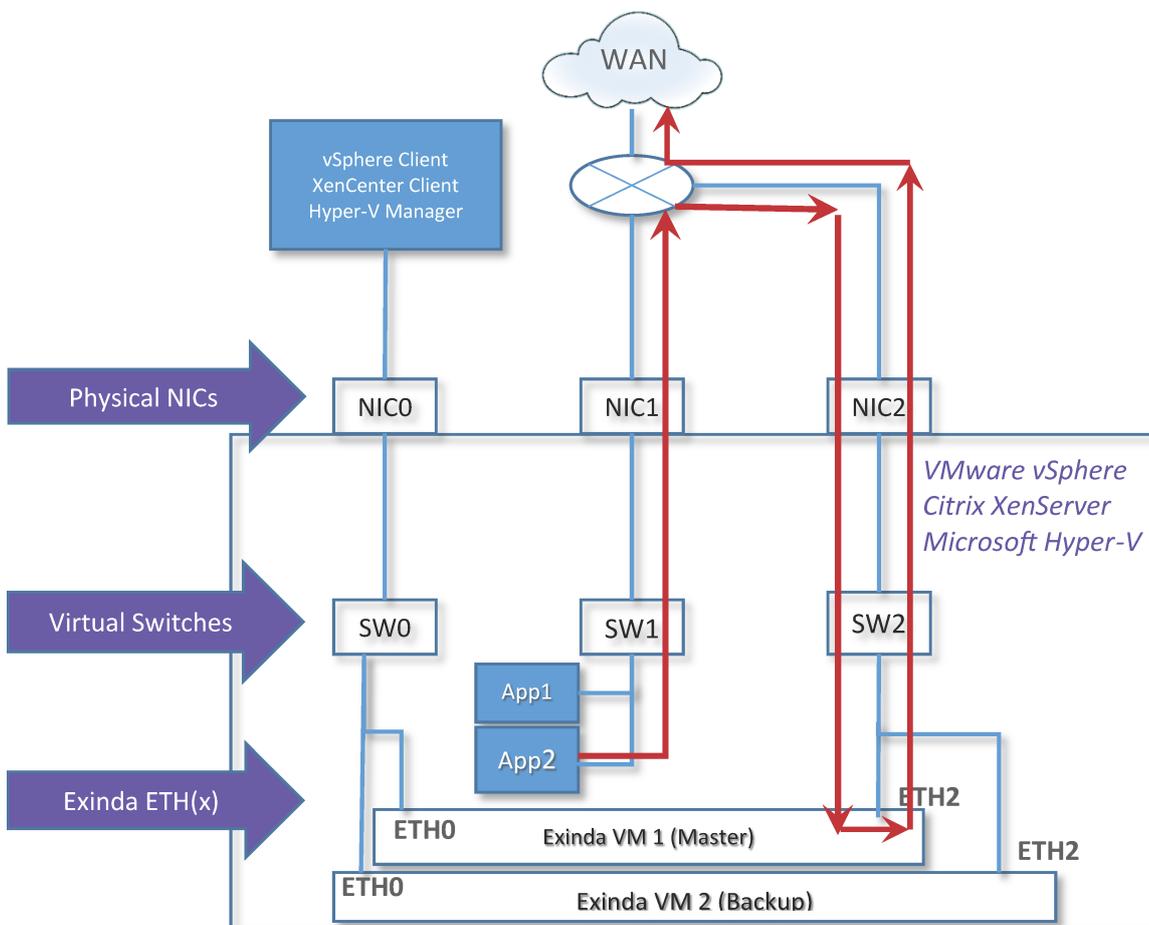
This configuration would work in either a branch office with virtual infrastructure, or in a data center where the application and local user traffic accessing the WAN is redirected to the Exinda virtual appliance for Diagnostics, Shaping, and Acceleration.

### Scenario

- » Branch users access virtualized, local services (print, DHCP, Active Directory, etc.).
- » One physical NIC interface is dedicated for management of the virtual machines and hypervisor.
- » One physical NIC interface is configured on the hypervisor for LAN access, and has a routed connection to the WAN router.

- » One physical NIC interface is configured and mapped to the virtual switch that is connected to the virtual AUX ports on all the Exinda Virtual Appliances.
- » User access from the branch office to the Data Center applications has two paths:
  - Path one is directly to the WAN router, with no traffic shaping or acceleration.
  - Path two is through the re-directed path invoked by the router (using PBR) to the Exinda virtual appliances. The traffic is received by the Exinda configured as the VRRP Master. In case of acceleration, the Exinda Virtual Appliances have the capability to share load, this happens automatically and requires no user intervention.
- » Traffic on the Data Center side has the same path as the branch side. Traffic that is selected to be optimized and accelerated is redirected to the Exinda Virtual Appliance through PBR on the WAN router.

## Use Case



Screenshot 38: Use-case for Out-of-band and High Availability (PBR/VRRP) Mode

In this diagram:

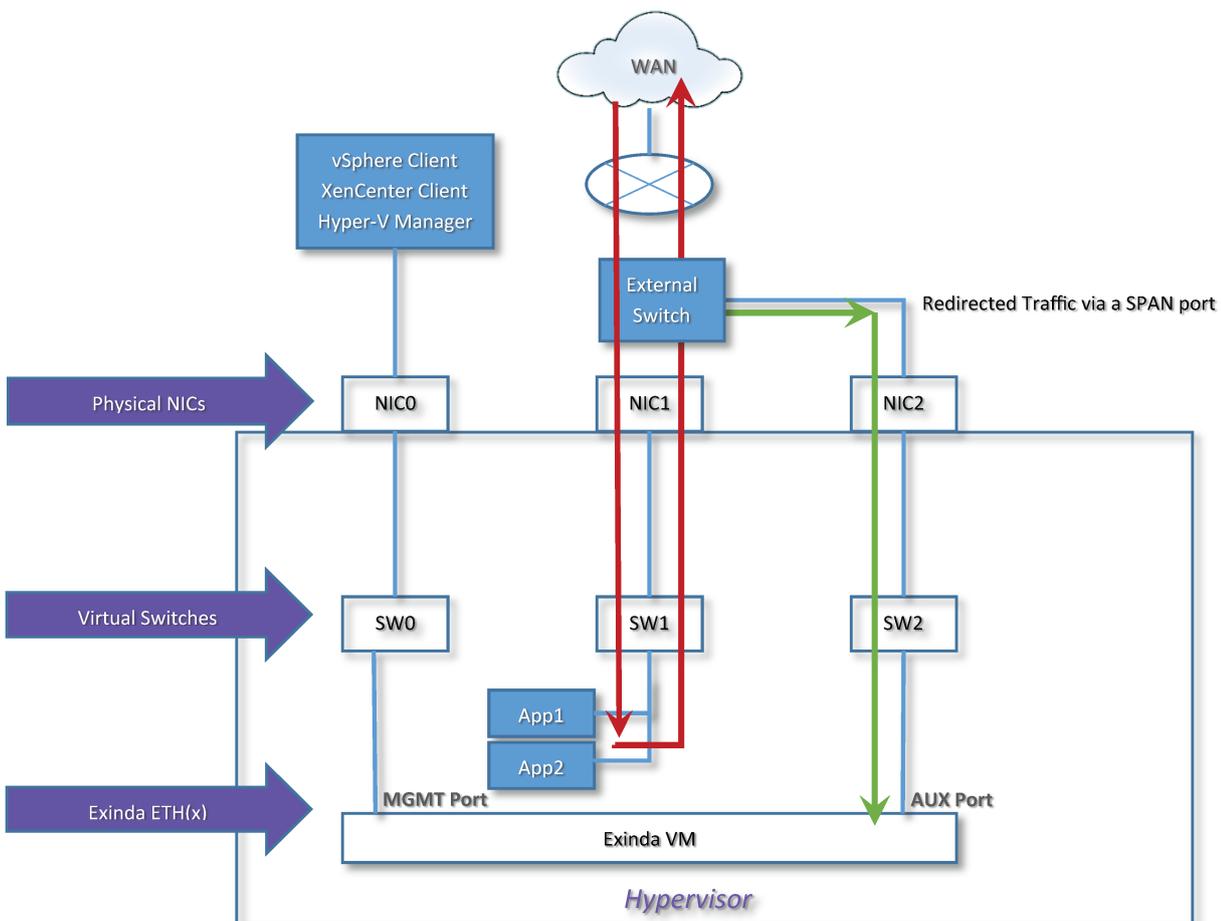
- » The virtual LAN with application servers (APP1 and APP2) are configured in the hypervisor on SW1 and mapped to NIC1. They have a direct path the WAN router.
- » For simplicity, both Exinda Virtual Appliances installed in the same host hypervisor. However, it is possible to install Exinda Virtual Appliances on separate hypervisors to achieve physical high availability
- » The host has three NICs; NIC 0 is dedicated for management of the system.
- » NIC 1 is dedicated to all virtual application workloads hosted on the hypervisor
- » All virtual application workloads are configured in the hypervisor to SW1.

- » The Exinda Virtual Appliances are configured on SW2 virtual switch and are mapped to the NIC2 interface.
- » The NIC2 interface has a direct connection to the WAN router, which configured for PBR redirection. For more information, refer to [Deployment options](#) (page 26).
- » Exinda Virtual Appliance firmware version = 7.4.2.

### Port mirroring/SPAN port Configuration

Port mirroring/SPAN port Configuration can be used to perform network audits as it provides great flexibility in restricted and complex network environments.

In this use case, the Exinda Virtual Appliance is set up to monitor and collect traffic for only reporting only, without installing the appliance in the inline mode. The appliance monitors and reports on all applications presented on the SPAN/mirror port.



Screenshot 39: Use-case for Port mirroring/SPAN port Configuration

In this use-case:

- » The Exinda Virtual Appliance is used for monitoring and reporting, and is configured in the hypervisor to use the dedicated NIC2 interface, and dedicated virtual switch SW2.
- » The Exinda Virtual Appliance management port is mapped to SW0 and NIC0.
- » The Exinda Virtual Appliance has three Ethernet interfaces (ETH0–2) with the APP1 and APP2 both mapped to SW1 and NIC1, and connected to the external Switch.
- » The AUX (ETH1) port is configured for SPAN port mirroring and management of the Exinda Virtual Appliance.

- » The NIC2 interface is connected to an external switch on a port that has been configured to support SPAN port mirroring. It is recommended that this port be dedicated so there is no impact to traffic performance.
- » Application virtual workloads (APP1 and APP2) are on a separate virtual SW1 and mapped to NIC1.
- » NIC1 is directly attached to the external switch.
- » The switch can also be a virtual one installed in the hypervisor.

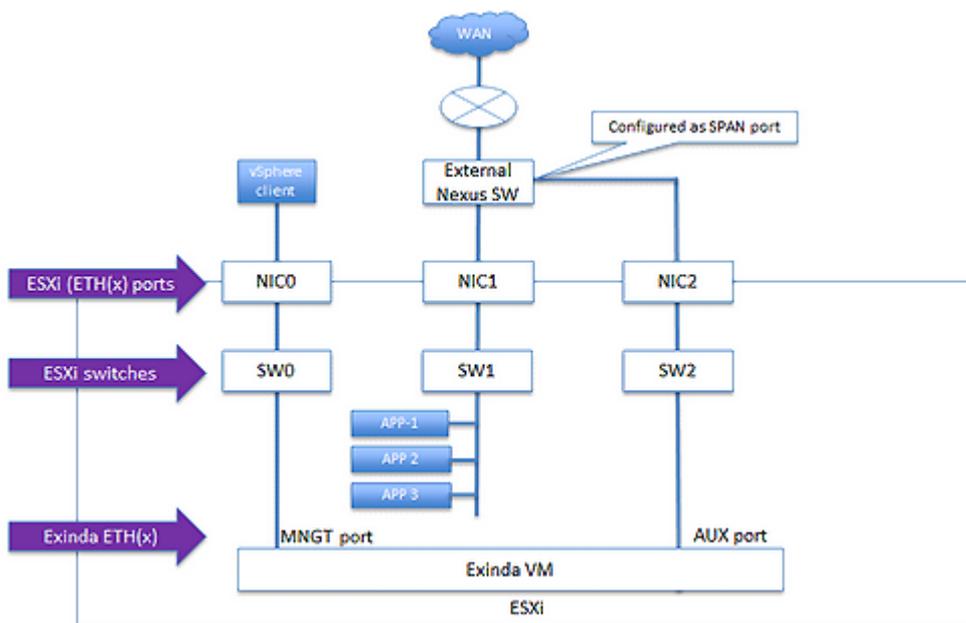
## Related topics

- » [Port mirroring with an external Nexus switch](#)
- » [Port mirroring with a virtual Nexus switch](#)

## Port mirroring with an external Nexus switch

### NOTE

in case of VMware, this configuration supports vMotion, HA, and Fault Tolerant.



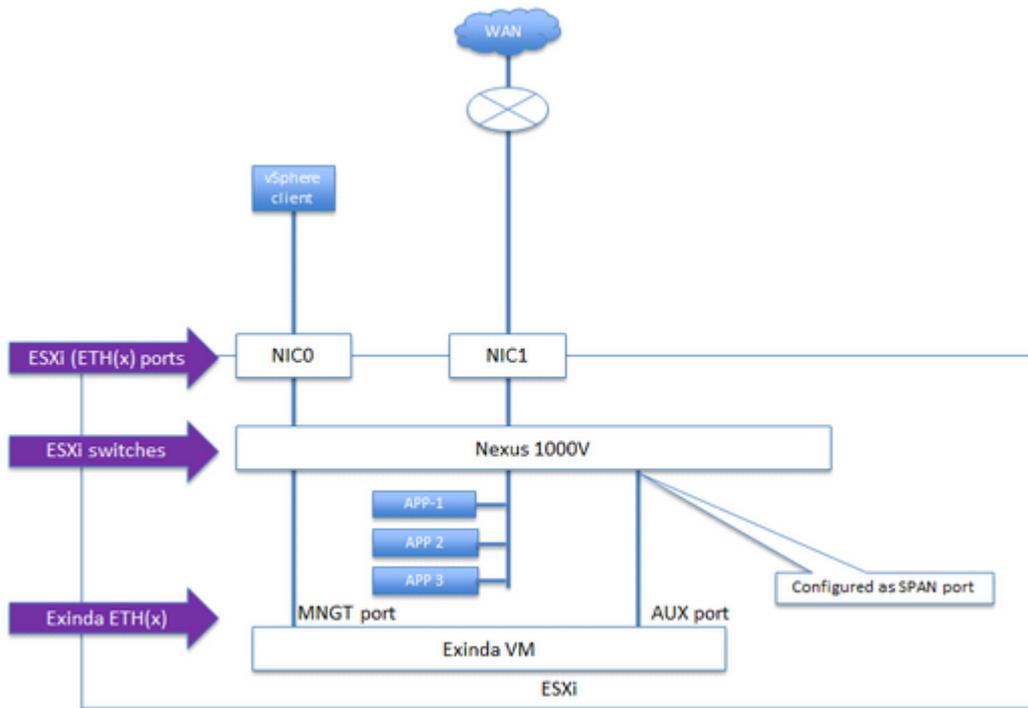
Screenshot 40: Use-case for Port mirroring with an external Nexus switch

In this use-case:

- » The customer has selected Exinda for its monitoring and reporting service. The EXN-V is configured as a virtual machine on a hypervisor or as an ESXi hypervisor (for VMware), on a dedicated NIC2 interface, and dedicated virtual switch SW2.
- » The EXN-V management port is mapped to SW0 and NIC0.
- » The host EXN-V has three four Ethernet interfaces (0-23) with the APP1-3 mapped to SW1 and NIC1 respectively, and connected to the external Nexus Switch.
- » The EXN-V with the AUX (ETH1) port is configured for SPAN Port Mirroring and management of the EXN-V.
- » The NIC2 interface is connected to an external switch on a port that has been configured to support SPAN port mirroring. It is recommended that this port be dedicated so there is no impact to traffic performance.

- » Application virtual workloads (APP1 to 3) are on a separate virtual SW1 and mapped to NIC1.
- » NIC1 is directly attached to the external Nexus switch.

### Port mirroring with a virtual Nexus switch



Screenshot 41: Use-case for Port mirroring with a virtual Nexus switch

In this use-case:

#### VMware

##### NOTE

In this configuration VMware supports vMotion, HA, and Fault Tolerant.

» The Cisco Nexus 1000V Series VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware Virtual Switch functionality. The VEM uses the VMware vNetwork Distributed Switch (vDS) API, which was developed jointly by Cisco and VMware, to provide advanced networking capability to virtual machines. This level of integration helps ensure that the Cisco Nexus 1000V Series is fully aware of all server virtualization events, such as VMware VMotion and Distributed Resource Scheduler (DRS). The VEM takes configuration information from the VSM and performs Layer 2 switching and advanced networking functions namely Monitoring:

- NetFlow
- Switch Port Analyzer (SPAN)
- Encapsulated Remote SPAN (ERSPAN)

» VMware versions - 5.5 and 6.0

#### XenServer

» EXN-V versions – 7.4.9

» The Nexus 1000V is configured and mapped to NIC1, which has a direct connection to the WAN router.

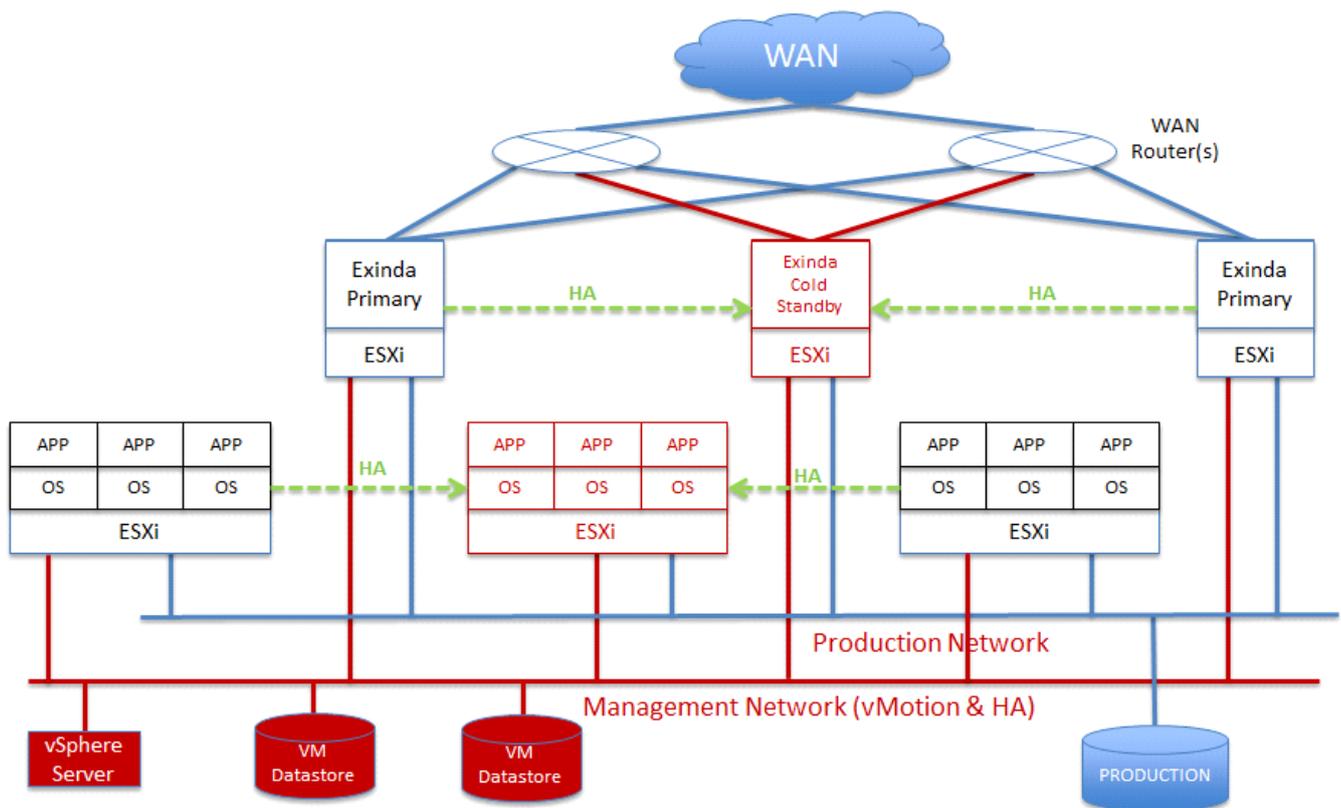
- » The EXN-V has two four Ethernet interfaces (ETH0 / ETH10-3) with the AUX (ETH1) configured for Mirroring and ETH0 for management and management of the EXN-V.
- » The EXN-V AUX port is configured to a port configured with SPAN port mirroring on the Nexus 1000V. This port should be dedicated to ensure there is no performance impact to data traffic.
- » The applications (APP1 to 3) are connected to a separate switch port on the Nexus 1000V.

### VMware High Availability (HA)

In this use case, we discuss the recommended configuration and best practices for installing the Exinda virtual appliance where the requirement is for:

- » Exinda firmware – 7.4.9
- » ESX/ESXi HA software versions 5.5 and 6.0
- » vMotion support,
- » vMotion = Yes for INLINE-line
- » VMware best practice recommends that at least three hosts are used for this configuration, and
- » Licensing for the Exinda Virtual Appliance units include two full licenses and one cold standby license.

Each Exinda Virtual Appliance must maintain network connectivity with the Exinda License server and will shut down the Exinda virtual appliance after 96 hours without a successful connection.



Screenshot 42: Use-case for High Availability (HA) on VMware

In this use-case:

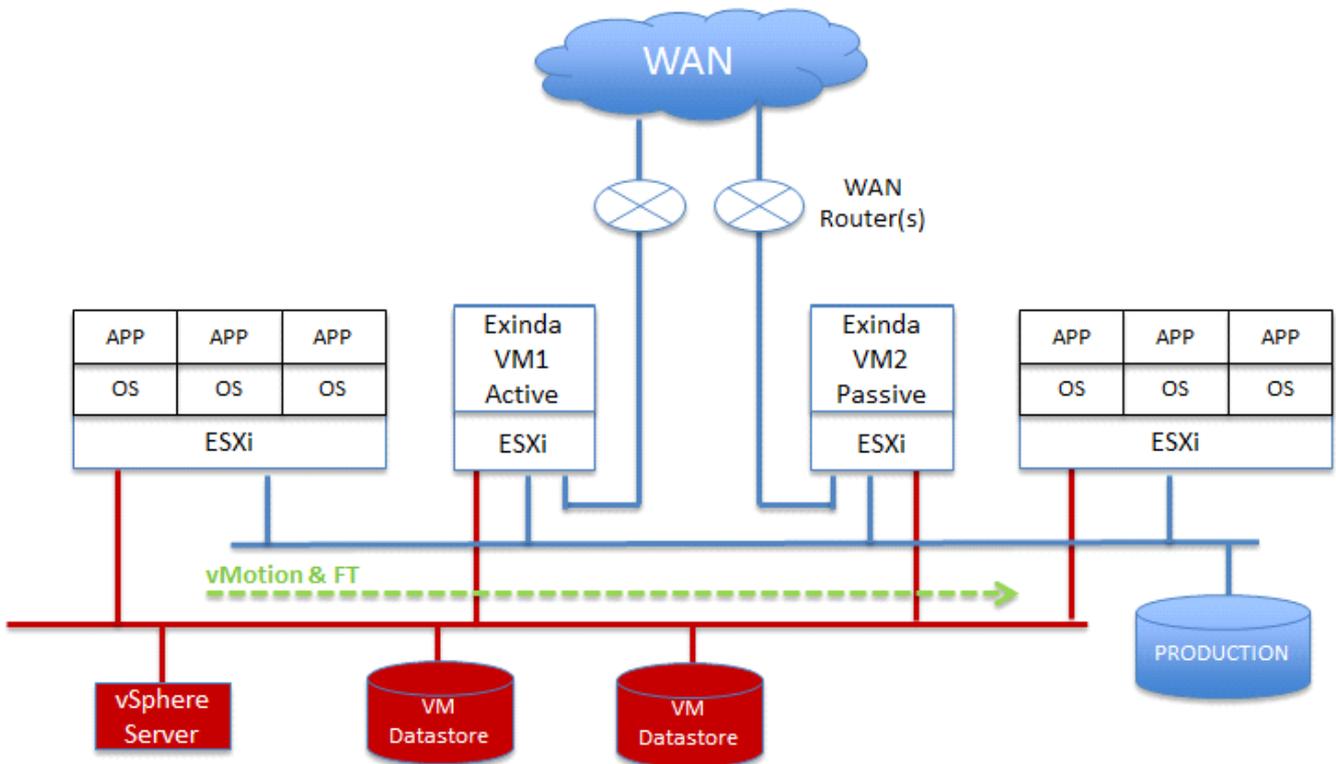
- » There are six ESX/ESXi hosts (can be done with three):
  - Two running virtual workloads, and a third as the backup HA system
  - Two running Exinda Virtual Appliances and a third running as a cold standby

- » Optionally you can move the Exinda Virtual Appliance to co-reside on the same hosts as the virtual workloads; however if vMotion is a requirement you must provide an external switch and separate NIC to pass the traffic between the workloads and the Exinda Virtual Appliance.
- » Having a separate host for the Exinda Virtual Appliance allows you to:
  - Segment other virtual appliances from the application workloads.
  - Support vMotion just for the application workloads and not for the host running the Exinda virtual appliance.
- » There are two networks:
  - A management network for vMotion and access to the external workload VMDK data stores.
  - A production network for data traffic to and from the applications and WAN.
- » This configuration assumes INLINE mode; optionally you can run in out of path mode, but WCCP is required.
- » Downtime for any workload in HA mode is for the duration of the virtual workload and/or the Exinda Virtual Appliance to reboot.

### VMware Fault Tolerance (FT) cluster

In this use case, we discuss the recommended configuration and best practices for installing the Exinda Virtual Appliance on VMware Fault Tolerance (FT) cluster. For this use-case, you require:

- » Exinda firmware version – 7.4.9
- » Hypervisor Fault Tolerance Cluster (software versions 5.5 and 6.0)
- » vMotion support
- » vMotion = Yes for INLINE
- » VMware best practice recommends that at least 4 hosts are used for this configuration
- » Licensing for the Exinda Virtual Appliance units will include one full license and one cold standby license.
  - Each Exinda Virtual Appliance must maintain network connectivity with the Exinda License server and will shut down the Exinda virtual appliance after 96 hours without a successful connection.
- » Optionally, the Exinda Virtual Appliance can co-reside with the Application workloads, but an external switch and an additional NIC is required for vMotion support.

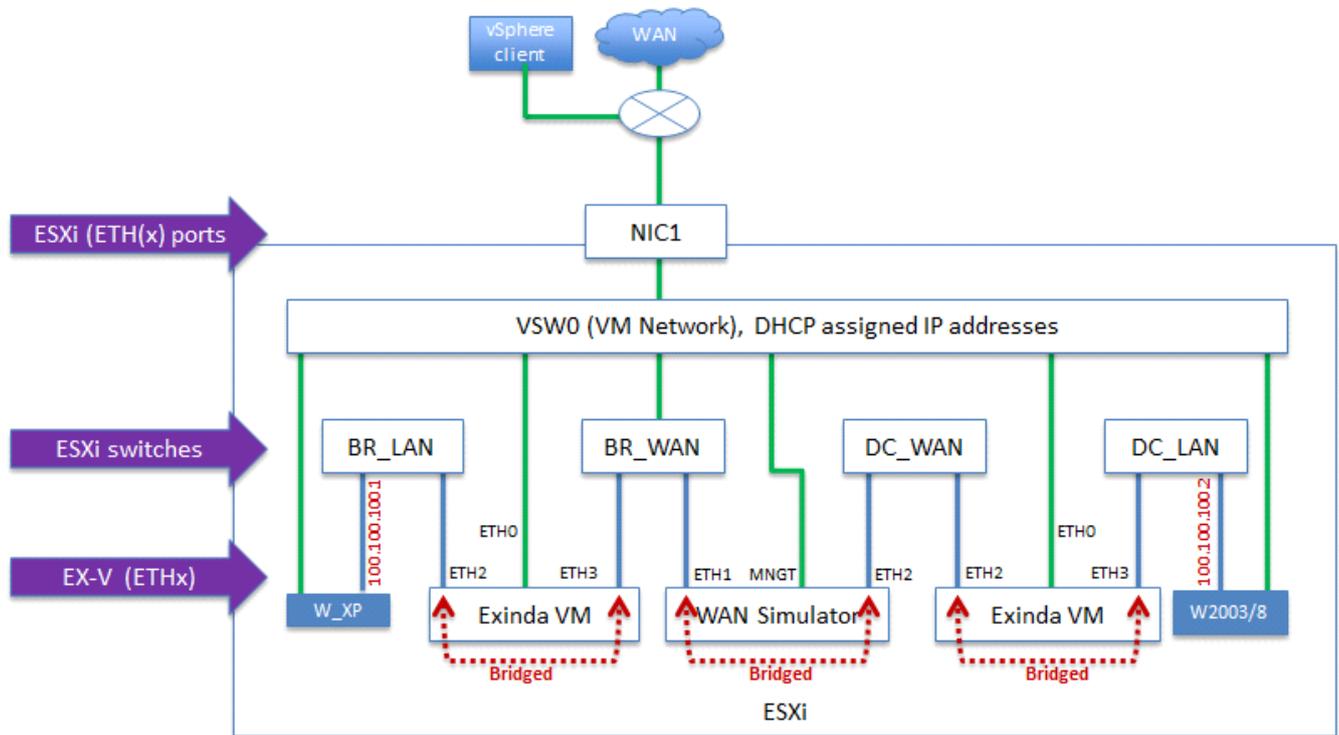


Screenshot 43: Use-case for VMware Fault Tolerance (FT) cluster

In this use-case:

- » There are four hypervisor hosts:
  - one running virtual workloads and a second as the Fault Tolerance system
  - one running Exinda Virtual Appliance active and a second running as a cold standby and Fault Tolerant.
- » Optionally, move the Exinda Virtual Appliance to co-reside on the same hosts as the virtual workloads; however if vMotion is a requirement you must provide an external switch and separate NIC to pass the traffic between the workloads and the Exinda Virtual Appliance.
- » Having a separate host for the Exinda Virtual Appliance allows you to:
  - Segment other virtual appliances from the application workloads.
  - Support vMotion just for the application workloads, and not for the host running the Exinda virtual appliance.
- » There are two networks:
  - A management network for vMotion and access to the external workload VMDK data stores
  - A production network for data traffic to and from the applications and WAN
- » This configuration assumes INLINE mode; optionally you can run in out of path mode which requires WCCP.

## Virtual WAN simulator in an isolated network (VMware ESXi)



Screenshot 44: Use-case for Virtual WAN simulator in an isolated network

In this use-case:

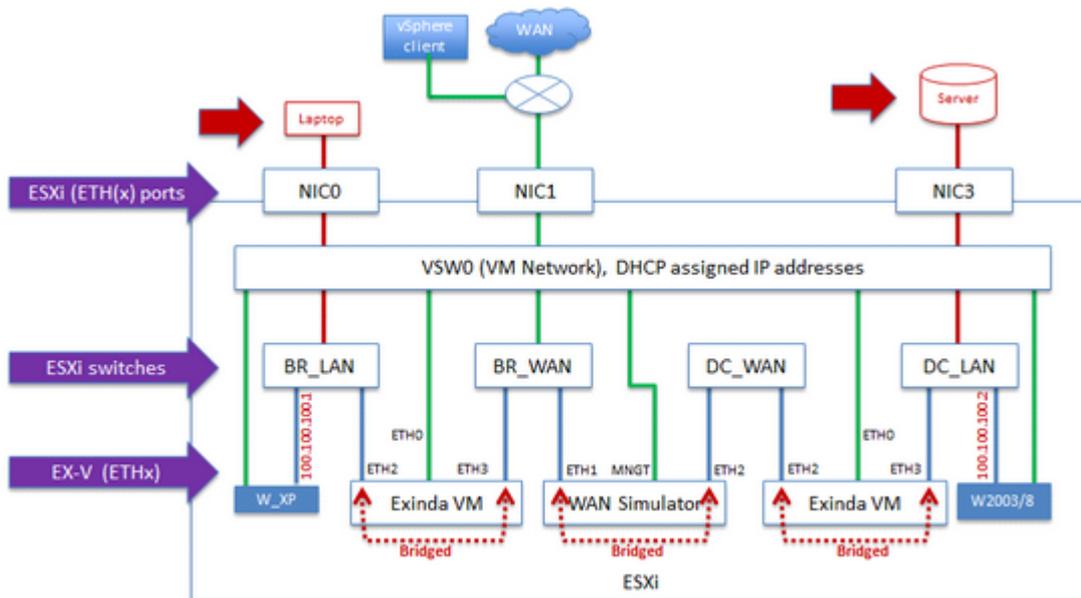
- » The hypervisor has one NIC:
  - NIC0, NIC2, and NIC3 are unused
  - NIC1 is connected to an external network and has access to the WAN for management and licensing of the virtual machines.

### NOTE

Any physical NIC interface can be used, NIC1 is used for illustration purposes.

- » VMware software version = 5.5 and 6.0
- » Exinda Virtual Appliance firmware version = 7.4.9
- » WANEM Virtual Simulator software = 2.3
  - <http://wanem.sourceforge.net/>
  - You can use your own WAN simulator of choice
- » Four virtual switches have been defined on the ESX/ESXi host:
  - BR\_LAN – branch side LAN switch
  - BR\_WAN – branch side WAN switch
  - DC\_WAN – data center side WAN switch
  - DC\_LAN – data center side LAN switch

- » Each Exinda Virtual Appliance is configured for INLINE Mode and a single management interface on ETH0.
- » DHCP is assumed on the network for management interfaces on the Exinda Virtual Appliance appliances.
- » Private network space is configured for the Windows Client and Server on the data path between them and a second Ethernet interface is configured for DHCP to manage each system through RDP.
- » Exinda Virtual Appliance and the WAN Simulator data path are bridged.
- » Optionally, you can configure the WAN Simulator as a router and change the default gateway of the client and server accordingly.
- » Optionally with a system that has at least 3 NIC interfaces you can attach an external workstation and server and pass traffic through the demo system.
  - You will need to configure on the ESX/ESXi host mapping BR\_LAN to NIC0 and DC\_LAN to NIC3 to connect the external workstation and server.
  - The benefit is you can test through the isolated virtual Exinda Virtual Appliance environment with no impact to a product network



Screenshot 45: WAN external workstation

## Hypervisor limitations

Consider these additional planning items when installing Exinda Virtual Appliance on other Hypervisors in the market including:

- » XenServer 6.2, and 6.5
- » Microsoft Hyper-V on Windows 2008 R2, 2012, and 2012 R2

### XenServer

- » No Silicom Hardware Bypass Card driver support. Exinda is working with our NIC vendor to provide this support.
- » Promiscuous mode is supported, and must be configured via the CLI. You can find the commands in the Exinda Virtualization How to Configure Guide, or refer to the Citrix site for XenServer hypervisor configuration guidance.
- » Traffic shaping, reporting, and optimization are supported for INLINE mode.
- » Reporting and optimization are supported for out of path (WCCP GRE) mode.

## VMware

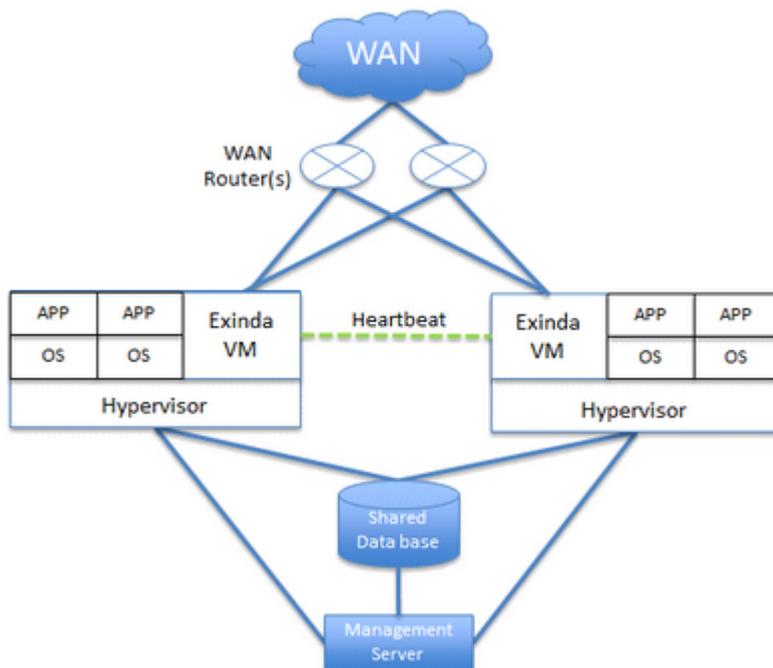
- » Bypass NIC card drivers are supported in 5.5 and 6.0.

## Hyper-V

- » No support for Parallel virtualization drivers, meaning the guest machine does not know it's virtualized.
- » Only out-of-band deployments are supported for Hyper-V.

External storage is supported and recommended for virtual machine workloads, and the Exinda virtual appliance

In the diagram below, the Exinda is running in Active/Active mode with a Heartbeat between the two systems. There must be a separate Virtual NIC configured for Heartbeat traffic to transit.



## Running on VMware vSphere (ESX and ESXi)

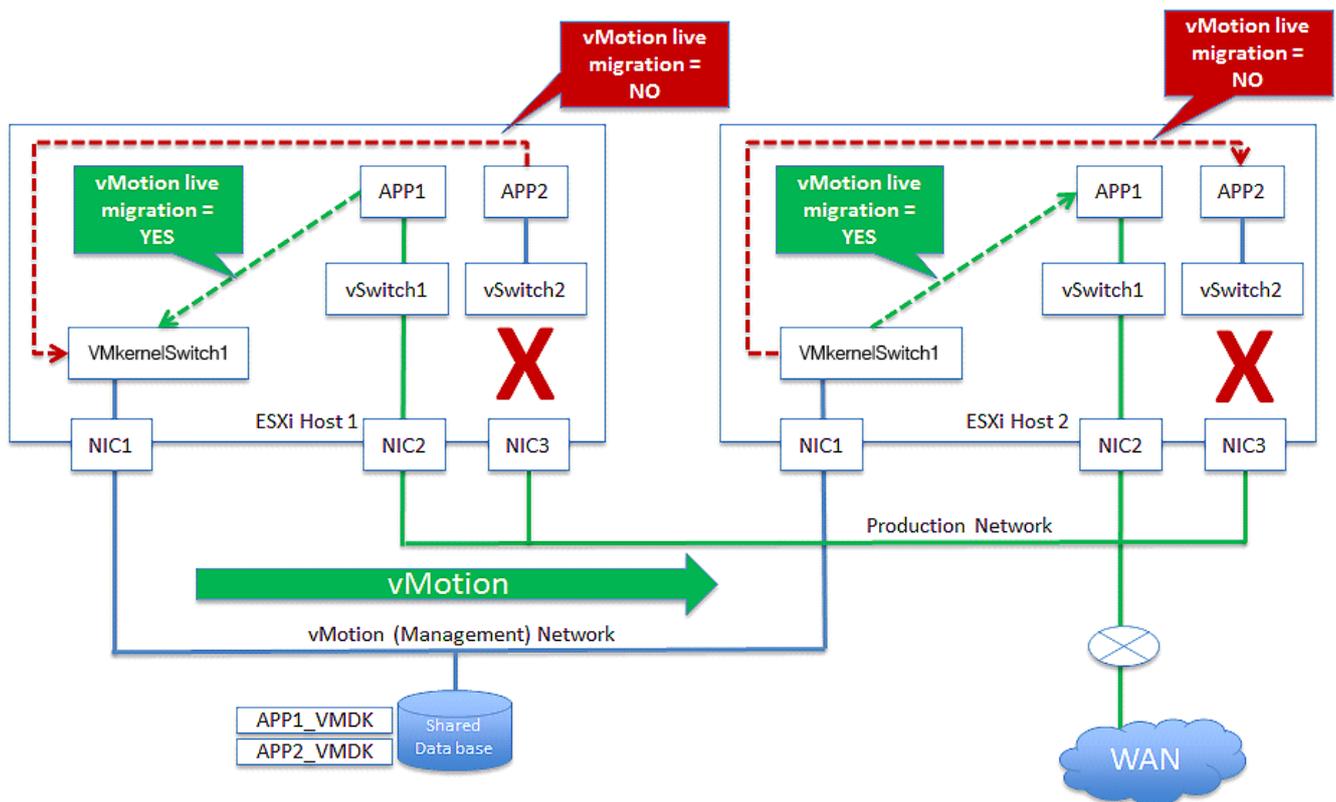
Learn how to run and customize the Exinda Virtual Appliance on VMware vSphere client.

### NOTE

VMware ESX/ESXi 5.5 (or later) is required.

## Understanding how vMotion works

For isolated virtual applications on the Exinda Virtual Appliance LAN port for inline mode, the VMware vMotion feature does not work. It is a requirement from VMware that any virtual switch must be mapped to a physical NIC and then to an external network. Below is a brief illustration of the process.



- » There are two types of virtual switches in the ESX/ESXi hypervisor: VMkernel Switch and vSwitch. The VMkernel Switch is used by the hypervisor exclusively. VMkernel is the bare metal hypervisor, and provides core and memory allocation, disk and network virtualization, and a driver to low level devices. The vSwitch is used by virtual machines, and behaves just like any external layer 2 switch. All virtual machines have a path to the external Data-store where each VMDK is stored through the hypervisor layer to the VMkernel Switch mapped to the NIC attached to the storage.
- » There are two networks:
  - a. Management network where vMotion moves workloads between ESXi hosts
  - b. Production network where the applications are accessed by the users
- » The VMkernelSwitch1 is mapped to external NIC1 and connected to the management network.
- » The vSwitch1 is mapped to NIC2 and connects APP1 to the production network.
- » The vSwitch2 is mapped to APP2 but does not have a mapping to external NIC3. The use case for this is that a network administrator may have one, or many, virtual workloads isolated on the host for testing purposes.
- » vMotion is executed for APP1 on ESXi 1 and moved over to ESXi 2 with no disruption to the application workload.
- » vMotion is executed for APP2 on ESXi 1 and fails because vSwitch2 mapped to APP2 is not mapped to an external NIC.
- » If an Exinda Virtual Appliance has at least one vSwitch mapped to it, and the vSwitch is not mapped to an external NIC interface, vMotion will not work for Exinda Virtual Appliance or workloads isolated behind it.

### Install the Virtual Appliance on VMware

1. Liaise with your local Exinda representative to obtain the download files.
2. Open the VMware vSphere client.
3. Select **File > Deploy OVF Template**.

4. Copy the URL of the latest release of the Exinda VMware Virtual Appliance from Exinda.com, and paste it into the **Deploy from...** field, and click **Next**.
5. Confirm the OVF template details are correct, and click **Next**.
6. Review and accept the End User License Agreement (EULA), and click **Next**.
7. Specify a name for the virtual appliance. If prompted, choose the location to deploy the virtual appliance, and click **Next**.
8. Choose the format to store the virtual disks for the virtual appliance. Exinda recommends **Thick Provisioning**(the default).

#### NOTE

By default, the Virtual Appliance is configured with a single 50GB disk. Additional storage can be added in the form of another disk after the Virtual Appliance has been deployed. See the [Additional Storage](#) section for more information.

9. Connect the network interfaces to the appropriate network, by doing the following:
  - a. Connect the Management interface to a network where you can manage the virtual appliance.
  - b. If you are configuring the virtual appliance for clustering, high availability, or out-of-path deployments, map the AUX interface to the appropriate network. This interface can be left disconnected if it is not required.
  - c. If you are deploying the virtual appliance in line, [add additional NICs](#).
4. Click **Next**.
5. Review the deployment settings, and click **Finish** to complete the deployment.

## 0.8.9 Related Topics

Review the following topics after completing the VM deployment:

- » [Adjusting the RAM available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Add Storage to the VMware Virtual Machine](#)

### Modifying the VMware Virtual Machine Configuration

To improve the performance of the virtual appliance, change the number of CPUs, the RAM, networking, and storage allocated to the virtual machine.

#### NOTE

You will need to shut the virtual appliance down before you can modify its configuration.

## 0.8.10 Related Topics

- » [Adjusting the number of CPUs available to the Virtual Machine](#)
- » [Adjusting the RAM available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Converting two NICs into a Bridge](#)

- » Allow Ports to Accept and Bridge Packets (Promiscuous Mode)
- » Add Storage to the VMware Virtual Machine

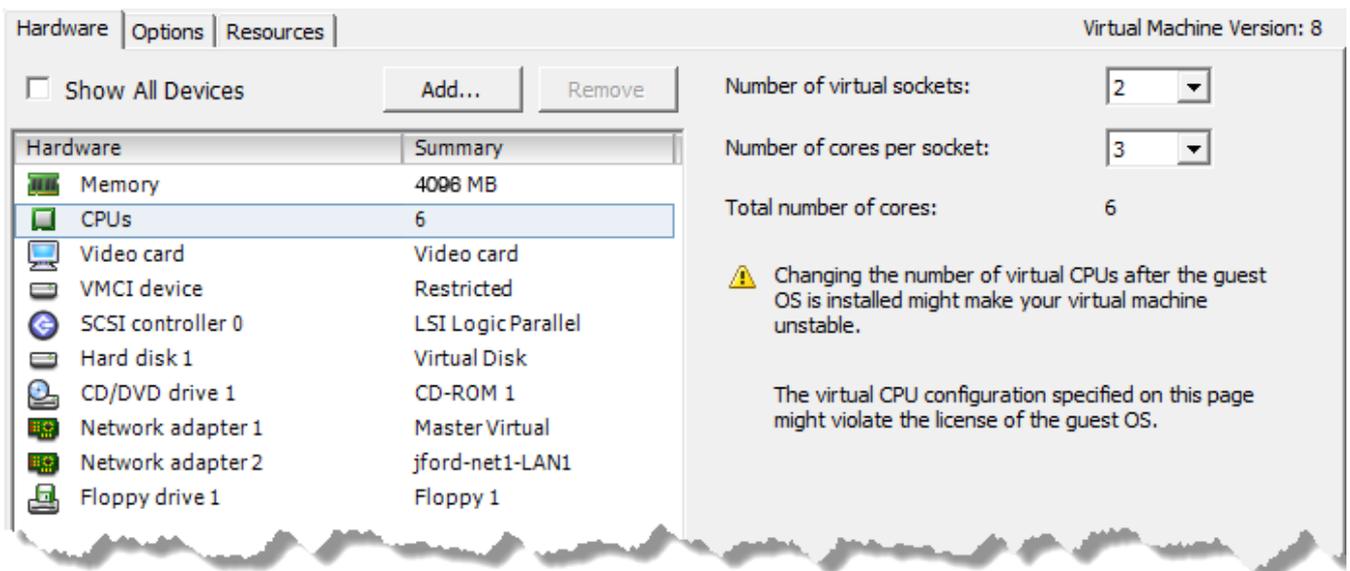
### Adjusting the number of CPUs available to the Virtual Machine

By default, all Virtual Appliances come configured with two virtual CPUs. Increase the number of CPUs to suit your requirements.

#### NOTE

If the memory or hard disk space needs to be adjusted, please contact Exinda Support.

1. Open the **VMware vSphere Client**.
2. Right-click on the Exinda Virtual Appliance, and select **Edit Settings**.
3. On the **Hardware** tab, select **CPUs**.
4. Select the **Number of virtual sockets**.
5. Select the **Number of cores per socket**. The resulting total number of cores is a number equal to or less than the number of logical CPUs on the host. For example, if the **Number of virtual sockets** is 2, and the **Number of cores per socket** is 3, the total number of cores will be 6. Show Image...

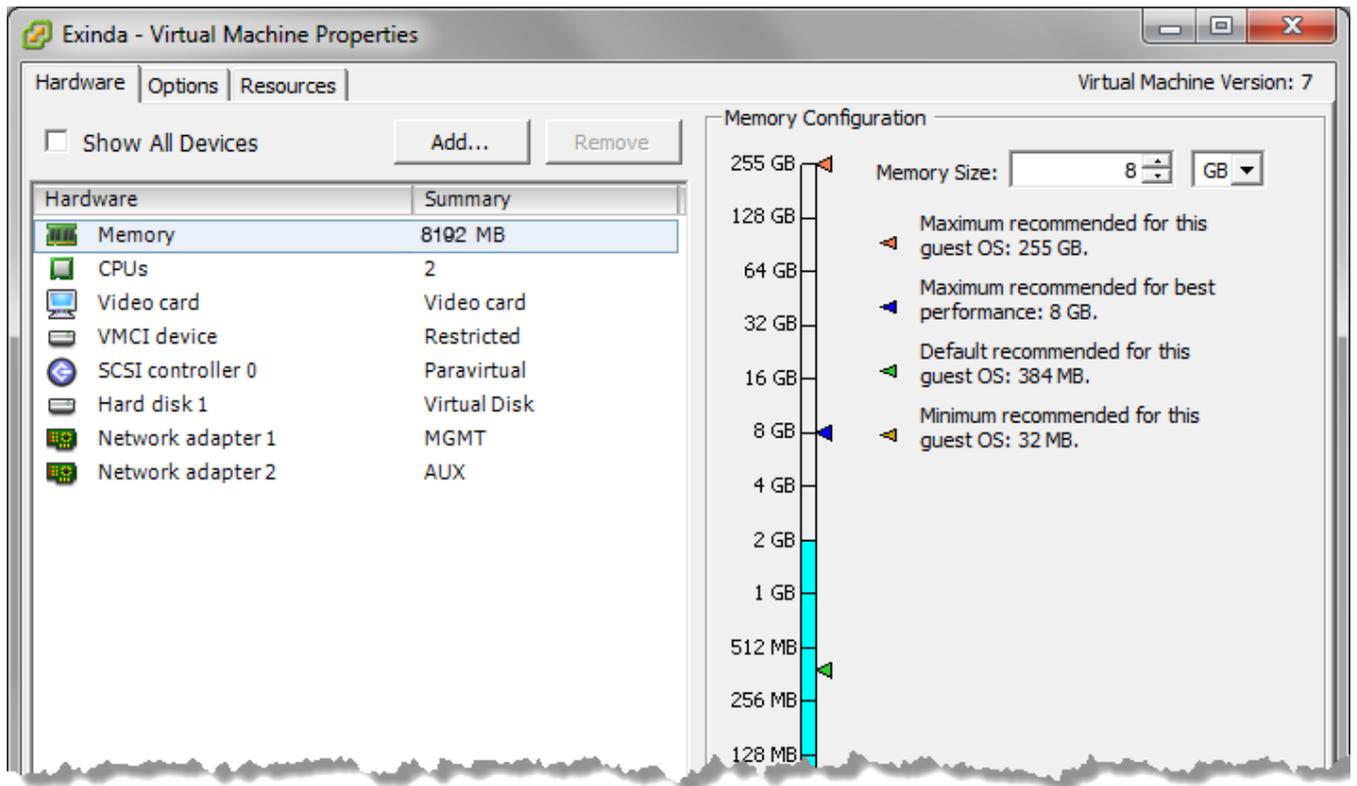


6. Click **OK**.

### Adjusting the RAM available to the Virtual Machine

By default, all Virtual Appliances come configured with 4GB of RAM. Increase the amount of RAM to suit your requirements.

1. Open the **VMware vSphere Client**.
2. Right-click the Exinda Virtual Appliance, and select **Edit Settings**.
3. On the **Hardware** tab, select **Memory**.
4. Click **OK**.
5. Select the desired **Memory Size**



### Adjusting the NICs available to the Virtual Machine

By default, all Exinda Virtual Appliances come with four NICs. Of these, the first NIC is the Management Interface (for managing the Virtual Appliance), the second NIC is the Auxiliary Interface (for HA topologies, clustering and out-of-path deployments), while the remaining two ports are bundled as a bridge for inline deployments.

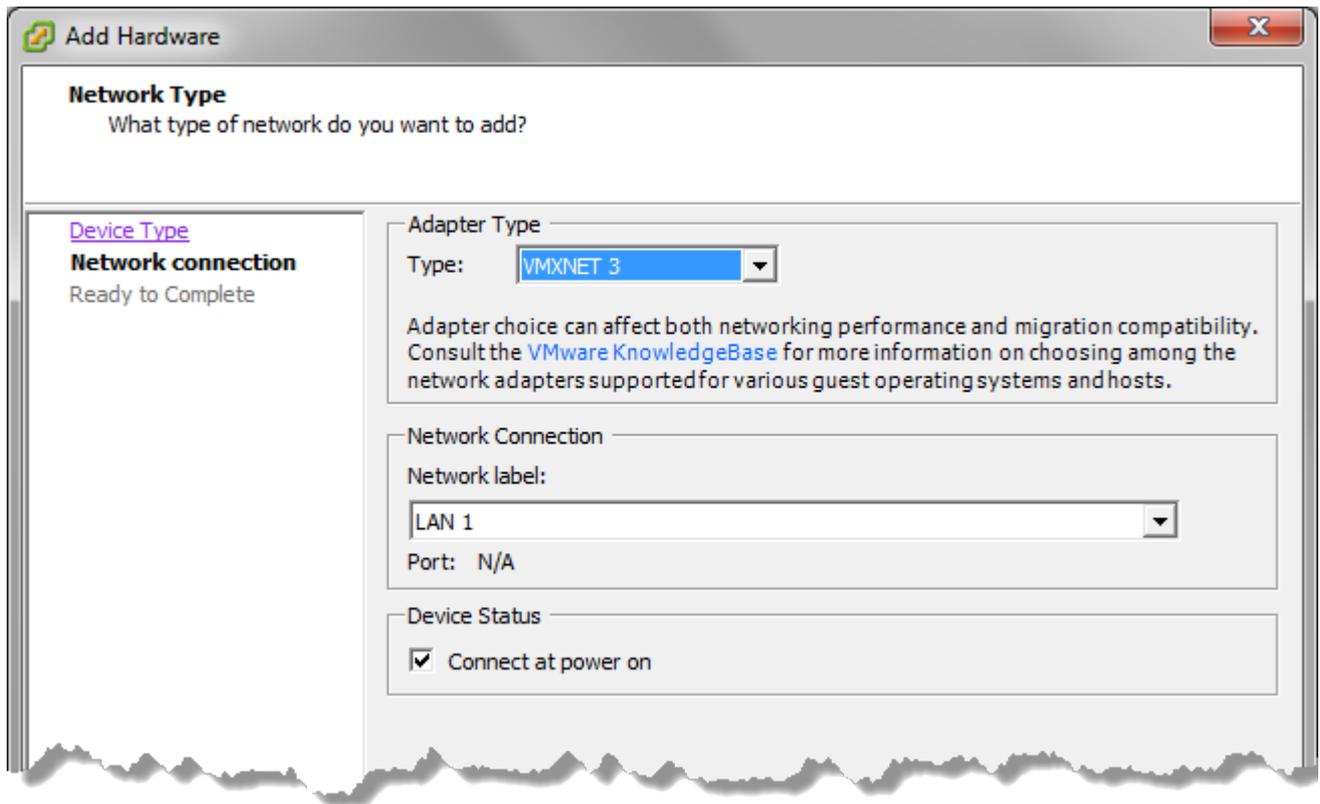
When placing the virtual appliance in line, you can add 2 extra NICs to be used as extra LAN and WAN ports for other circuits. The additional NIC pairs are bridged by default. For more information, refer to [Converting two NICs into a Bridge](#) (page 100)..

The following steps describe how to add extra NICs to the Virtual Appliance. You need to add extra NICs in pairs, in order to create LAN/WAN bridges.

#### NOTE

Even though there is no limitation on the number of bridges a given Virtual Exinda appliance can have, the number of connections can affect the performance of the VM. For more information, refer to [Sizing and resource requirements](#) (page 74). Please ensure that the virtual hardware is appropriate to handle the number of expected connections.

1. Open the **VMware vSphere Client**.
2. Right-click the Exinda Virtual Appliance, and select **Properties**.
3. Switch to the **Hardware** tab.
4. Click **Add**.
5. From the Device Type list, select **Ethernet Adaptor** and click **Next**.
6. In the Adapter Type list, select **VMXNET 3**.
7. Select the network to map the NIC to.



8. Click **Next**.
9. Review the information and click **Finish** to add the NIC.
10. Restart the virtual appliance. The new NICs are automatically detected and any additional NIC pairs are bridged.

### Converting two NICs into a Bridge

Convert the first two NICs into a bridge so the Management Interface becomes a LAN Interface, and the Auxiliary Interface becomes a WAN Interface.

Start the virtual appliance and then...

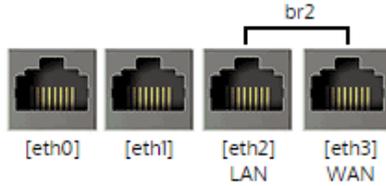
1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Navigate to **Configuration > System > Network > IP Address**.
6. To bridge the two NICs together, select the bridge number you would like to convert and click **Apply Changes**.
7. To manage the Virtual Appliance, in the **IPv4** or **IPv6** field specify an IP Address for the bridge.

# Network Setup

- NICs
- IP Address
- Routes
- DNS
- HTTP Proxy
- Email
- SNMP
- Active Directory

Choose the mode of operation and enter IP Address and Netmask details for that mode. Also specify a Default Route.

**Note:** Take care when making changes here as this appliance may become unreachable.



| Interface Settings                             |                                                                                                                             |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>br0</b> <input type="checkbox"/>            |                                                                                                                             |
| Role:                                          | <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR |
| Autoconf:                                      | IPv4: <input checked="" type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                         |
| Dynamic Addresses:                             | 10.10.6.53/16<br>fe80::20c:29ff:fe0c:47ca/64                                                                                |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| <b>eth0</b>                                    |                                                                                                                             |
| Role:                                          | <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR |
| Autoconf:                                      | IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                    |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| <b>eth1</b>                                    |                                                                                                                             |
| Role:                                          | <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR |
| Autoconf:                                      | IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                    |
| Dynamic Addresses:                             | fe80::20c:29ff:fe0c:47de/64                                                                                                 |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| <b>br2</b> <input checked="" type="checkbox"/> |                                                                                                                             |
| Role:                                          | <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR |
| Autoconf:                                      | IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                    |
| Dynamic Addresses:                             | fe80::20c:29ff:fe0c:47de/64                                                                                                 |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |

**Gateway Settings**

IPv4:

IPv6:

Apply Changes

## NOTE

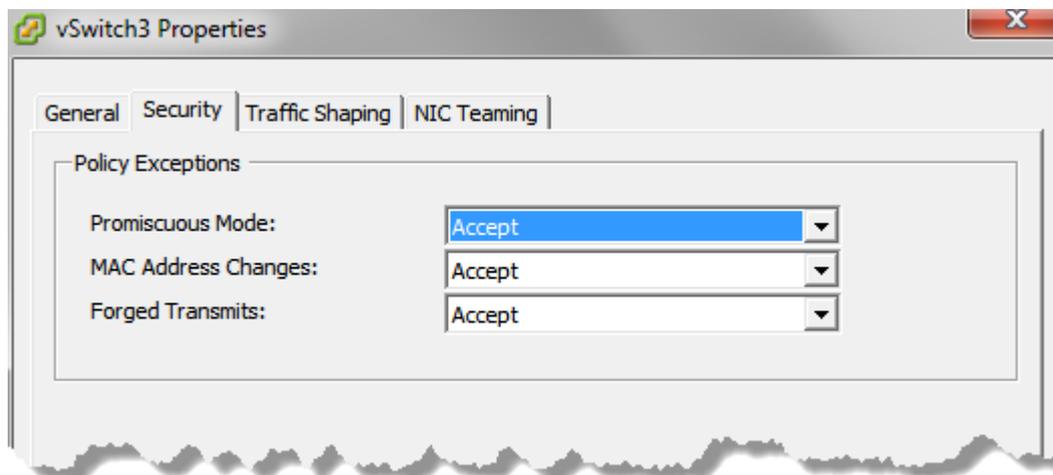
For inline deployments to work correctly under VMware, the virtual switches need to allow promiscuous mode. For more information, refer to [Allow Ports to Accept and Bridge Packets \(Promiscuous Mode\)](#) (page 101).

### Allow Ports to Accept and Bridge Packets (Promiscuous Mode)

Any VMware virtual NIC used to deploy the virtual appliance in line must be configured to allow promiscuous mode, ensuring the LAN and WAN ports are capable of accepting and bridging packets that are not destined for them.

1. Open the **VMware vSphere Client**.
2. Select the ESXi server, and switch to the **Configuration** tab.

3. In the list of Hardware configuration options, select **Networking**.
4. Beside the switch name, click **Properties**.
5. In the switch properties, switch to the **Security** tab.



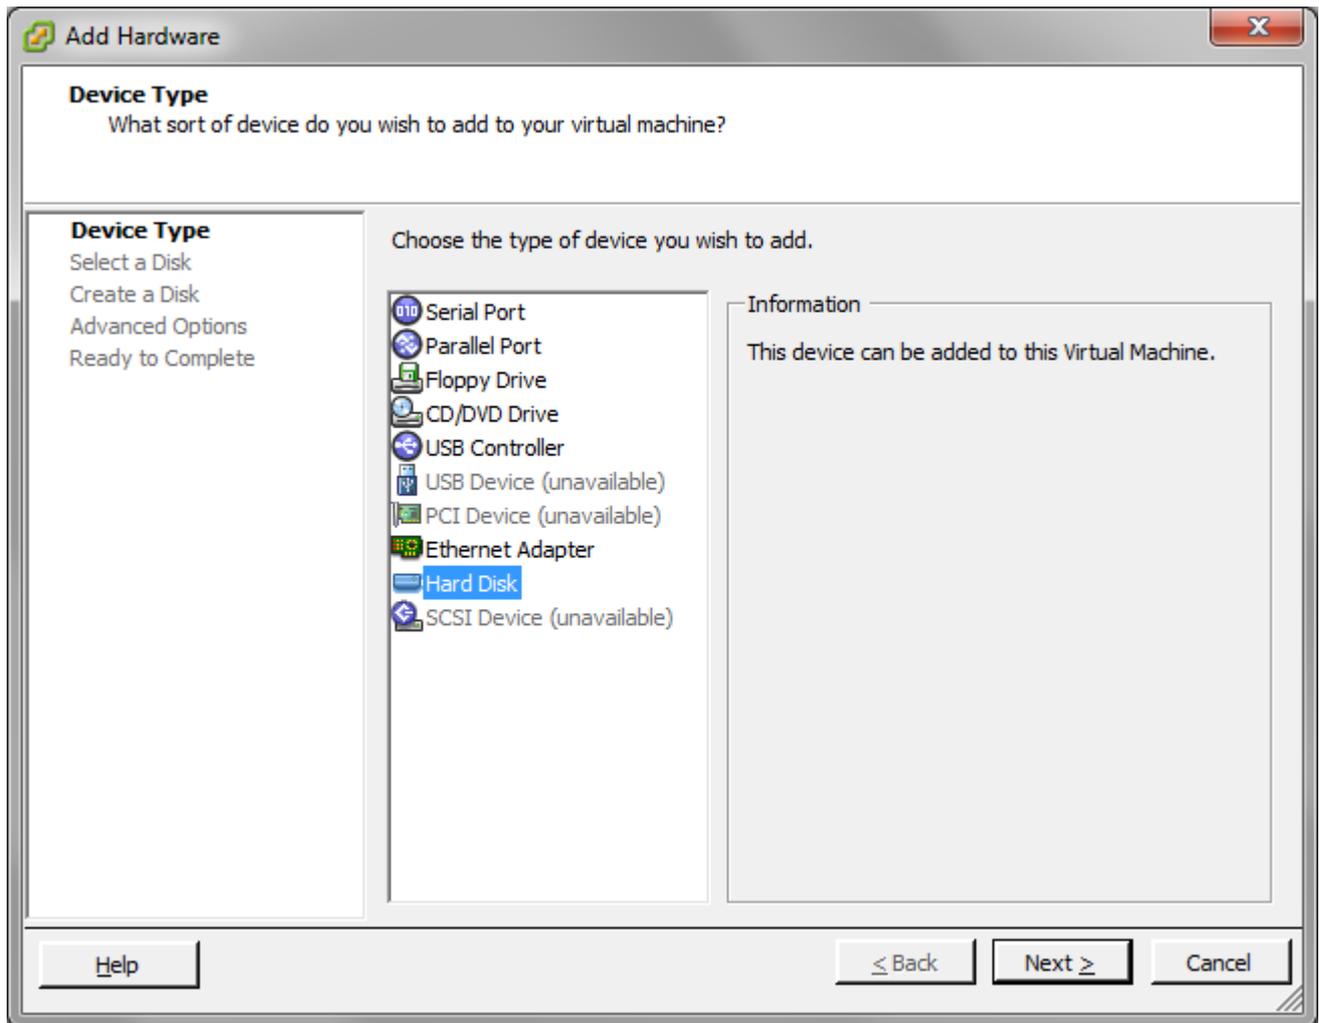
6. Set **Promiscuous Mode** to **Accept**.
7. Click **OK**.
8. Repeat these steps for each virtual switch that is attached to a NIC used in an inline deployment.

### Add Storage to the VMware Virtual Machine

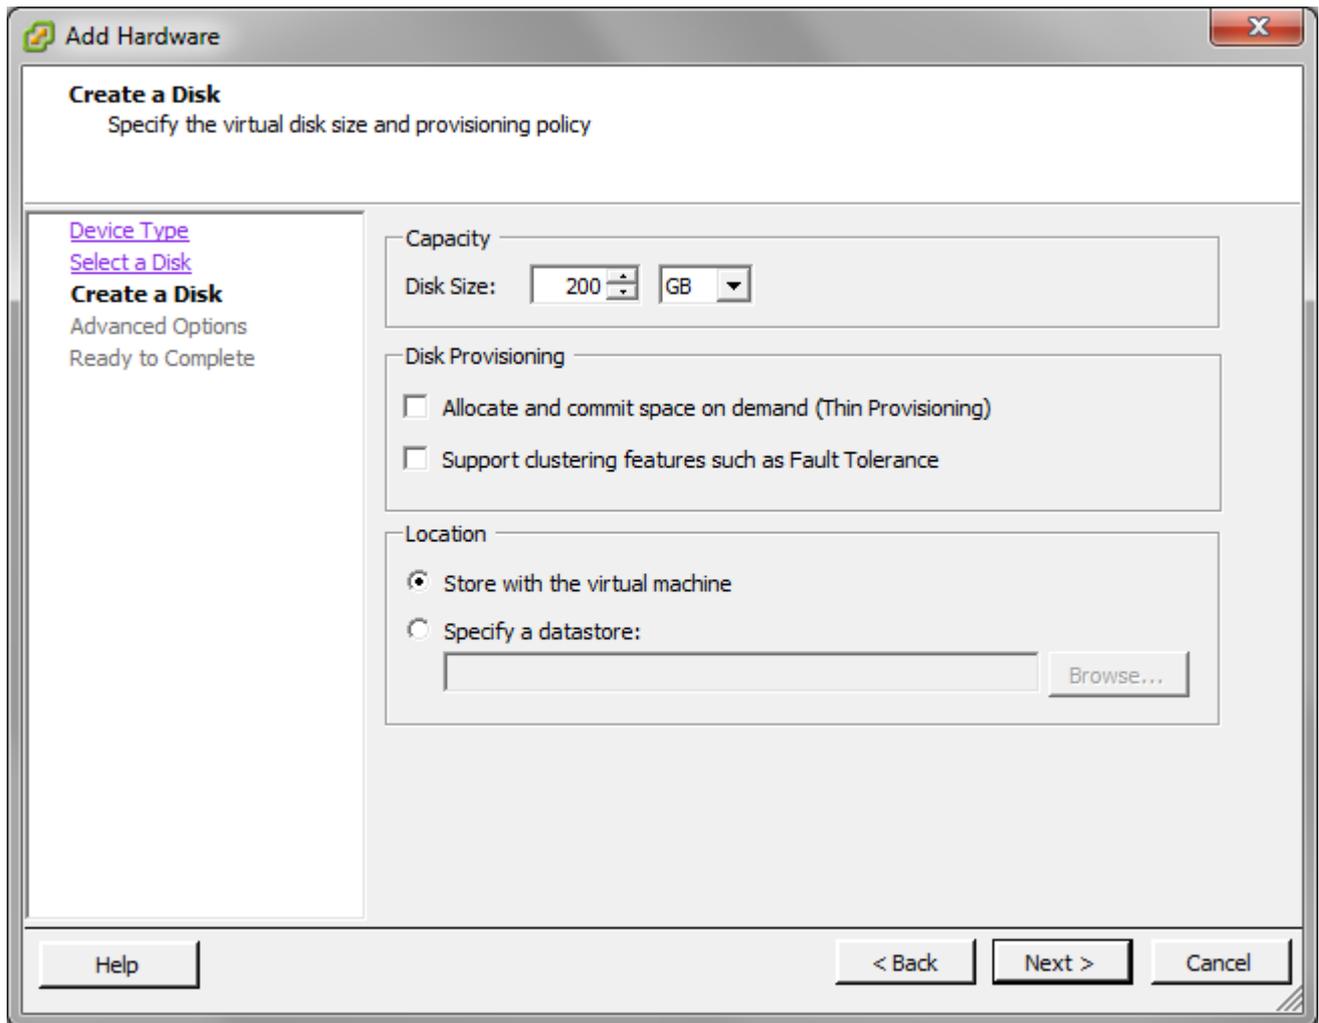
By default, all Exinda Virtual Appliances come with a single 50GB (fixed-size) disk. Usually, you will want more storage for features such as WAN Memory and Edge Cache. This is achieved by adding an additional disk to the Virtual Appliance.

The size of the disk you should add largely depends on the amount of RAM allocated to the Virtual Appliance. As a general rule, you should add a maximum of 100GB of disk storage per 1GB of RAM. So if you have given 4GB of RAM to your Virtual Appliance, you can add up to 400GB of extra storage.

1. Open the **VMware vSphere Client**.
2. From the Hardware tab in the Exinda Virtual Appliance Properties screen, click **Add**.
3. Select **Hard Disk**, then click **Next**.

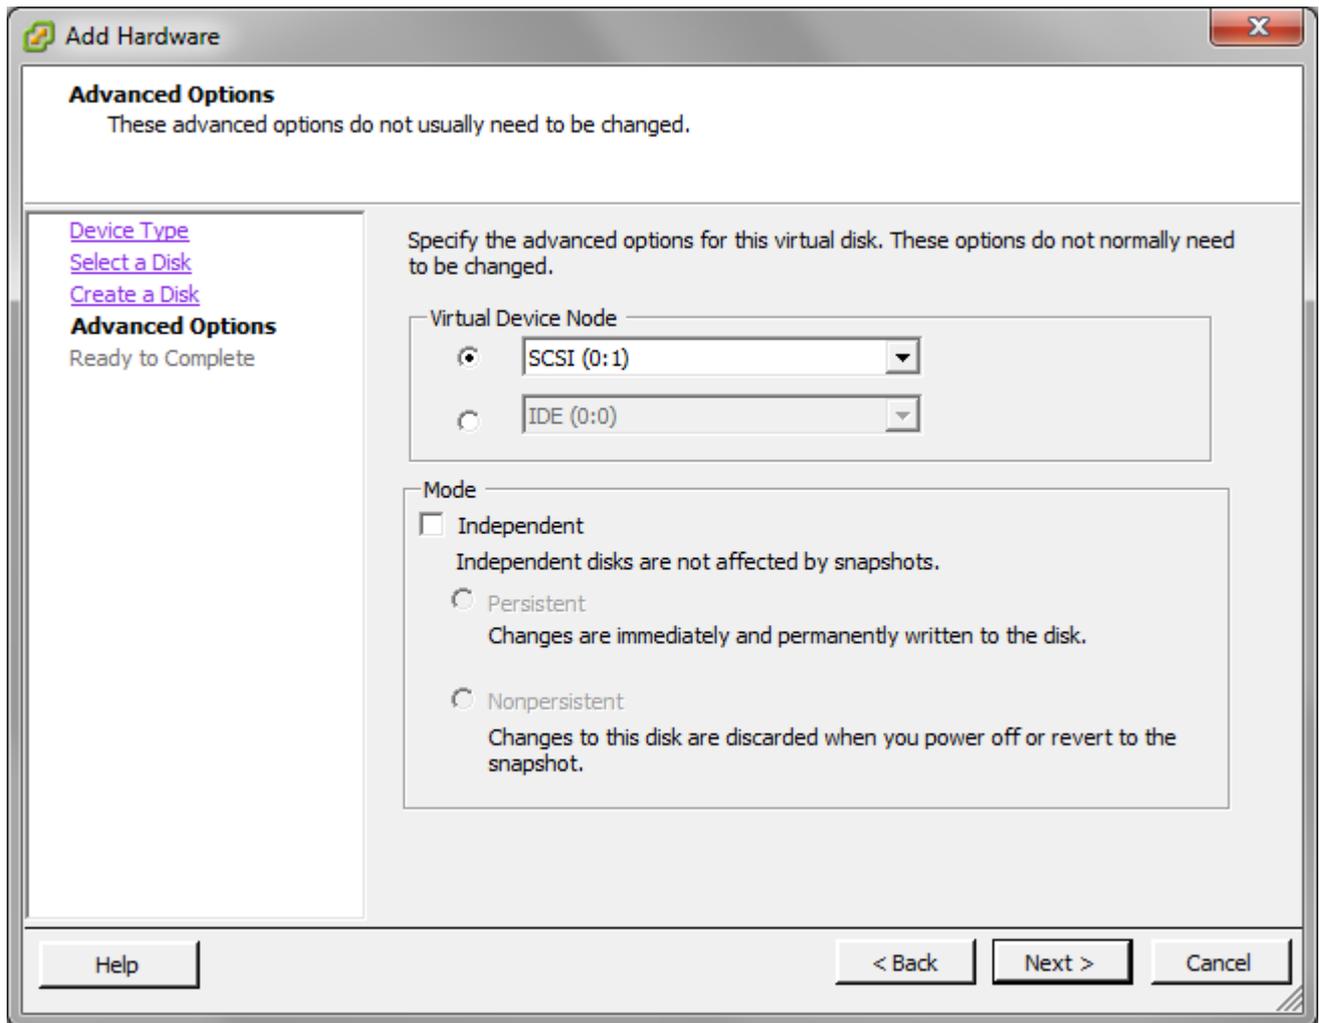


4. Specify the size of the additional disk to create. This space will be added to the default 50GB that comes with the Virtual Appliance. So if you add a 200GB disk here, the total storage for the Virtual Appliance will be 250GB.



5. Click **Next**.

6. Attach the new disk to the next available SCSI node for best performance.



7. Click **Next**.
8. Review the information and click **Finish** to add the disk.
9. When the Virtual Appliance is next booted, you can use the storage commands in the CLI to provision the new storage. The **show storage** command lists the current storage allocations as well as the Virtual Appliance's disks.

```
(config) # show storage
Services:
 cifs: available - 3743.46M free of 3876M total
 edge-cache: available - 3723.53M free of 3872M total
 monitor: available - 9882.83M free of 10G total
 users: available - 974.62M free of 1024M total
 wan-memory: available - 17.21G free of 17.65G total

Disks:
 sda10(internal): in use - 36.22 GB
 sdb: not in use - 214.7 GB

Total: 36.22
Unallocated: 0
```

10. The output shows that our new 200G disk is called 'sdb' and it's currently not in use. **The storage disk add** command is used to provision the new disk.

```
(config) # storage disk add sdb
This will erase all data on the disk. Do you really want to do this (Y/N)? [N] Y
```

11. After this command has executed, another look at **show storage** shows that the new disk is now in use and our 200G is ready for allocation.

```
(config) # show storage
Services:
  cifs: available - 3743.46M free of 3876M total
  edge-cache: available - 3723.53M free of 3872M total
  monitor: available - 9882.83M free of 10G total
  users: available - 974.62M free of 1024M total
  wan-memory: available - 17.21G free of 17.65G total

Disks:
  sda10(internal): in use - 36.22 GB
  sdb: in use - 200.00 GB

Total: 236.21G
Unallocated: 200G
```

## Starting the VMware Virtual Appliance

When you are ready to start the virtual appliance for the first time, Power it on. The Virtual Appliance boots, and displays a login prompt on the VMware console. At this point, you can login with the default username admin and password exinda.

If the first NIC is connected to a network that provides addresses using DHCP, the Virtual Appliance should have picked up an IP address. On the Virtual Appliance summary screen, VMware tools should display the IP address that the Virtual Appliance has obtained.

```
VMware Tools:    Unmanaged
IP Addresses:    192.168.0.221
DNS Name:        exinda-aab541
```

### NOTE

The VMware Tools state 'Unmanaged' is normal. This simply means that VMware Tools are installed and running, but are managed by the guest (the Exinda Virtual Appliance) rather than the host.

If the first NIC is not able to obtain an address using DHCP, you'll need to use the VMware console to enter the following CLI commands to set a static IP address.

```
> en
# conf t
(config) # interface eth0 ip address <ip> <netmask>
(config) # ip default-gateway <default gateway>
(config) # ip name-server <dns server>
```

Once you have determined the IP address or set a static IP address, you can access the web-based user interface by navigating to <https://<ip address>>.

## 0.8.11 Related Topics

At this point, the following tasks should be completed before using the Virtual Appliance:

- » Add extra [NICs](#) (if required) and deploy the Virtual Appliance either in line or out-of-path.
- » Add and provision extra [storage](#) (if required).
- » Obtain a [license](#) for this Virtual Appliance.

### Install the Silicom Bypass Driver on ESXi 5.5 and 6.0

If your ESX/ESXi server has a Silicom network interface card (NIC), you must install the Silicom bypass driver.

1. Enable SSH on your ESX system.
  - a. Enable SSH through the CLI
    - i. In the `/etc/ssh/sshd_config` modify the following variable: `PermitRootLogin yes`
    - ii. Restart the sshd service: `# service sshd restart`
  - b. Enable local or remote TSM from the Direct Console User Interface (DCUI)
    - i. At the DCUI of the ESXi host, press F2 and provide credentials when prompted.
    - ii. Scroll to **Troubleshooting Options**, and press **Enter**.
    - iii. If you want to enable local TSM, select **Local Tech Support** and press **Enter** once. This allows users to login on the virtual console of the ESXi host.
    - iv. If you want to enable remote TSM, select **Remote Tech Support (SSH)** and press **Enter** once. This allows users to login via SSH on the virtual console of the ESXi host.

#### RECOMMENDATION

Have your virtual Exinda already installed with the number of interfaces already set, keep it turned off, we will work with it later.

2. Query the existing VIBs. Make sure you are in maintenance mode: `# vim-cmd /hostsvc/maintenance_mode_enter`

#### NOTE

If the VIB you are deploying exist, you must first remove the existing VIB.

3. Run the following command to determine if any of the existing VIBs match the VIB you are deploying: `# esxcli software vib list | grep bpvm`

#### NOTE

If there are no matches with your VIB, skip the next step.

4. If necessary, remove the existing VIB.

```
# esxcli software vib remove -n net-bpvm
# reboot
```

- Download the Silicom Driver for ESXi 5.1, 5.5 and 6.0 from: [https://updates.exinda.com/exos/virtual/vmware/bypass/5.1-5.5-6.0/net-bpvm-2.0.1.15-1OEM.510.0.0.802205.x86\\_64.vib](https://updates.exinda.com/exos/virtual/vmware/bypass/5.1-5.5-6.0/net-bpvm-2.0.1.15-1OEM.510.0.0.802205.x86_64.vib)
- Copy the driver into the ESX system with SCP or SFTP, drop it on the **/tmp** directory: `# scp net-bpvm-2.0.1.15-1OEM.510.0.0.802205.x86_64.vib root@<esx-serverip>:/tmp`
- Deploy the VIB on the ESX system: `# esxcli software vib install -v /tmp/net-bpvm-2.0.1.15-1OEM.510.0.0.802205.x86_64.vib --no-sig-check`

**NOTE**

Ensure that you specify the full path to the .vib file.

- Reboot the appliance: `# reboot`
- When the ESXi server comes back, verify that a new network adapter named "bpvm0" is listed under **Configuration > Network Adapters**:

| Device                                                               | Speed     | Configured | Switch   | MAC Address       | Observed IP ranges | Wake on LAN |
|----------------------------------------------------------------------|-----------|------------|----------|-------------------|--------------------|-------------|
| <b>bpvm0</b>                                                         | 1000 Full | 1000 Full  | vSwitch4 | 00:e0:ed:18:75:8a | None               | No          |
| <b>Broadcom Corporation Broadcom NetXtreme II BCM5709 1000Base-T</b> |           |            |          |                   |                    |             |
| vmnic3                                                               | Down      | Negotiate  | vSwitch7 | 78:2b:cb:35:7b:2e | None               | Yes         |
| vmnic2                                                               | 100 Full  | Negotiate  | vSwitch5 | 78:2b:cb:35:7b:2e | None               | Yes         |

**NOTE**

Other network interfaces of the silicom card will now show up with the following duplex/speed settings if disconnected:

|               |                   |           |      |                   |      |     |
|---------------|-------------------|-----------|------|-------------------|------|-----|
| <b>vmnic7</b> | <b>65535 Half</b> | Negotiate | None | 00:e0:ed:1e:85:15 | None | Yes |
| <b>vmnic6</b> | <b>65535 Half</b> | Negotiate | None | 00:e0:ed:1e:85:14 | None | Yes |

- Create two standalone vSwitches, and assign the LAN interface of the bridge to one of them and the WAN interface of the bridge to the other. (Configure both standalone switches with Promiscuous Mode and as accepting all VLANs (4065).
- Look around in the **Configuration > Networking** configuration to see if the bpvm0 adapter is already attached to a standalone vSwitch that is not one of the ones created in the previous two steps. If that is the case, simply disconnect that bpvm0 adapter from it and assign it to the vSwitch that is currently connected to the LAN interface of the virtual Exinda by using the following command: `# esxcli network vswitch standard uplink remove -L bpvm0 vSwitch <NUMBER>`

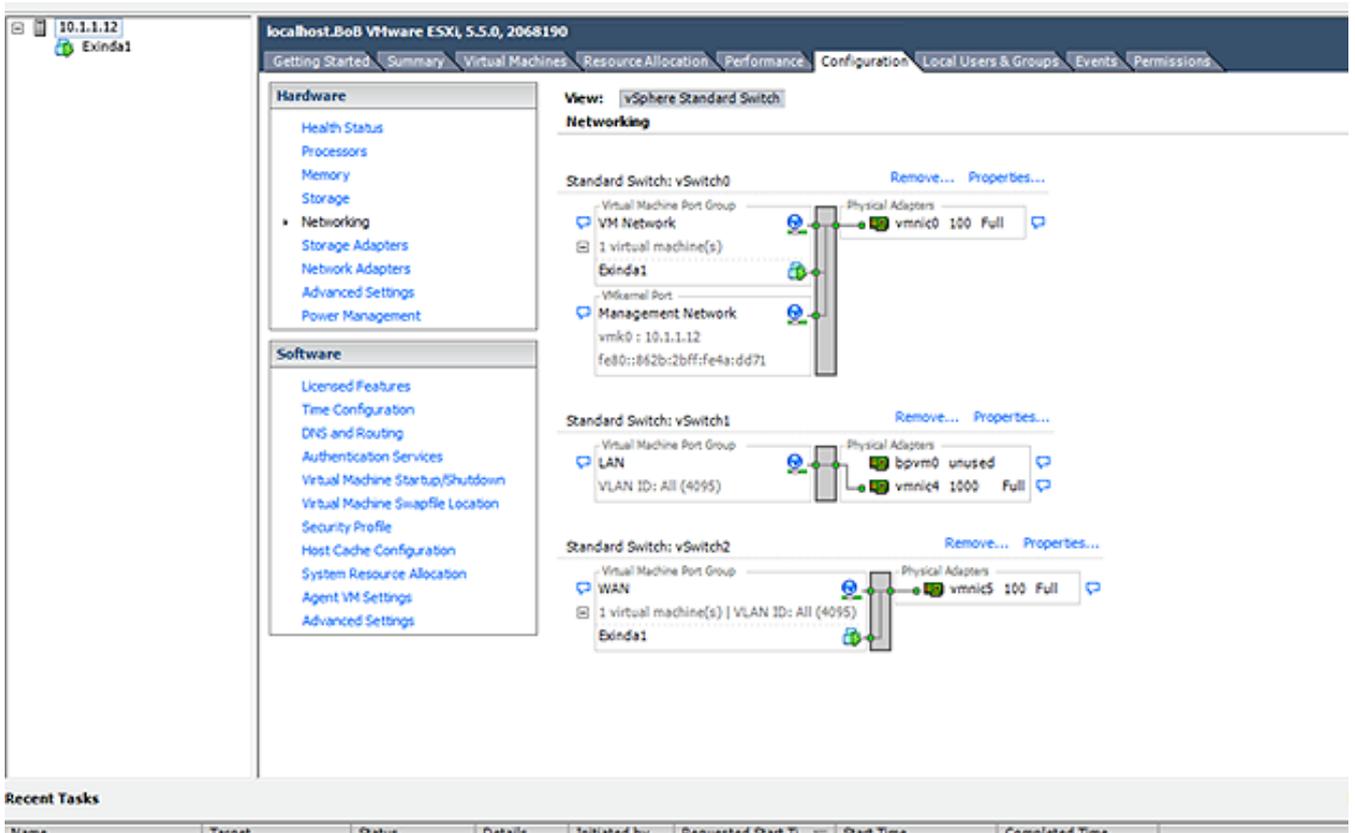
**NOTE**

It is possible that the above command could fail saying that the bpvm0 uplink already exists or that the device is busy. If this happens, reboot the ESXi server one more time with the `reboot` command.

- Connect the physical interfaces of the bridge to its peers (usually to the core switch and to the router/firewall).
- Exit maintenance mode (You can right click the name of the ESXi server for this purpose).
- Turn the Exinda ON

When the Exinda Appliance comes back, you should see the bypass capability available and the duplex/speed negotiations pointing to the right values. One vSwitch should have both the LAN interface of the bridge and the bpvm0 driver while the other vSwitch will have the WAN interface of the bridge, refer to the below pictures (In the below

example, the virtual exinda is configured with 4 interfaces, the first two are standalone interfaces while the last two are for bridging purposes):



#### NOTE

From ESXI v6.0, it is possible that after assigning the `bpxvm0` driver to the “LAN” switch, the driver will not show up as a Physical Adapter (unused) as in the above screenshot. If this is the case, you can continue”

**Network Setup**

**NICs** | IP Address | Routes | DNS | HTTP Proxy | Email | SNMP | Active Directory

Use the form below to set the speed/duplex and MTU of the System NICs. In most cases the default settings will work, however, sometimes explicitly setting speed/duplex is required.

Note: Ensure that the devices connected to the Exinda appliance have the same speed/duplex settings for their network interfaces (autonegotiation is acceptable). If they are different, and the Exinda appliance is in bypass mode, the devices may not communicate. It is recommended you set all your devices, including the Exinda, to either auto-negotiate OR fixed to the same speed/duplex mode.

[View NIC Diagnostics...](#)

| Interface | Media        | HW Address        | Speed | Duplex | MTU  | Link Status                                       |
|-----------|--------------|-------------------|-------|--------|------|---------------------------------------------------|
| eth0      | Twisted Pair | 00:0C:29:16:E3:80 | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 10000Mb/s, Duplex: Full |
| eth1      | Twisted Pair | 00:0C:29:16:E3:8A | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 10000Mb/s, Duplex: Full |
| eth2      | Twisted Pair | 00:0C:29:16:E3:94 | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 10000Mb/s, Duplex: Full |
| eth3      | Twisted Pair | 00:0C:29:16:E3:9E | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 10000Mb/s, Duplex: Full |

Apply Changes

Use the form below to configure bypass state and failover settings.

Note: Take care when making changes here as this appliance may become unreachable.

| Bridge | Status | Running Mode | Enable Failover                     | On Failover |
|--------|--------|--------------|-------------------------------------|-------------|
| br2    | Active | Active       | <input checked="" type="checkbox"/> | Bypass      |

Apply Changes

Use the form below to configure Link State Mirroring. This feature will automatically bring down the second port of a bridge pair if the first port goes down.

**Link State Mirroring**

Link State Mirroring  Enable

Apply Changes

10.1.1.12 - vSphere Client

File Edit View Inventory Administration Plug-ins Help

Home Inventory Inventory

10.1.1.12 Exinda1

Exinda1 - Virtual Machine Properties

Virtual Machine Version: 7

Hardware | Options | Resources

Show All Devices Add... Remove

| Hardware          | Summary      |
|-------------------|--------------|
| Memory            | 4096 MB      |
| CPUs              | 2            |
| Video card        | Video card   |
| VMCI device       | Restricted   |
| SCSI controller 0 | Paravirtual  |
| Hard disk 1       | Virtual Disk |
| Network adapter 1 | VM Network   |
| Network adapter 2 | VM Network   |
| Network adapter 3 | LAN          |
| Network adapter 4 | WAN          |

Memory Configuration

Memory Size: 4 GB

255 GB  
128 GB  
64 GB  
32 GB  
16 GB  
8 GB  
4 GB  
2 GB  
1 GB  
512 MB  
256 MB  
128 MB  
64 MB  
32 MB  
16 MB  
8 MB  
4 MB

Maximum recommended for this guest OS: 255 GB.  
Maximum recommended for best performance: 20472 MB.  
Default recommended for this guest OS: 2 GB.  
Minimum recommended for this guest OS: 512 MB.

Help OK Cancel

Recent Tasks Name, Target or Status

## Monitor IOPS in VMware vSphere

IOPS (Input/output Operations per Second; pronounced "eye-ops") is a common performance measurement used to benchmark computer storage devices like hard disk drives (HDD), solid state drives (SSD), and storage area networks (SAN). As with any benchmark, IOPS numbers published by storage device manufacturers do not guarantee real-world application performance. IOPS are measured in both Commands per Second (IO operations per second) or Throughput (Megabytes per Second).

In the sizing charts for the Exinda virtual appliance (EXN-V) we have represented the measurement in Commands per Second. There are three numbers for IOPS:

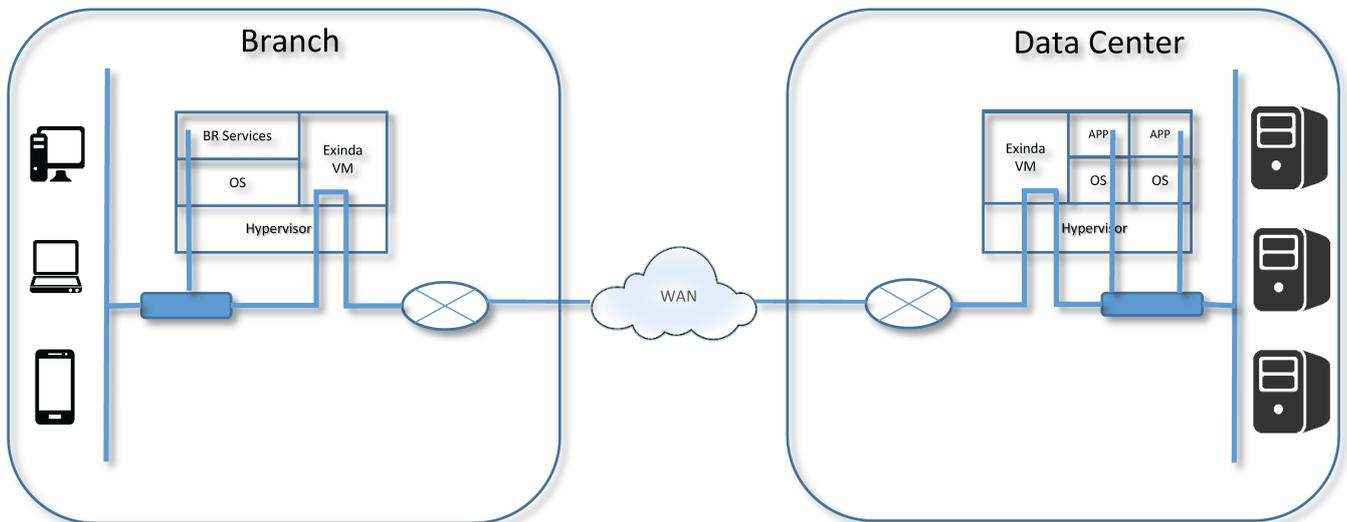
- » EC IOPS for Edge Cache IOPS
- » Monitoring IOPS
- » Average IOPS for Optimization IOPS

The formula to calculate the IOPS for EXN-V you will add the IOPS for each service:

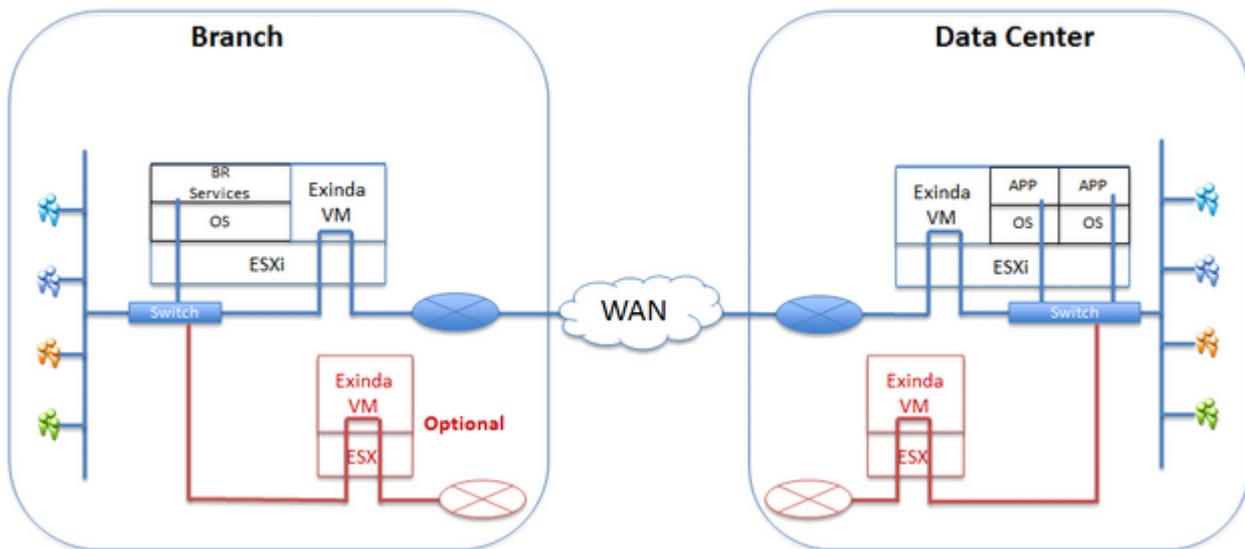
$$\text{Edge Cache IOPS} + \text{Monitoring IOPS} + \text{Average Optimization IOPS} = \text{Total IOPS}$$

| Example: Virtual Model - 2061 | IOPS |
|-------------------------------|------|
| Edge Cache IOPS               | 30   |
| Monitoring IOPS               | 140  |
| Average Optimization IOPS     | 200  |
| Total IOPS                    | 370  |

1. On the Custom Performance Chart for the EXN-V, select **Virtual disk > Real-time**.



2. Select **Average write requests per second (inbound and outbound)**. The report indicates the Minimum, Maximum, and Average Commands per Second.



## Running on Citrix XenServer

Learn how to deploy Exinda Virtual Appliance as well as customize the virtual hardware to suit your requirements. Exinda Virtual Appliances are available for Citrix XenServer hypervisors.

### NOTE

The Exinda Virtual Appliance must be run on either Citrix XenServer 6.2 or 6.5.

## Installing the Virtual Appliance on XenServer

1. Liaise with your local Exinda representative to obtain the download files.
2. Download the Virtual Appliance XVA file.
3. Open your Citrix XenCenter client and select **File > Import...**
4. Select the Virtual Appliance XVA file, and click **Next**.
5. Select the target XenServer to deploy the Virtual Appliance, and click **Next**.
6. Choose the storage location for the Virtual Appliance. By default, the Virtual Appliance comes with a single, 50GB disk. Additional storage can be added in the form of another disk after the Virtual Appliance has been deployed. For more information, refer to [Add storage to the XenServer virtual appliance](#) (page 118). Then click **Next**.
7. Choose the NIC mapping. By default, the Virtual Appliance comes with 4 NICs. The first NIC is the Management Interface, and you should connect it to a network that allows you to manage the Virtual Appliance. The second NIC is an AUX Interface, and is usually used for clustering, high availability or out-of-path deployments. This interface can be left disconnected if not required. In order to fully deploy the Virtual Appliance in line, you may need the additional NICs beyond the four in the configuration. For more information, refer to [Additional NICs](#) (page 115). Then click **Next**.
8. Review the information and clear the 'Start VM(s) after import' box if you want to add extra NICs or storage, and click **Finish** to deploy the Virtual Appliance.
9. Select the Exinda virtual machine you are importing, and switch to the Log tab to see the progress and the completion notification.

## TIP

It is highly recommend that you import the virtual machine on a Gigabit network connection or local storage, as the import file is large in size and installation is affected by slowly performing networks.

10. Right-click on the imported Exinda and select **Start the Exinda virtual appliance**. You will see the progress bar screen below in the Log tab indicating you have successfully started the virtual appliance.

11. On the **XenCenter Console** tab of the Exinda virtual machine, type the credentials and the default parameters as part of the first time wizard setup. The default user name is `admin`, and the password is `exinda`.

12. Press Enter to read the EULA agreement. Press Ctrl-C to get to the EULA agreement question.

13. Press Y to accept the EULA agreement and press **Enter**.

14. You will be prompted with a series of questions as part of the initial configuration Wizard. It is recommended you accept the defaults, as you have the option to configure the system later from the Exinda GUI. Press **Yes**. Use the following defaults to complete the wizard configuration.

- Select **No** to disable IPv6.
- Select **Yes** to configure ETH0 for management access. This will disable the BR0 bridge.
- Select **Yes** to use DHCP on ETH0.
- Select **null** to default to the Exinda hostname.
- Select **null** for SMTP server address.
- Select **null** for email address for reports and alerts.
- Select **null** to use the default password which is "exinda".
- Select **Yes** to change the interface speed.
- Select **AUTO** to configure the interface speed on ETH0 (assumes a gigabit NIC).

You have successfully completed the wizard setup.

15. Determine the IP address of your Exinda virtual appliance on the XenServer Network tab of the Exinda virtual machine and note the IP address assigned by default to NIC 0.

16. Browse to the Dashboard tab and find the Host-ID that the XenServer host created for this virtual machine.

## 0.8.12 Related Topics

Once the appliance is deployed, review the following sections:

- » [Custom Settings](#)
- » [Additional NICs](#)
- » [Add storage to the XenServer virtual appliance](#)

### Modifying the XenServer Virtual Machine Configuration

To improve the performance of the virtual appliance, change the number of CPUs, the RAM, networking, and storage allocated to the virtual machine.

## TIP

You will need to shut the virtual appliance down before you can modify its configuration.

## 0.8.13 Related Topics

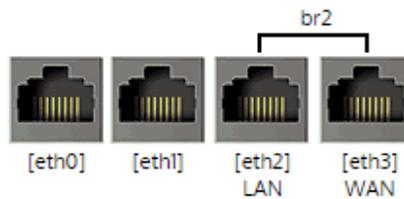
- » Custom Settings
- » Additional NICs
- » Add storage to the XenServer virtual appliance

### Custom Settings

By default, all Exinda virtual appliances come with four network interface cards (NICs):

#### TIP

Before you can make changes to the virtual appliance, you will need to shut it down.



| Interface Settings                             |                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>br0</b> <input type="checkbox"/>            | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                                                                   |
| <b>eth0</b>                                    | Dynamic Addresses: 10.10.1.179/16<br>fe80::20c:29ff:fe2:4d11/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                                         |
| <b>eth1</b>                                    | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/> |
| <b>br2</b> <input checked="" type="checkbox"/> | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                      |
| Gateway Settings                               |                                                                                                                                                                                                                                                                                                                           |
| IPv4:                                          | <input type="text"/>                                                                                                                                                                                                                                                                                                      |
| IPv6:                                          | <input type="text"/>                                                                                                                                                                                                                                                                                                      |

If more interfaces are needed, please follow the next procedure (the Exinda appliance will recognize that if two NICs are added they can then be bridged). The following steps describe how to add extra NICs to the Virtual Appliance. In order to create LAN/WAN bridges, you need to add extra NICs in pairs, .

1. From the **Networking** tab in the Exinda Virtual Appliance settings, click **Add** Interface.
2. Choose the network to which to map this new NIC, then click **Add**.

### Additional NICs

By default, all Exinda Virtual Appliances come with four NICs. The first NIC is the Management Interface (for managing the Virtual Appliance) and the second NIC is the Auxiliary Interface (for use with HA, clustering, and out-of-path deployments).

There are 2 options when it comes to placing the Virtual Appliance in line:

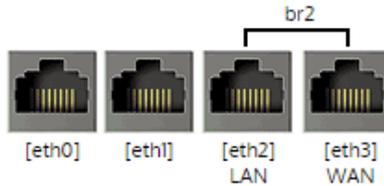
- » Convert the first two NICs into a bridge, so that the Management Interface becomes a LAN Interface and the Auxiliary Interface becomes a WAN Interface. This is achieved by booting into the Virtual Appliance and navigating to the **Configuration > System > Network > IP Address** page on the Web UI, advanced mode. From this page, you can select the **br0** checkbox to bridge the first two NICs together. In order to manage the Virtual Appliance, an IP Address must be specified for this bridge. Stay cautious when using this option as this will cause the first two NICs to be bridged.

## Network Setup

NICs | IP Address | Routes | DNS | HTTP Proxy | Email | SNMP | Active Directory

Choose the mode of operation and enter IP Address and Netmask details for that mode. Also specify a Default Route.

**Note:** Take care when making changes here as this appliance may become unreachable.



| Interface Settings                             |                                                                                                                             |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>br0</b> <input type="checkbox"/>            |                                                                                                                             |
| Role:                                          | <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR |
| Autoconf:                                      | IPv4: <input checked="" type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                         |
| Dynamic Addresses:                             | 10.10.6.53/16<br>fe80::20c:29ff:fe0c:47ca/64                                                                                |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| <b>eth0</b>                                    |                                                                                                                             |
| Role:                                          | <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR |
| Autoconf:                                      | IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                    |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| <b>eth1</b>                                    |                                                                                                                             |
| Autoconf:                                      | IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                    |
| Dynamic Addresses:                             | fe80::20c:29ff:fe0c:47de/64                                                                                                 |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| <b>br2</b> <input checked="" type="checkbox"/> |                                                                                                                             |
| Autoconf:                                      | IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                    |
| Dynamic Addresses:                             | fe80::20c:29ff:fe0c:47de/64                                                                                                 |
| Static Addresses:                              | <input type="text"/> / <input type="text"/>                                                                                 |
| Comment:                                       | <input type="text"/>                                                                                                        |
| Gateway Settings                               |                                                                                                                             |
| IPv4:                                          | <input type="text"/>                                                                                                        |
| IPv6:                                          | <input type="text"/>                                                                                                        |

Apply Changes

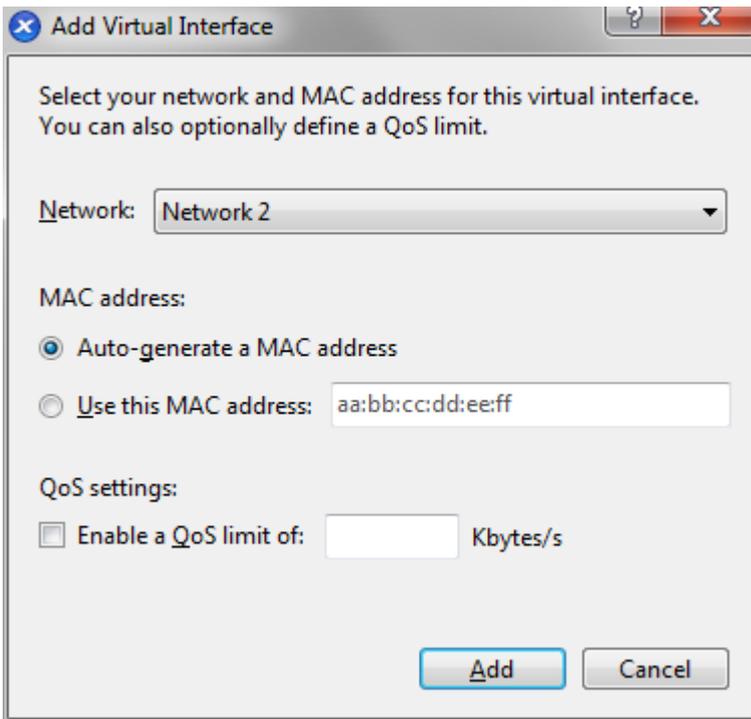
### IMPORTANT

You must power off the virtual appliance while changing the virtual machine configuration.

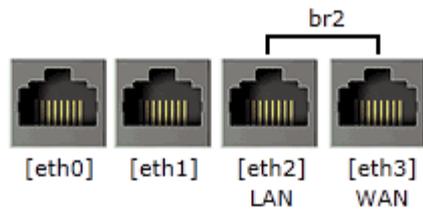
» Use the 2 extra NICs as LAN and WAN ports. These 2 additional NICs can be bridged to allow the Virtual Appliance to be placed in line.

The following steps describe how to add extra NICs to the Virtual Appliance. You need to add extra NICs in pairs, in order to create LAN/WAN bridges.

1. From the Networking tab in the Exinda Virtual Appliance settings, click **Add Interface**.
2. Choose the network to map this new NIC to, then click **Add**.



3. Take care when using this option as this will cause the first two NICs to be bridged.
4. When the Virtual Appliance is next booted, the new NICs will be automatically detected and any additional NIC pairs will be bridged. Below is what the **System -> Network -> IP Address** page on the Web UI looks like after 2 extra NICs have been added.



| Interface Settings                             |                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>br0</b> <input type="checkbox"/>            | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP<br>Autoconf: IPv4: <input checked="" type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: 192.168.0.225/24<br>fe80::10e0:9ff:fe0d:3021/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/> |
| <b>eth0</b>                                    |                                                                                                                                                                                                                                                                                                                                                                               |
| <b>eth1</b>                                    | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::4430:b2ff:fe2a:c4a5/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                               |
| <b>br2</b> <input checked="" type="checkbox"/> | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::2c5d:7fff:fe84:9221/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                       |
| Gateway Settings                               |                                                                                                                                                                                                                                                                                                                                                                               |
| IPv4:                                          | <input type="text"/>                                                                                                                                                                                                                                                                                                                                                          |
| IPv6:                                          | <input type="text"/>                                                                                                                                                                                                                                                                                                                                                          |

### Add storage to the XenServer virtual appliance

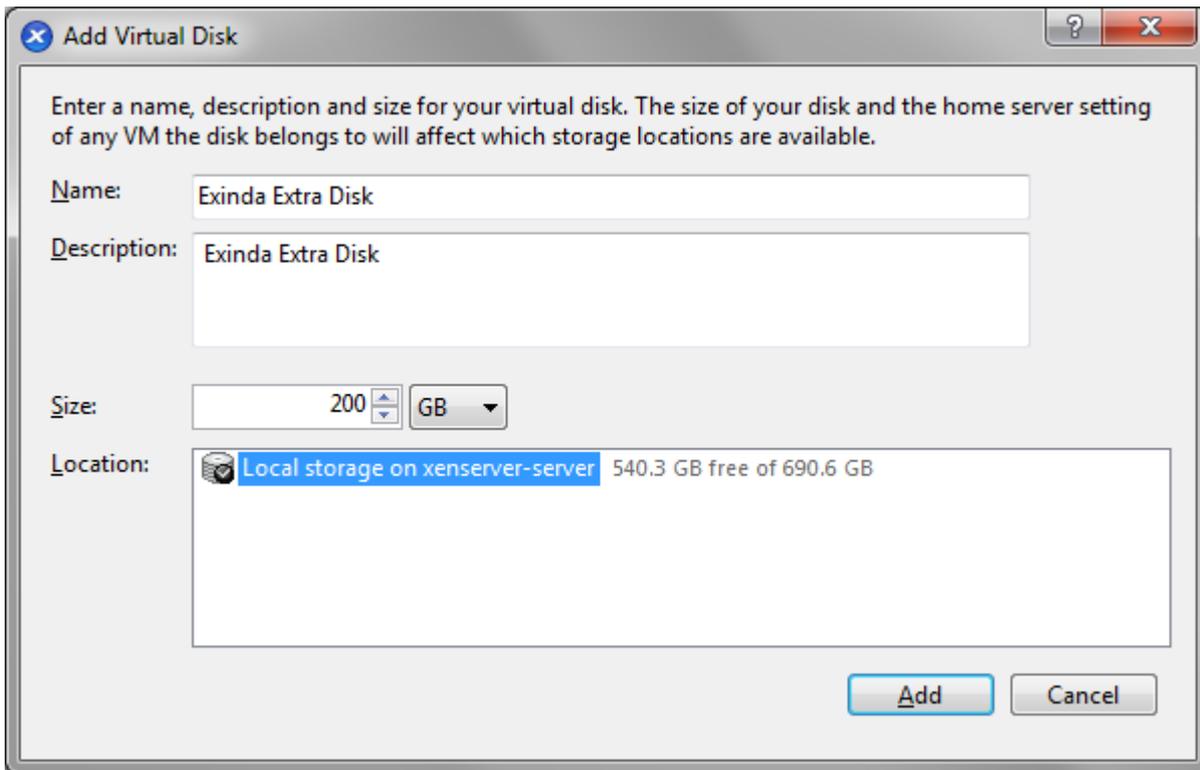
By default, all Exinda Virtual Appliances come with a single 50GB (fixed-size) disk. Usually, you will want more storage for features such as WAN Memory and Edge Cache. This is achieved by adding an additional disk to the Virtual Appliance.

The size of the disk you should add largely depends on the amount of RAM allocated to the Virtual Appliance. As a general rule, you should add a maximum of 100GB of disk storage per 1GB of RAM. So if you have given 4GB of RAM to your Virtual Appliance, you can add up to 400GB of extra storage.

#### IMPORTANT

You must power off the virtual appliance while changing the virtual machine configuration.

1. From the Storage tab in the Exinda Virtual Appliance settings, click **Add**.



2. Specify the size of the additional disk to create. This space will be added to the default 50GB that comes with the Virtual Appliance. So if you add a 200GB disk here, the total storage for the Virtual Appliance will be 250GB.

3. Then click Add. The Virtual Appliance storage should look something like this:

| Position | Name              | Description       | SR                                | Size   | Read Only | Priority    | Active | Device Path |
|----------|-------------------|-------------------|-----------------------------------|--------|-----------|-------------|--------|-------------|
| 0        | Exinda Base Disk  | Exinda Base Disk  | Local storage on xenserver-server | 50 GB  | No        | 0 (Highest) | No     | <unknown>   |
| 1        | Exinda Extra Disk | Exinda Extra Disk | Local storage on xenserver-server | 200 GB | No        | 0 (Highest) | No     | <unknown>   |

4. When the Virtual Appliance is next booted, you can use the storage commands in the CLI to provision the new storage. The **show storage** command lists the current storage allocations as well as the Virtual Appliance's disks.

```
(config) # show storage
Services:
  cifs: available - 3743.46M free of 3876M total
  edge-cache: available - 3723.53M free of 3872M total
  monitor: available - 9882.83M free of 10G total
  users: available - 974.62M free of 1024M total
  wan-memory: available - 17.21G free of 17.65G total

Disks:
  xvda10(internal): in use - 36.22 GB
  xvdb: not in use - 214.7 GB

Total: 36.22
Unallocated: 0
```

5. The output shows that our new 200G disk is called 'xvdb' and it's currently not in use. The **storage disk add** command is used to provision the new disk.

```
(config) # storage disk add xvdb
This will erase all data on the disk. Do you really want to do this (Y/N)? [N] Y
```

6. After this command has executed, another look at **show storage** shows that the new disk is now in use and our 200G is ready for allocation.

```
(config) # show storage
Services:
  cifs: available - 3743.46M free of 3876M total
  edge-cache: available - 3723.53M free of 3872M total
  monitor: available - 9882.83M free of 10G total
  users: available - 974.62M free of 1024M total
  wan-memory: available - 17.21G free of 17.65G total

Disks:
  xvda10(internal): in use - 36.22 GB
  xvdb: in use - 200.00 GB

Total: 236.21G
Unallocated: 200G
```

For more information on adding disks in general and allocating storage, see [Configuring Storage](#).

## Starting the XenServer Virtual Appliance

When you are ready to boot the Virtual Appliance for the first time, Power it on. The Virtual Appliance will boot, and when ready, will display a login prompt on the XenCenter console.

At this point, you can login with the default username **admin** and password **Exinda**.

If the first NIC is connected to a network that provides addresses using DHCP, the Virtual Appliance should have picked up an IP address. On the Virtual Appliance Networking screen, XenCenter should display the IP address that the Virtual Appliance has obtained.

### Networks

| Device | MAC               | Limit | Network   | IP Address    | Active |
|--------|-------------------|-------|-----------|---------------|--------|
| 0      | 12:e0:09:cd:30:21 |       | Network 0 | 192.168.0.225 | Yes    |
| 1      | 46:30:b2:2a:c4:a5 |       | Network 1 | Unknown       | Yes    |

If the first NIC is not able to obtain an address using DHCP, you'll need to use the XenCenter console to enter the following CLI commands to set a static IP address.

```
> en
# con t
(config) # interface eth0 ip address <ip> <netmask>
(config) # ip default-gateway <default gateway>
(config) # ip name-server <dns server>
```

Once you have determined the IP address or set a static IP address, you can access the web-based user interface by navigating to `https://<ip address>`.

## 0.8.14 Related Topics

At this point, the following tasks should be completed before using the Virtual Appliance:

- » Obtain a license for this Virtual Appliance.
- » Add and provision extra [storage](#) (if required).
- » Add extra [NICs](#) (if required) and deploy the Virtual Appliance either in line or out-of-path.

### Running on Microsoft Hyper-V

The following sections describe how to deploy Exinda Virtual Appliance as well as to customize the virtual hardware to suit your requirements.

The Exinda Virtual Appliance are available for Microsoft Hyper-V hypervisors.

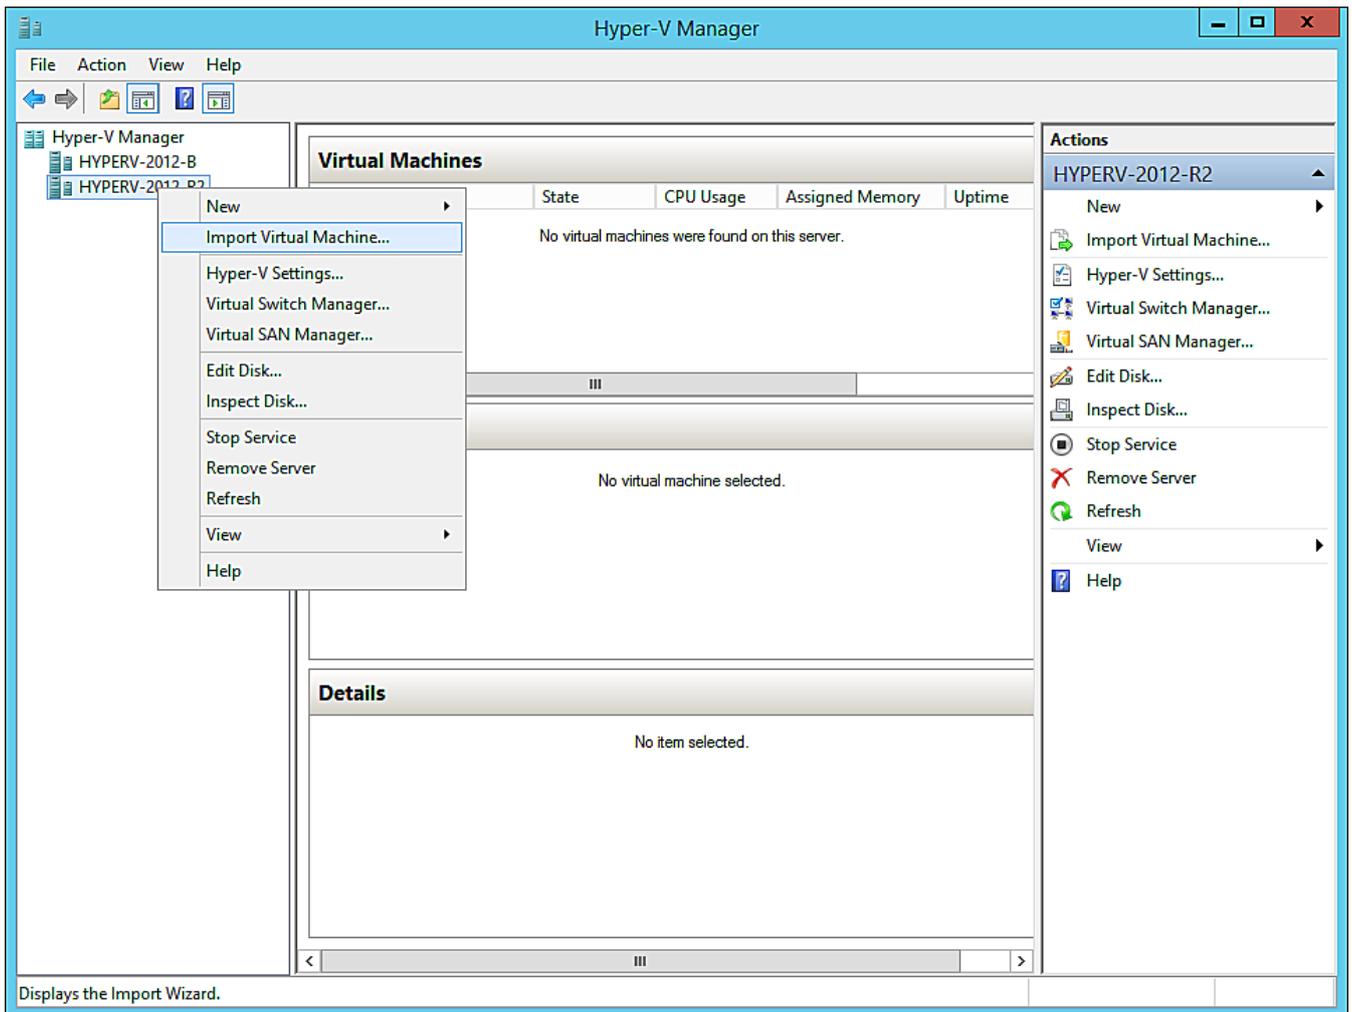
---

---

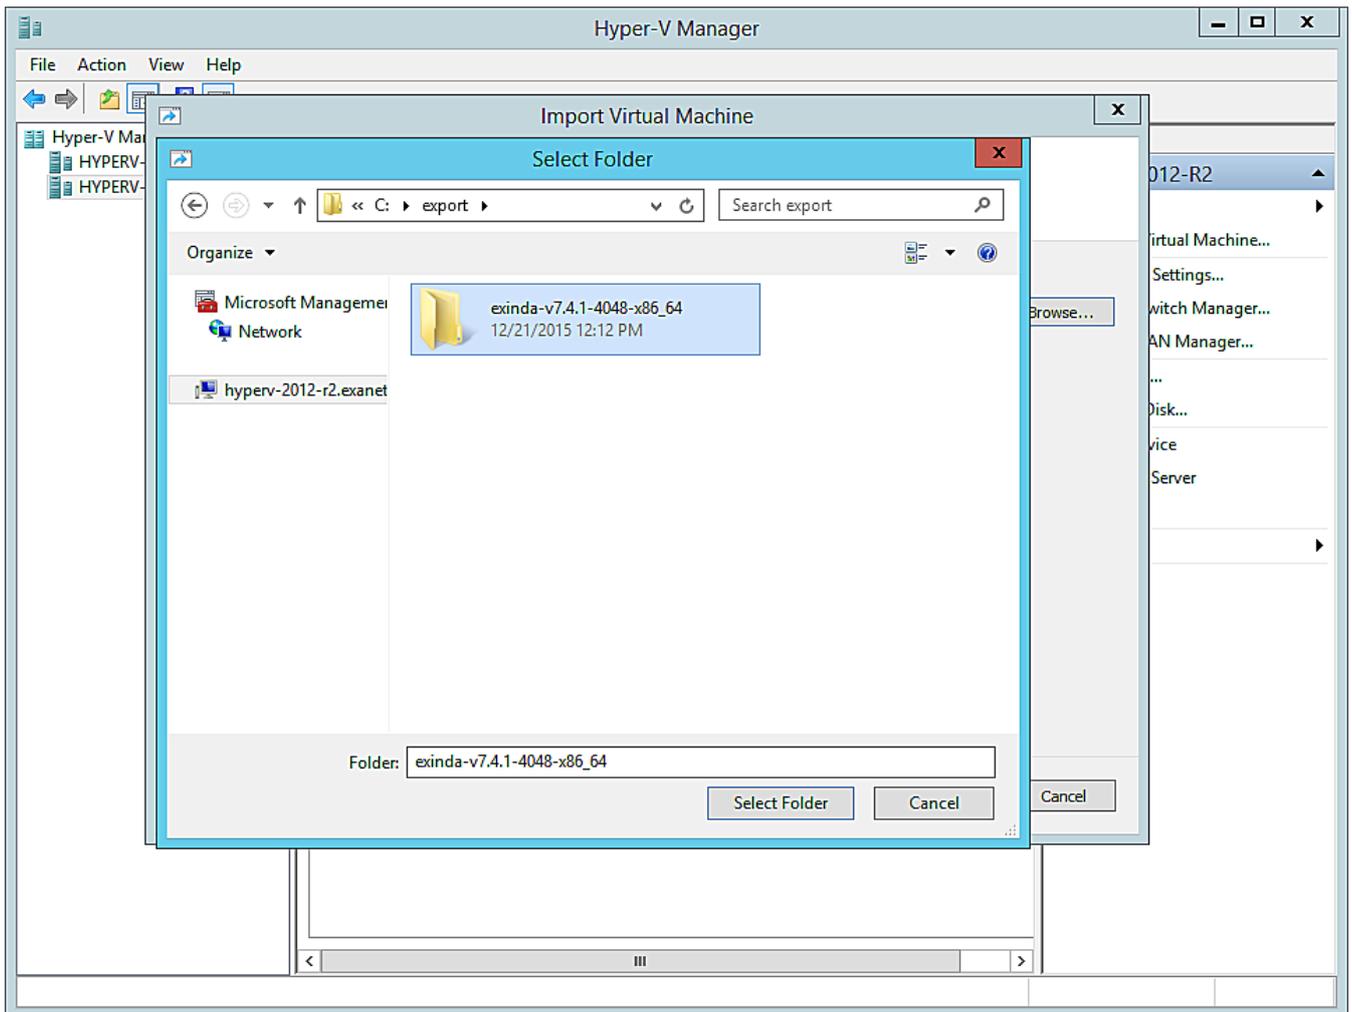
#### Install the Virtual Machine on Hyper-V

The Exinda Virtual Appliances have been prepared to run in a variety of virtual environments. Hyper-V provides support for hosting the Exinda Virtual Appliances in Microsoft Server 2012 and 2012 R2.

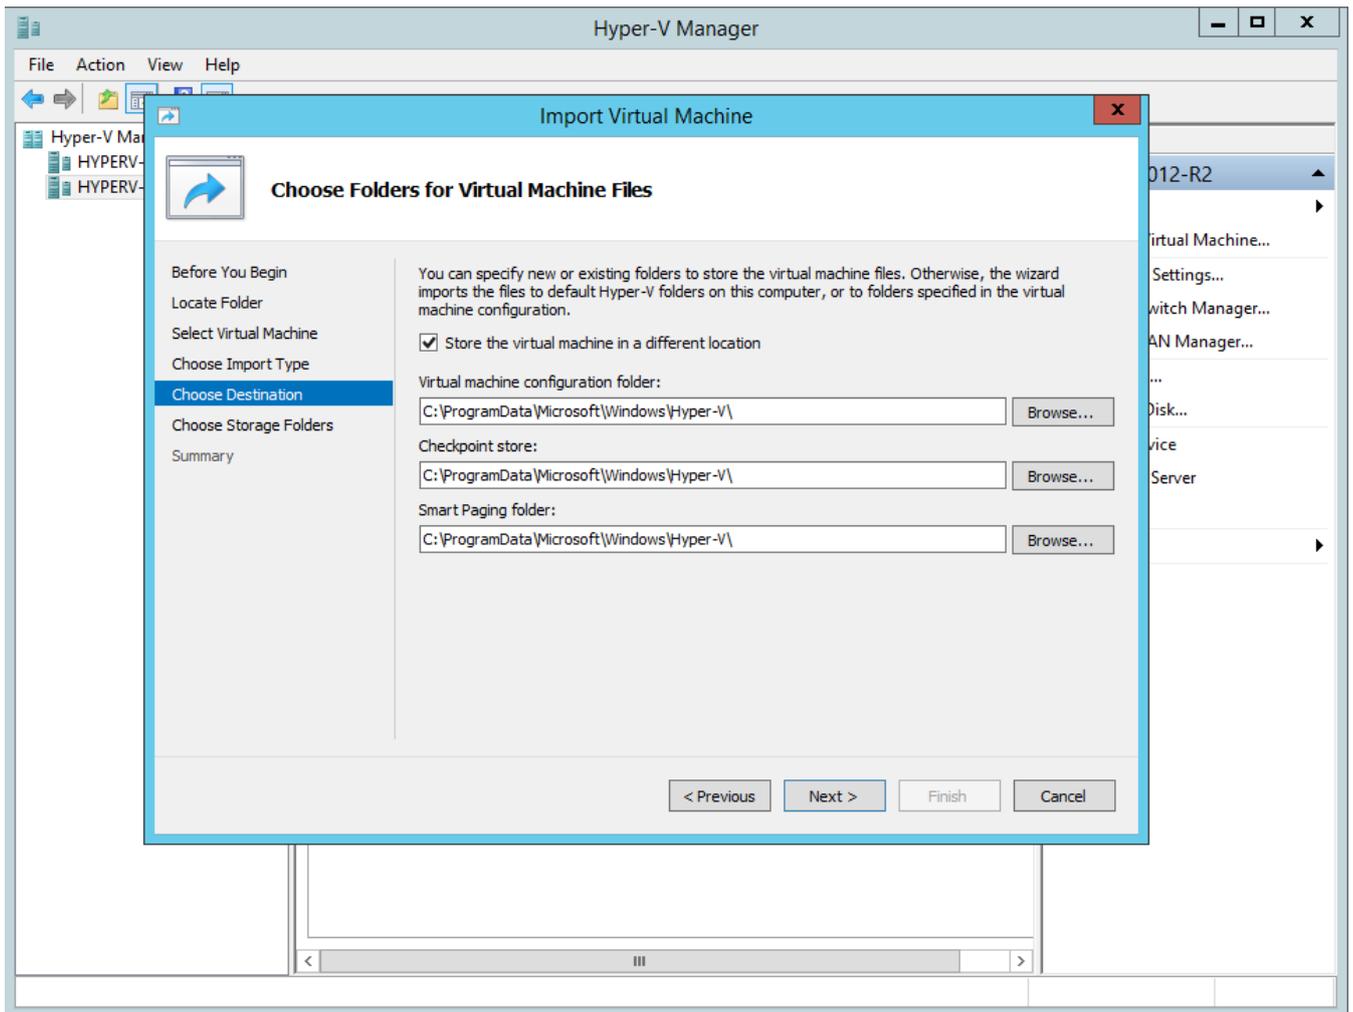
1. Liaise with your local Exinda representative to obtain the download files.
2. Unzip the ZIP file into a local folder. The ZIP archive contains three folders, which contain the following:
  - Snapshots – this folder is empty.
  - Virtual Hard Disks – contains the virtual machine file in VHDX format.
  - Virtual Machines – contains an XML file that in turn contains the configuration of the virtual machine.
3. Open the Hyper-V Manager.
4. In the left pane, right-click on the host machine and select **Import Virtual Machine**. The "Import Virtual Machine" wizard opens.



5. Click the **Browse** button and navigate to and select the local folder where you unzipped the downloaded file. The wizard then recognizes the virtual machine.

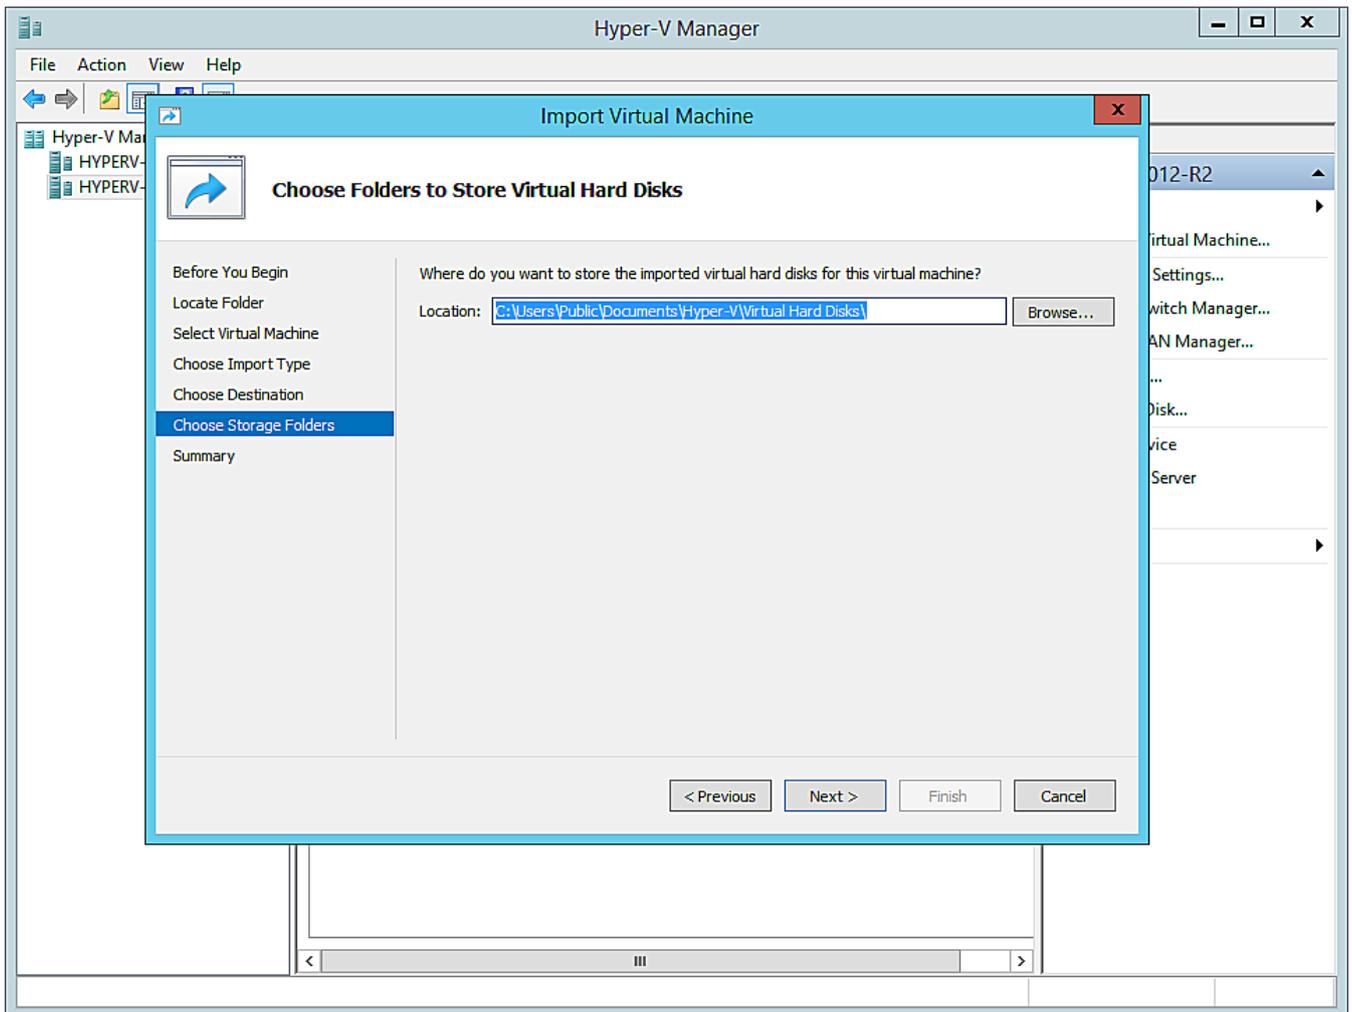


6. On the **Import Virtual Machine** dialog box, click **Next**. The "Choose Import Type" page of the wizard opens.
7. Select the **Copy the virtual machine...** radio button. The "Choose Folder for Virtual Machine Files" page of the wizard opens.
8. If you prefer not to use the default folders, select the **Store the virtual machine...** checkbox and for each of the three folder options browse to and select your preferred folder.



9. Click **Next**. The "Choose Folders to Store Virtual Disks" page of the wizard opens.

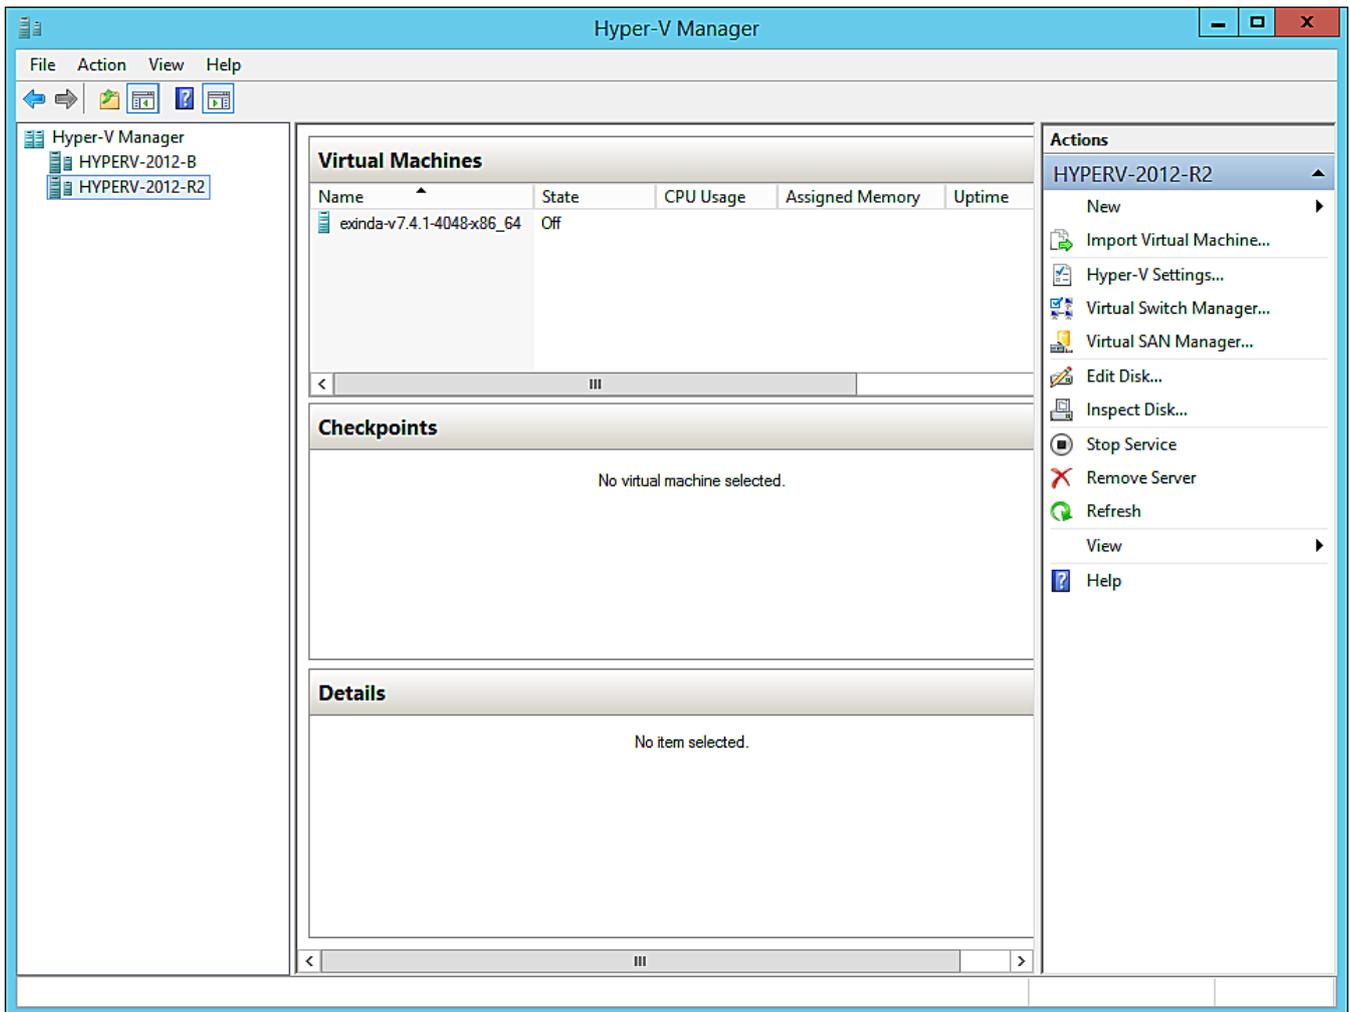
10. Click the **Browse** button and select the folder to use.



11. Click **Next**. The "Completing Impoport Wizard" page opens.

12. Review the settings in the right pane. If they are correct, click **Finish**. The installation proceeds.

When the installation is complete, an entry for the new virtual machine appears in the **Virtual Machines** pane in the **Hyper-V Manager**.



The Virtual Machine, as supplied by Exinda, may not have all of the configuration options you prefer. For example, the disk storage is confined to 50GB, which is unlikely to be sufficient for your needs. When preparing the Exinda Virtual Appliance for download, it is not possible to know just what hardware is available on the host machine. After you have installed the virtual machine, you will need to make some adjustments to the configuration using the controls in the Hyper-V Manager. See the following related tasks.

### IMPORTANT

Before powering on your Exinda Virtual Appliance for the first time, you need to make sure that the virtual configuration is what you need. See the following Related Tasks to fully configure your VM.

## 0.8.15 Related Topics

- » [Adjusting the number of CPUs available to the Virtual Machine](#)
- » [Adjusting the RAM available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Increasing the size of the virtual hard drive](#)

### Modifying the Hyper-V Virtual Machine Configuration

As supplied, the Exinda Virtual Appliances will require some configuration changes before you introduce them to your network. For example, the virtual hard drives are limited to 50GB, which would be unlikely to be sufficient for your needs.

The virtual machines available are sized with minimal configuration as it is not possible to know just what hardware is available on any host machine. To edit the configuration, you need to open the settings for the virtual machine in the Hyper-V Manager.

There are many settings that you can change, but for the purposes of configuring the Exinda Virtual Appliance, these task instructions are limited to what is necessary for bringing the appliance into an operational state. If you need more information, please consult the documentation for Hyper-V. This topic deals with changes to the configuration related to the number of CPUs, the available RAM, the NICs, and adjusting the storage for the virtual machine.

The configuration changes are required before your initial use of the virtual machine. You can also make further changes to your virtual machine at any later time. If over time you require more resources for the virtual machine, as long as those resources are available on the host, you can make them available to the guest.

#### **NOTE**

You will need to shut the virtual appliance down before you can modify its configuration.

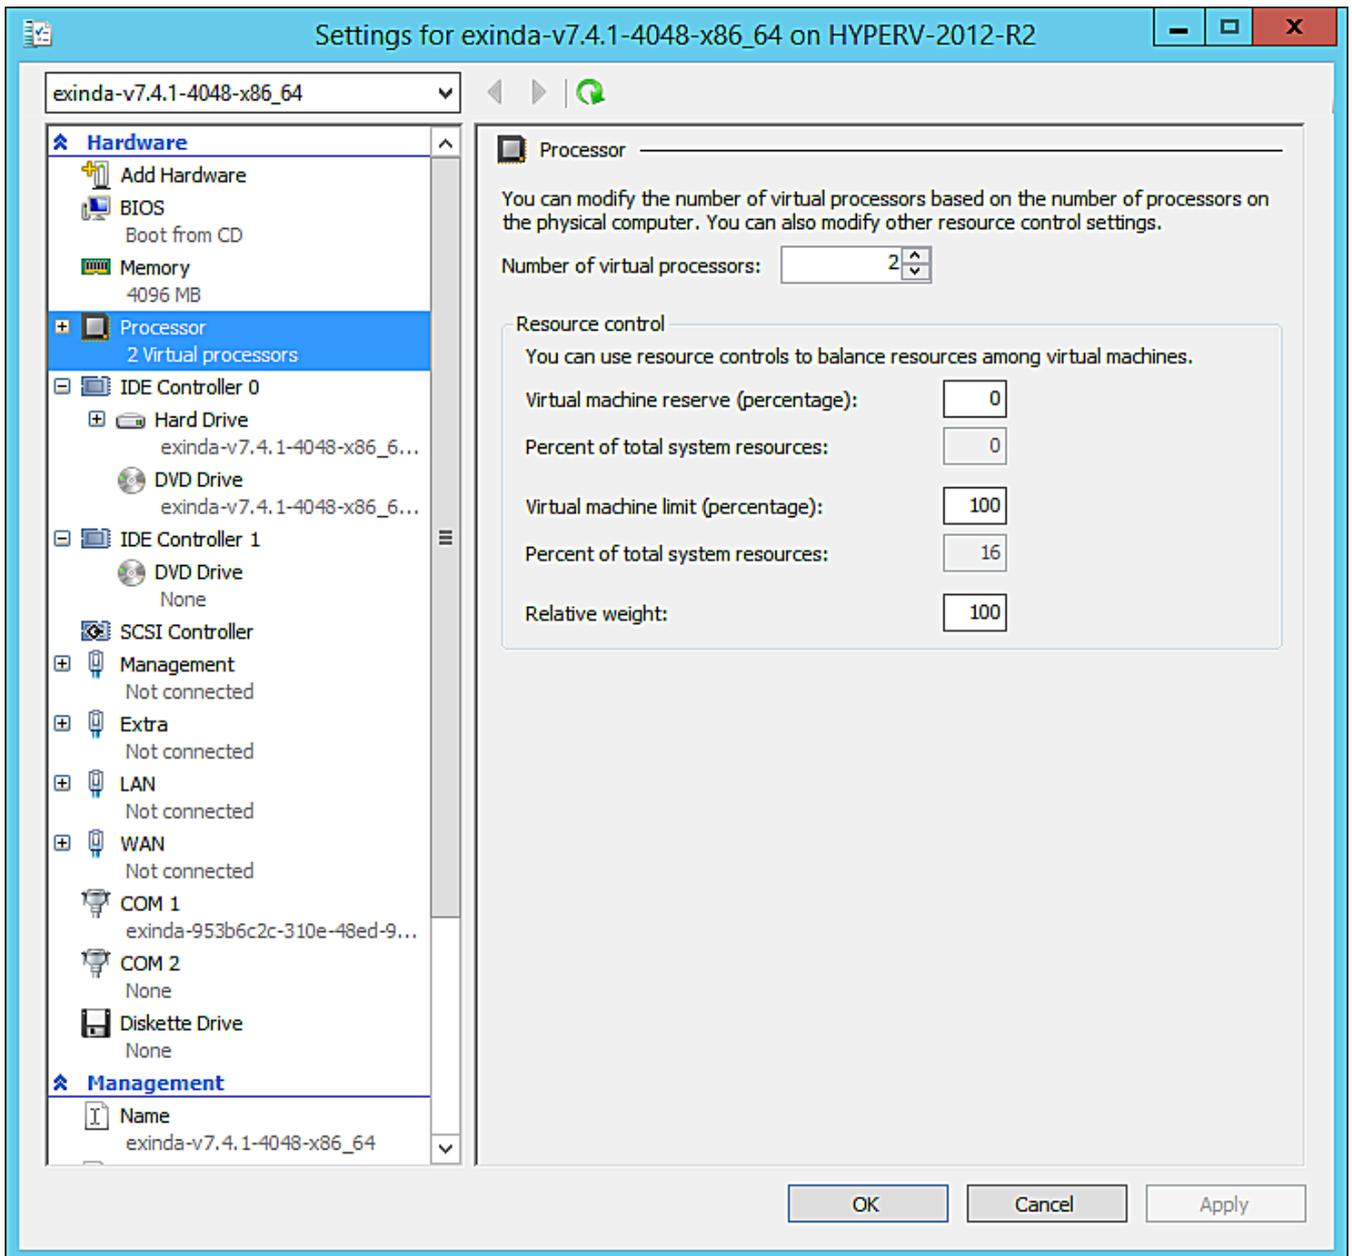
### **0.8.16 Related Topics**

- » [Adjusting the number of CPUs available to the Virtual Machine](#)
- » [Adjusting the RAM available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Increase storage by adding new virtual drives](#)

#### [Adjusting the number of CPUs available to the Virtual Machine](#)

After installing the virtual machine, you may need to adjust the number of CPUs that are available to the Exinda Virtual Appliance. The basic virtual machine configuration includes a minimal number of CPUs, but if you have spare CPUs on the host machine, you may want to make these available to the virtual machine. You make adjustments to the number of CPUs in the Hyper-V Manager.

1. Open the Hyper-V Manager.
2. In the left pane, right-click on the virtual machine you need to edit and select **Settings**. The Settings dialog box for the virtual machine opens.
3. In the left pane, under **Hardware**, select the **Processor** item. The processor settings open in the right pane.



4. In the **Number of virtual processors** spinbox, click the up- or down-arrows to adjust the number of CPUs.

#### NOTE

In this pane you can also adjust several other settings to balance resources among any other virtual machines. Consult the Hyper-V documentation for more information on these settings.

5. Click **OK**. The number of CPUs available to the virtual machine is immediately adjusted.

#### NOTE

These instructions also apply to changing the configuration after the virtual appliance has entered service.

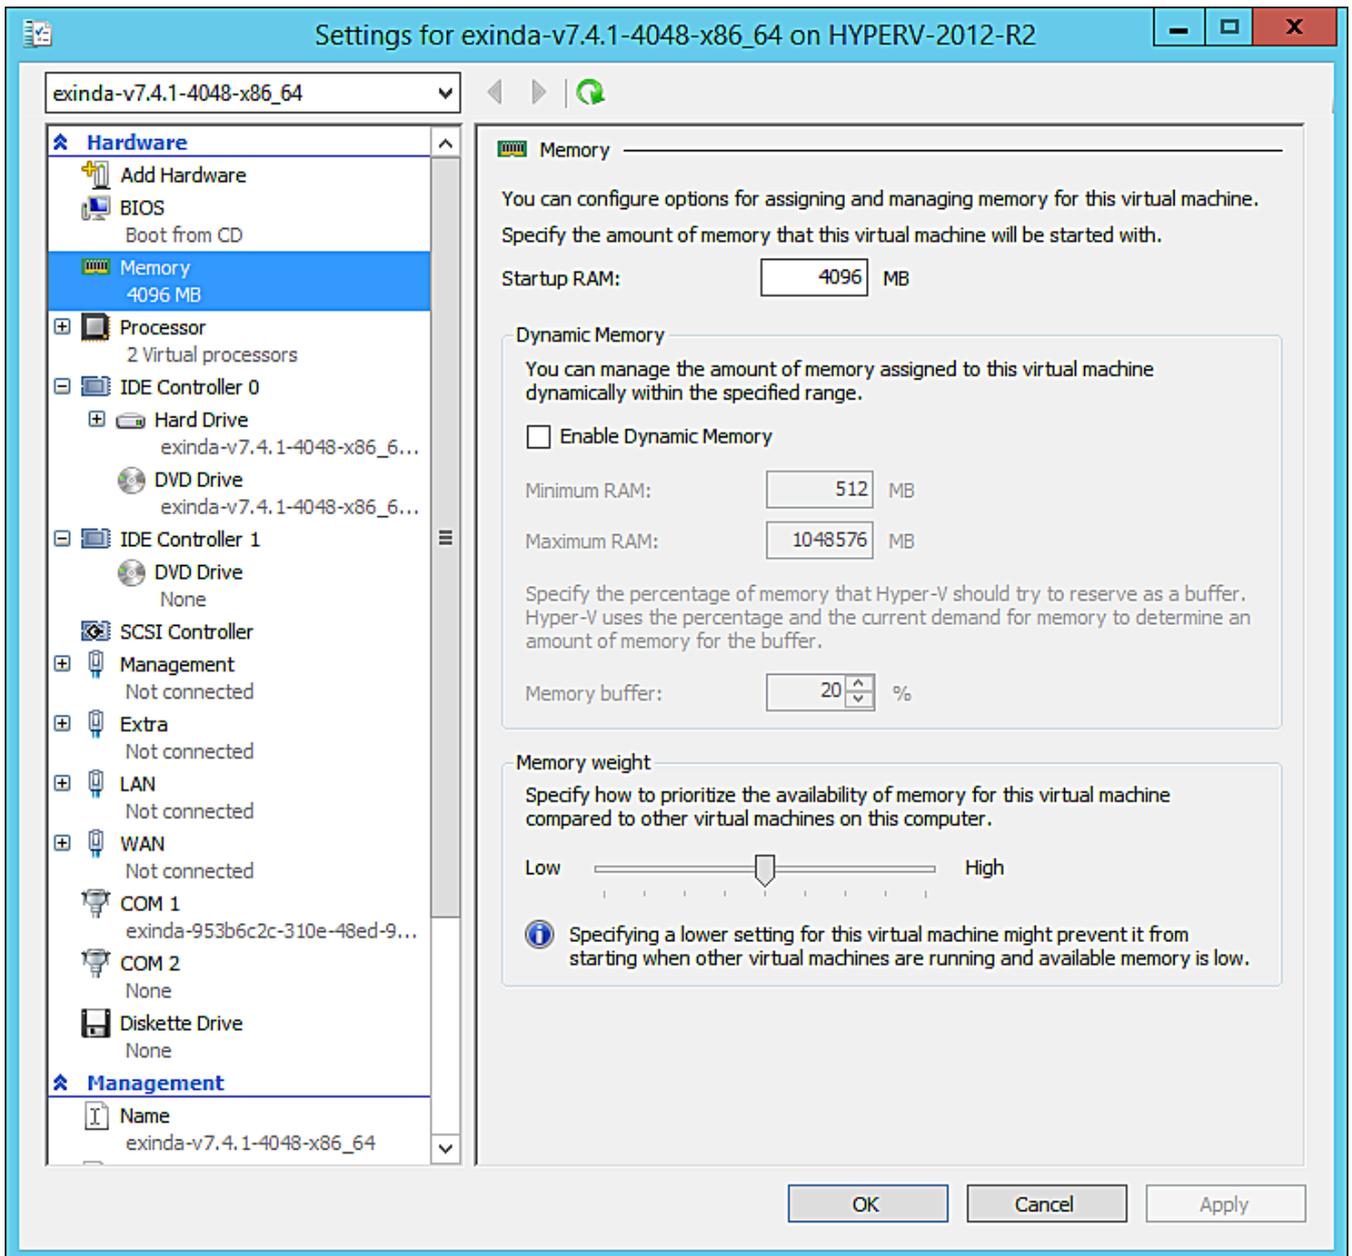
## 0.8.17 Related Topics

- » [Adjusting the RAM available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Increasing the size of the virtual hard drive](#)

### Adjusting the RAM available to the Virtual Machine

After installing the virtual machine, you may need to adjust the amount of RAM that is available to the Exinda Virtual Appliance. There is a basic amount of RAM provided in the Exinda Virtual Appliance, but if you have spare RAM on the host machine, you may want to make this available to the virtual machine. You make adjustments to the amount of RAM in the Hyper-V Manager.

1. Open the Hyper-V Manager.
2. In the left pane, right-click on the virtual machine you need to edit and select **Settings**. The Settings dialog box for the virtual machine opens.
3. In the left pane, under **Hardware**, select the **Memory** item. The memory settings open in the right pane.



4. In the **Startup RAM** field, type a new amount for the quantity of RAM.

**TIP**

In this pane you can also Enable Dynamic Memory and specify amounts, and adjust Memory weight. Consult the Hyper-V documentation for more information on these settings.

5. Click **OK**. The amount of RAM available to the virtual machine is immediately adjusted.

**NOTE**

These instructions also apply to changing the configuration after the virtual appliance has entered service.

## 0.8.18 Related Topics

- » [Adjusting the number of CPUs available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Increasing the size of the virtual hard drive](#)

### Adjusting the NICs available to the Virtual Machine

After installing the virtual machine, you will need to enable the NICs that are available to the Exinda Virtual Appliance. The basic virtual machine configuration does not include the configuration necessary to link the virtual devices to the physical hardware on the host machine. You make adjustments to the number of CPUs in the Hyper-V Manager.

1. Open the Hyper-V Manager.
2. In the left pane, right-click on the virtual machine you need to edit and select **Settings**. The Settings dialog box for the virtual machine opens.
3. In the left pane, under **Hardware**, select the network adapter to update. The network adapter settings open in the right pane.
4. In the **Virtual switch** drop-down list, select the network to which to map the NIC.

#### TIP

In this pane you can also enable Bandwidth Management. Consult the Hyper-V documentation for more information on these settings.

5. Click **OK**. The NIC configuration is immediately modified for when the virtual machine is started.

#### NOTE

These instructions also apply to changing the configuration after the virtual appliance has entered service.

## 0.8.19 Related Topics

- » [Adjusting the number of CPUs available to the Virtual Machine](#)

### Increase storage by adding new virtual drives

During the process of installing the virtual machine, you needed to connect the virtual hard drive (VHD) to the Exinda Virtual Appliance. Prior to powering the VM on for the first time, it is likely that you should need to increase the size of the VHD. You can also make this adjustment after bringing the Exinda Virtual Appliance into service. You make adjustments to the size of the VHD in the Hyper-V Manager by adding additional hard drives to the VM.

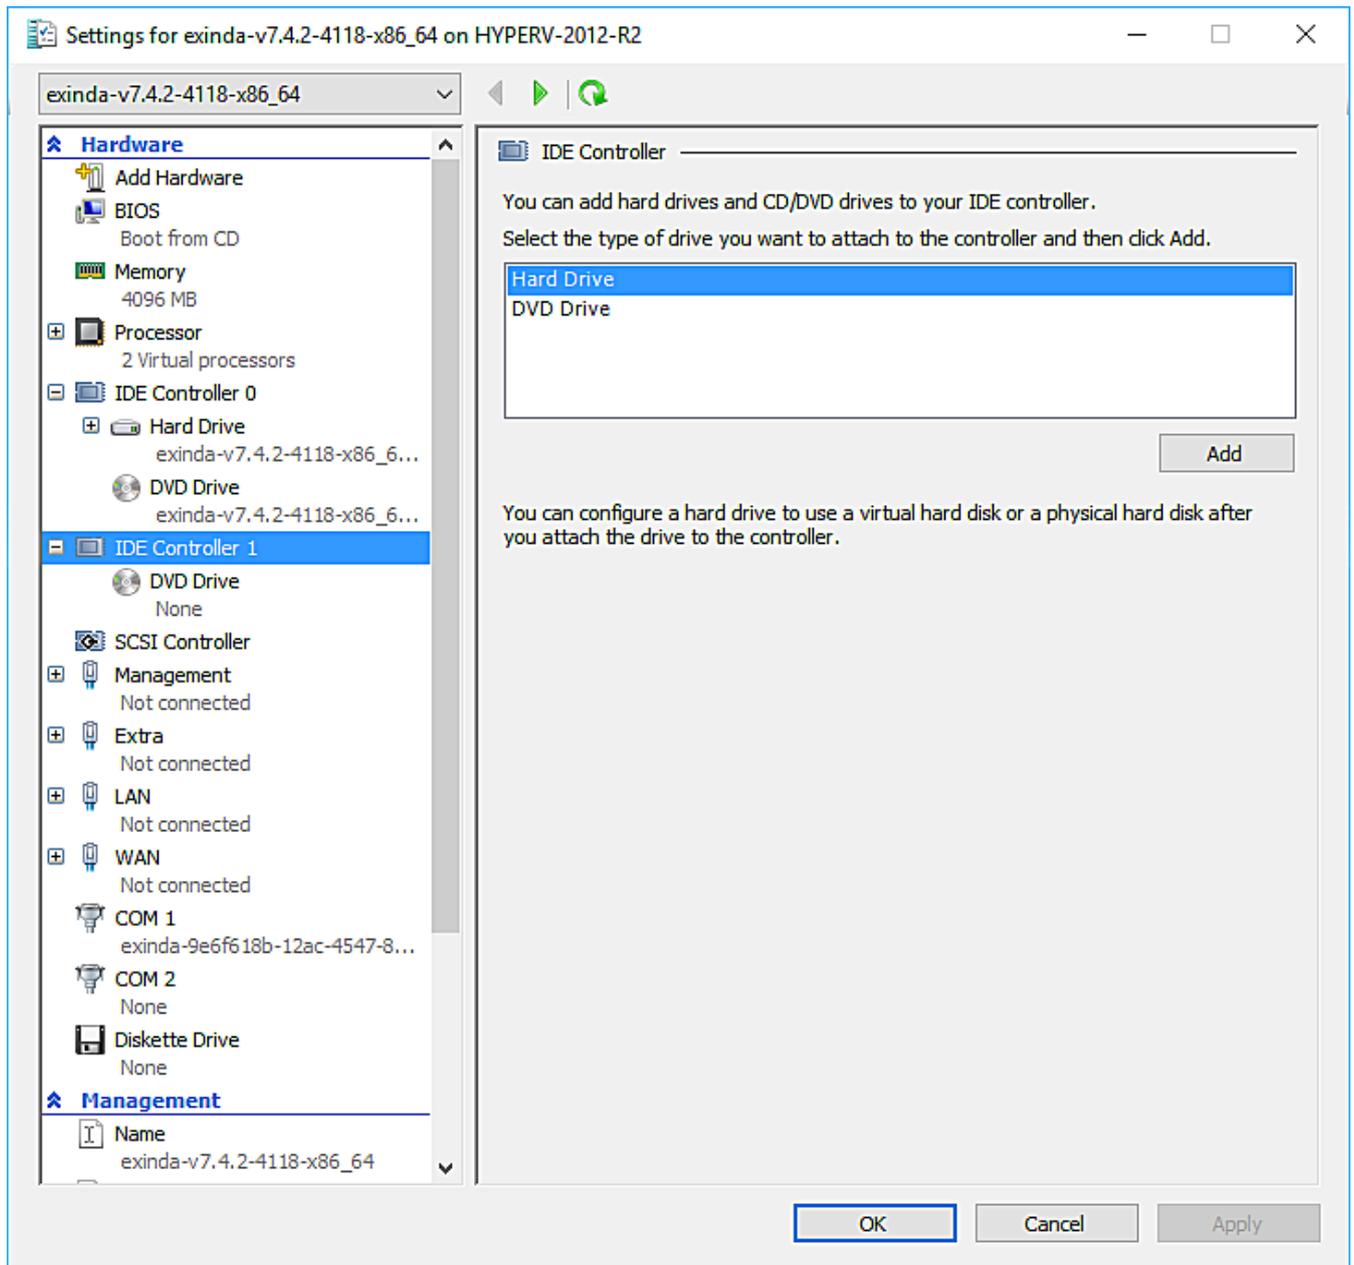
#### Prerequisites

Before starting this task, ensure that the virtual machine is switched off.

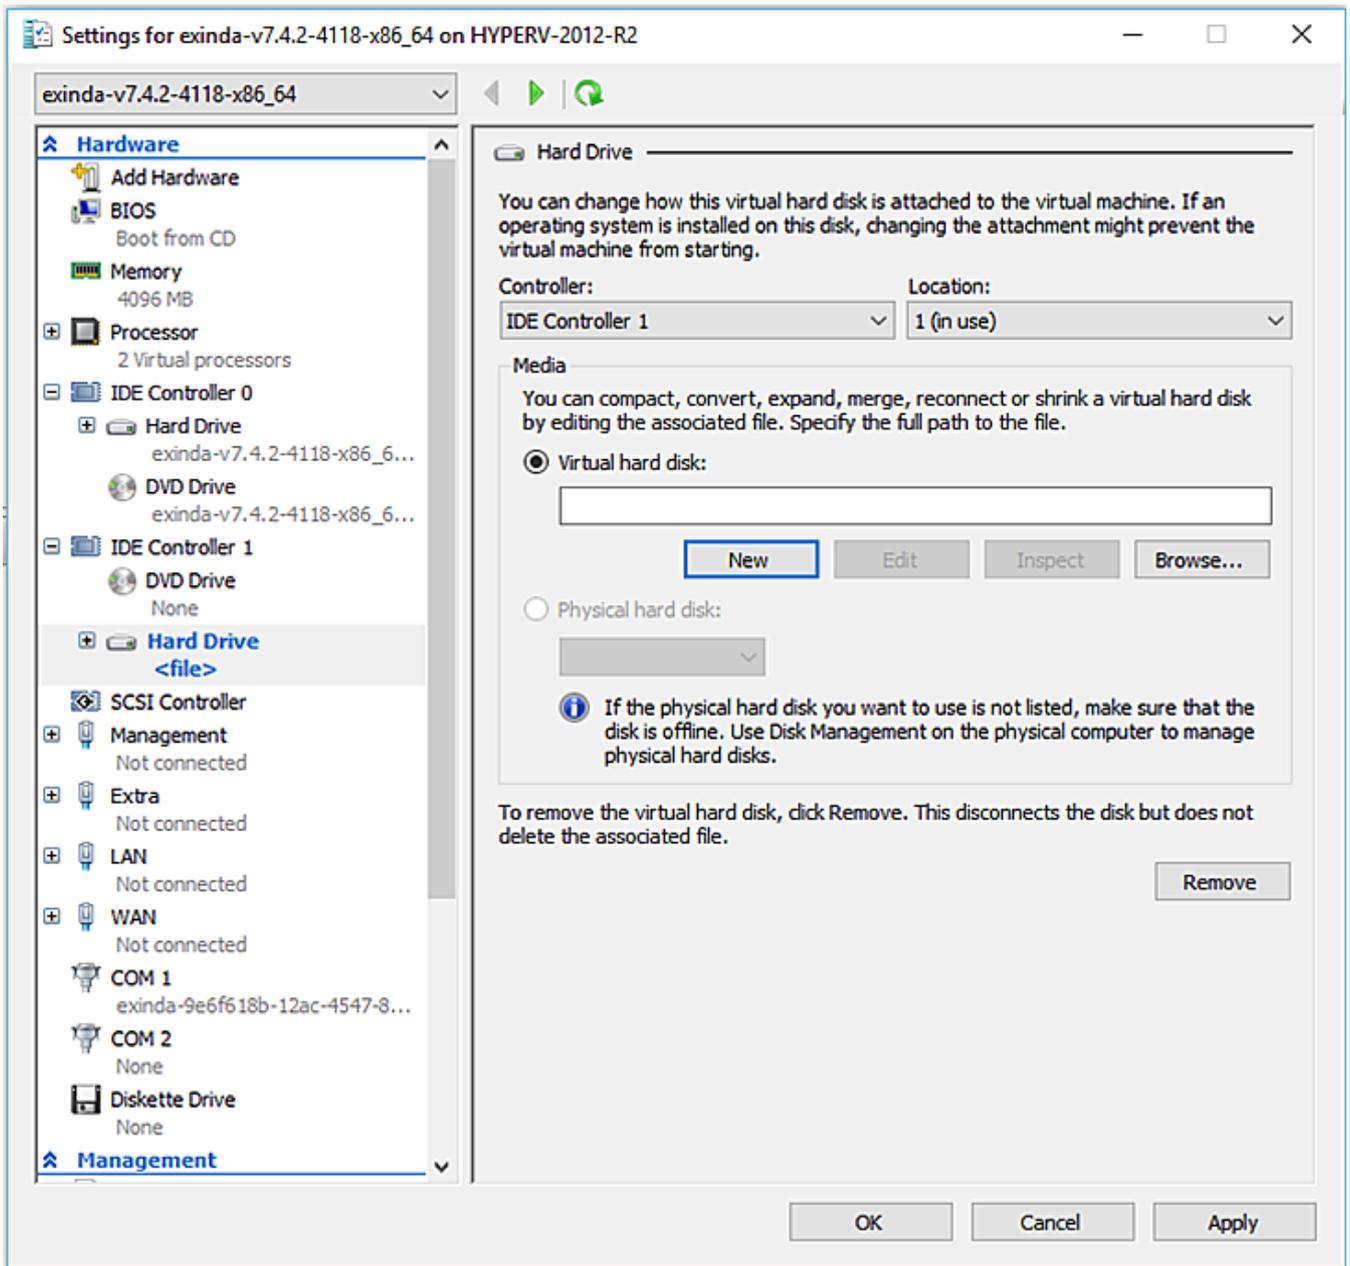
#### Procedure

1. Open the Hyper-V Manager.
2. In the left pane, right-click on the virtual machine you need to edit and select **Settings**. The Settings dialog box for the virtual machine opens.

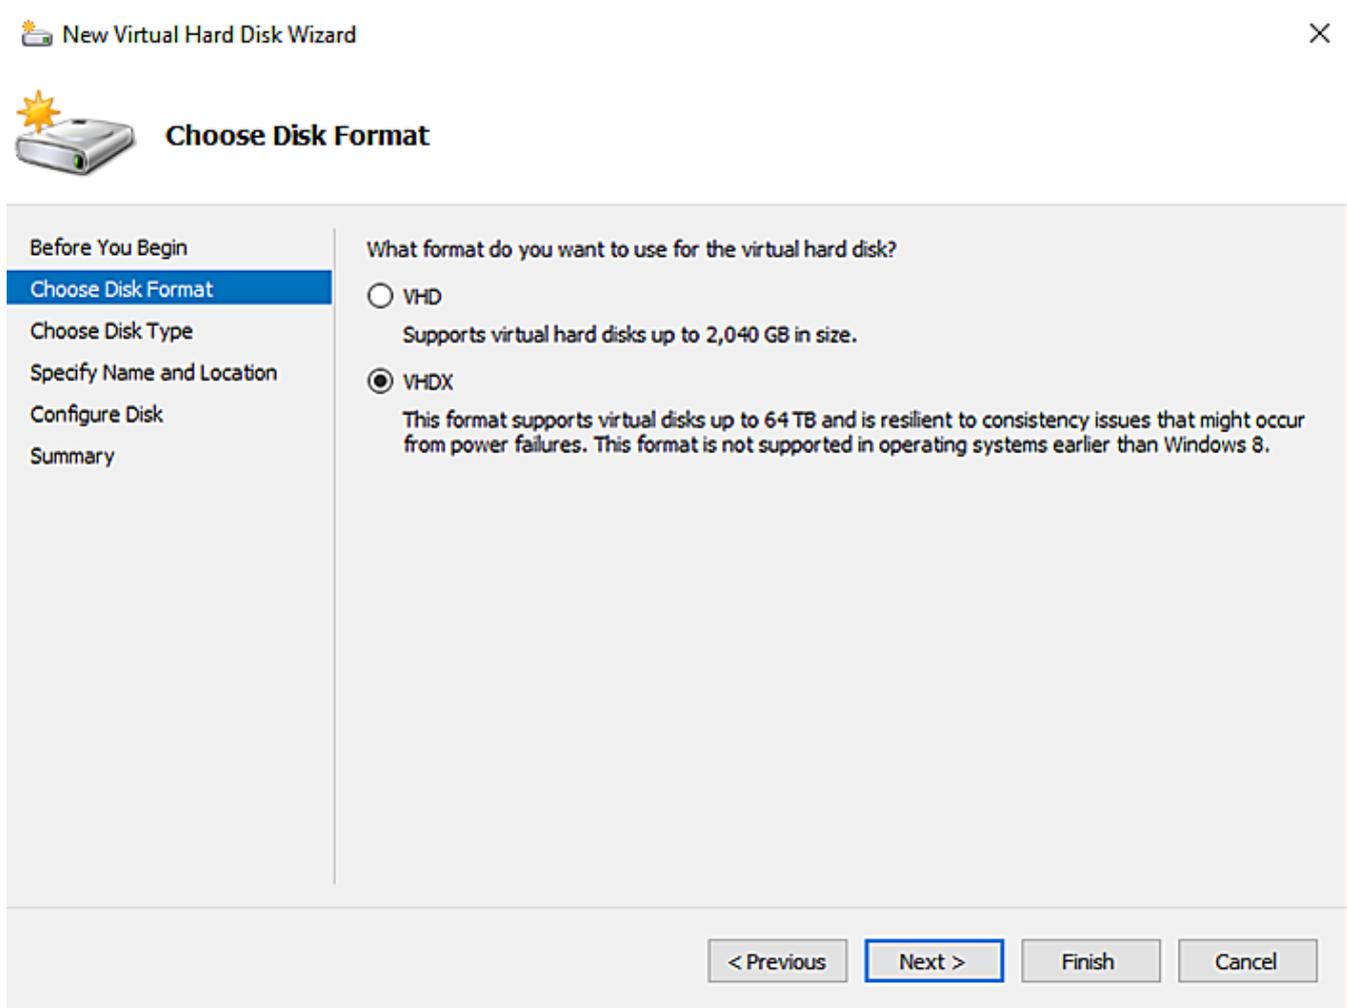
3. In the left pane, under **Hardware**, select any IDE Controller item. The Hard Drive settings open in the right pane.
4. Select the **Hard Drive** option in the right panel and click **Add**.



5. In the Hard Drive section, select "IDE Controller 1" as the Controller and "1 (in use)" as the location. By default, this is the only slot available in the virtual machine to which to insert a new Virtual Hard Drive. However, if more hard drives are needed in the future, you could remove the DVD Drives present by default given that these are not needed in the appliance. In such a case, Controller 0: Location 1 and Controller 1: Location 0 will also be available for further use.



6. Click **New**. The New Virtual Hard Disk wizard opens.



7. Select VHDX as the Disk Format type and click **Next**.



## Choose Disk Type

Before You Begin

Choose Disk Format

**Choose Disk Type**

Specify Name and Location

Configure Disk

Summary

What type of virtual hard disk do you want to create?

Fixed size

This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The virtual hard disk file that is created initially uses the size of the virtual hard disk and does not change when data is deleted or added.

Dynamically expanding

This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The virtual hard disk file that is created is small initially and changes as data is added.

Differencing

This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily. All children must have the same virtual hard disk format as the parent (VHD or VHDX).

< Previous   Next >   Finish   Cancel

8. In the **Choose Disk Type** section, select the **Fixed Size** option and click **Next**.



### Specify Name and Location

Specify the name and location of the virtual hard disk file.

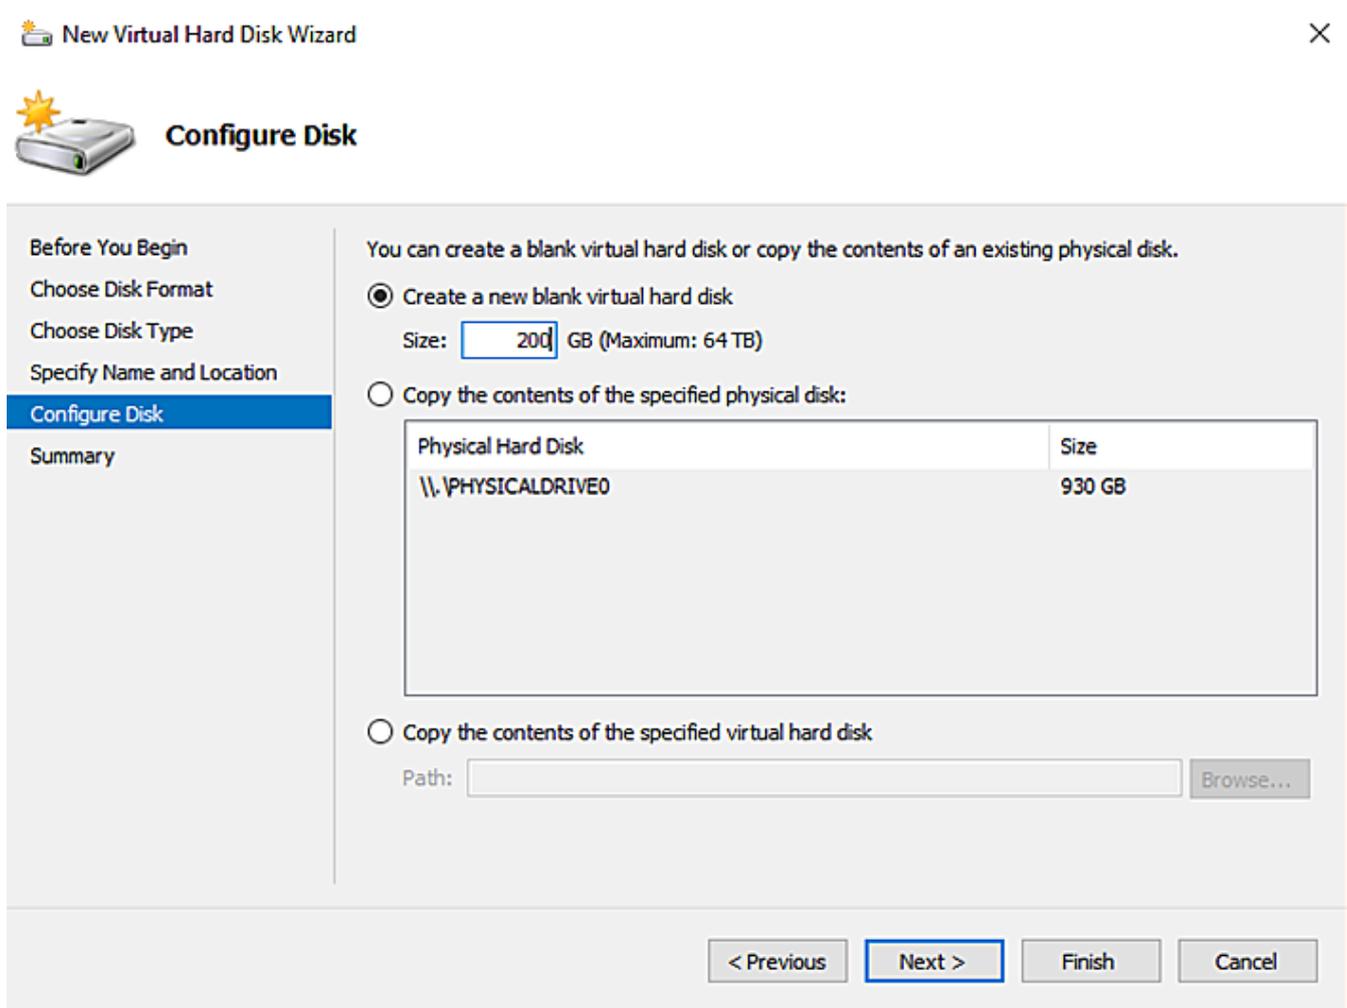
Before You Begin  
Choose Disk Format  
Choose Disk Type  
**Specify Name and Location**  
Configure Disk  
Summary

Name:

Location:

< Previous    **Next >**    Finish    Cancel

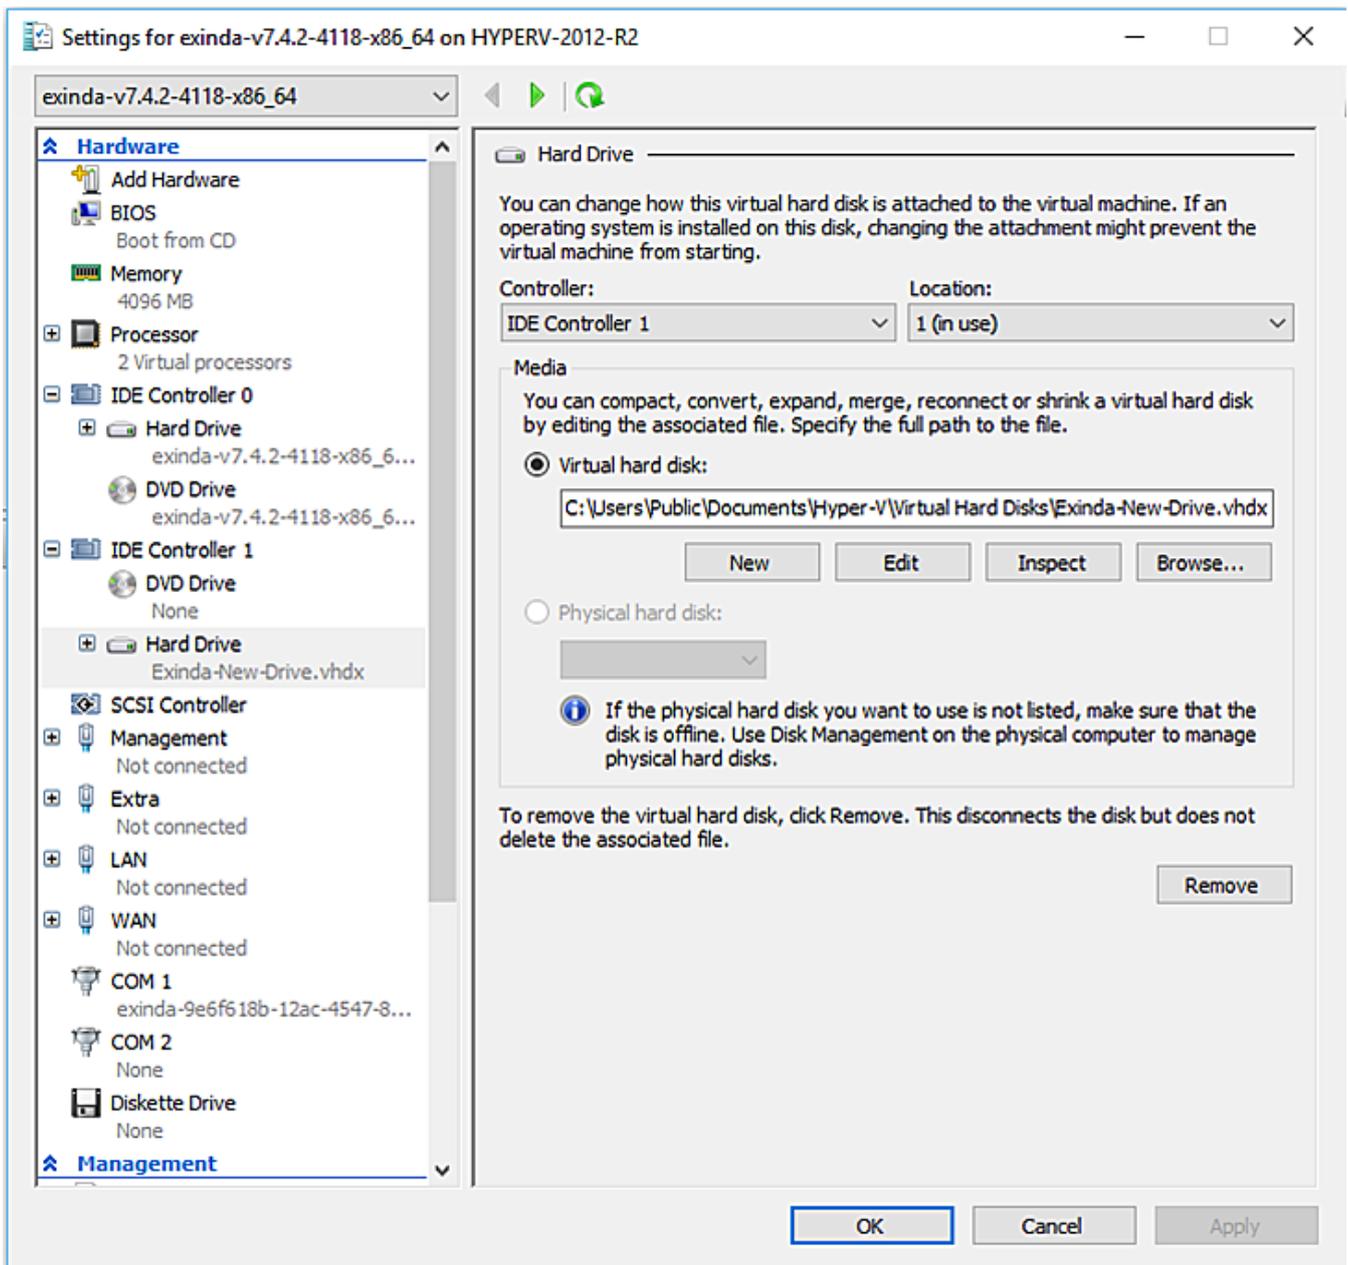
9. Specify a **Name** and **Location** for the virtual hard drive, and click **Next**.



10. Set the **Disk Size** based on the license you have acquired and then click **Next**. The minimum recommended sizes are the following:

- EXNV-VM Small: 250 GB Total (Add a 200GB Disk)
- EXNV-VM Medium: 250 GB Total (Add a 200GB Disk)
- EXNV-VM Large: 500 GB Total (Add a 450GB Disk)
- EXNV-VM Extra Large: 500 GB Total (Add a 450GB Disk)
- EXNV-VM Extra Extra Large: 500 GB Total (Add a 450GB Disk)

6. Click **Finish** to create the hard drive. This can take a few minutes.



7. When the Hard Drive settings page for the newly created drive opens, click **OK**.

8. Start the virtual machine. When the VM starts, it will automatically recognize the new drive, but the new storage must be manually added the virtual appliance.

#### NOTE

Before connecting, the management interface must already have been configured with an IP address or will obtain an IP address using DHCP. You need to make sure that the Management Interface is connected to the proper Virtual Switch in your Hyper-V environment.

9. Find the IP address assigned to the management interface by right-clicking on the VM and selecting the **Connect** option. This provides console access.

10. Log on to the appliance using the default credentials (username: admin, password: exinda). You might need to accept the EULA before proceeding.

11. Apply the following commands. The output contains the IP address you need to access the appliance web user interface.

```
exinda> en
```

```
exinda># show int eth0
```

12. Connect through HTTPS to the Exinda appliance using a browser.

13. Once logged on, click **Configuration > System > Setup > Storage**.

14. Add the new drive.

## NOTE

The following screenshot assumes that the chosen controller was **1**, and the location was **0**, so the new drive is **sdd**.

The screenshot displays the 'Storage Configuration' section of the Exinda web interface. It includes a 'Storage Configuration' table and a 'Disk Configuration' table. The 'Storage Configuration' table shows the following data:

| Service                         | Status    | Free         | Size           | Minimum  | Encrypted | Operation             |
|---------------------------------|-----------|--------------|----------------|----------|-----------|-----------------------|
| cifs                            | available | 5117.65M 96% | 5340.00M       | 1024.00M | ✘         | Resize Format Encrypt |
| edge-cache                      | available | 2184.37M 95% | 2304.00M       | 1024.00M | ✘         | Resize Format Encrypt |
| monitor                         | available | 9657.20M 94% | 10.00G         | 10.00G   |           | Resize Format         |
| users                           | available | 974.62M 95%  | 1024.00M       | 512.00M  |           | Resize Format         |
| wan-memory                      | available | 8168.84M 97% | 8448.00M       | 5120.00M | ✘         | Resize Format Encrypt |
| unallocated storage             |           |              | 200.00G        |          |           |                       |
| <b>Total Available Storage:</b> |           |              | <b>228.71G</b> |          |           |                       |

The 'Disk Configuration' table shows the following data:

| Disk       | Status        | Size              | Operation  |
|------------|---------------|-------------------|------------|
| sda        | in-use        |                   |            |
| sdb        | unused        | 214.7 GB          | Add        |
| sdcl0      | in-use        | 28.72 GiB         | Remove     |
| <b>sdd</b> | <b>in-use</b> | <b>200.00 GiB</b> | <b>Add</b> |

The new space appears as "unallocated storage" inside the "Storage Configuration" section.

## NOTE

Exinda recommends that you resize the "monitor" partition to at least 100GB. If you are licensed for acceleration, you should allocate most of the remaining storage in "wan memory" partition (cache partition for all TCP protocols), but if you are accelerating CIFS/SMB protocols, allow some storage in the "cifs" partition.

15. Allocate the storage as appropriate.

## 0.8.20 Related Topics

- » Adjusting the number of CPUs available to the Virtual Machine
- » Adjusting the RAM available to the Virtual Machine

- » [Adjusting the NICs available to the Virtual Machine](#)

## Customizing a Hyper-V Virtual Machine

As supplied, the Exinda Virtual Appliances will require some configuration changes before you introduce them to your network. For example, the virtual hard drives are limited to 50GB, which would be unlikely to be sufficient for your needs. The virtual machines available are sized with minimal configuration as it is not possible to know just what hardware is available on any host machine. To edit the configuration, you need to open the settings for the virtual machine in the Hyper-V Manager.

There are many settings that you can change, but for the purposes of configuring the Exinda Virtual Appliance, these task instructions are limited to what is necessary for bringing the appliance into an operational state. If you need more information, please consult the documentation for Hyper-V. This topic deals with changes to the configuration related to the number of CPUs, the available RAM, the NICs, and adjusting the storage for the virtual machine.

The configuration changes are required before your initial use of the virtual machine. You can also make further changes to your virtual machine at any later time. If over time you require more resources for the virtual machine, as long as those resources are available on the host, you can make them available to the guest.

### 0.8.21 Related Topics

- » [Adjusting the number of CPUs available to the Virtual Machine](#)
- » [Adjusting the RAM available to the Virtual Machine](#)
- » [Adjusting the NICs available to the Virtual Machine](#)
- » [Increasing the size of the virtual hard drive](#)

## Running on Linux KVM

**KVM** (Kernel-based Virtual Machine) is a complete virtualization solution for Linux on x86 hardware.

It contains virtualization extensions (Intel VT or AMD-V) and a loadable kernel module, **kvm.ko**. This provides the core virtualization infrastructure; and a processor specific module **kvm-intel.ko** or **kvm-amd.ko**.

By using KVM, you can run multiple virtual machines running unmodified Linux or Windows images. Each virtual machine has private virtualized hardware: a network card, disk, graphics adapter, etc.

See the following for more information:

---

---

## Installing KVM on Dedicated Hardware

[Learn to install KVM on dedicated hardware with Ubuntu server.](#)

You can Install KVM by running the below command:

```
sudo apt-get install qemu-kvm libvirt-bin virtinst bridge-utils cpu-checker
```

Verify that KVM is correctly installed by running the below command:

```
kvm-ok
```

If you get the below output then the server hardware is capable of running virtual machines. If not, you have to enable this from the BIOS.

*INFO: /dev/kvm exists*

*KVM acceleration can be used*

## Network Setup

### Learn to setup the network

The network setup depends on a particular configuration. Bridges are used in KVM to add network interfaces to a particular VM. Each Exinda appliance requires at least 2 network interfaces.

### Installing the VM

To install the VM, run the below command:

```
sudo virt-install \  
--virt-type=kvm \  
--name <EXINDA_VM_NAME> \  
--ram 4096 \  
--vcpus=2 \  
--os-variant=centos6.3 \  
--virt-type=kvm \  
--hvm \  
--cdrom=<PATH_TO_ISO> \  
--network=bridge=<BRIDGE_1>,model=virtio \  
--network=bridge=<BRIDGE_2>,model=virtio \  
--disk path=/var/lib/libvirt/images/<EXINDA_VM_\  
NAME>.qcow2,size=40,bus=virtio,format=qcow2
```

The above command will create the VM and also allocate the disk stored in qcow2 format.

#### Note

You can get the installation media ISO from your Exinda partner.

The domain might take some time to be created, as the script that installs Exinda OS is being executed. You can view the progress by opening another console and running the below command:

```
sudo sudo virsh console $VM_ID
```

Where `$VM_ID` is the unique identifier for the running VM. You can view running VMs by running `virsh list` and `virsh list --all` for all VMs (started and stopped).

#### Note

For Autostart run, `sudo virsh autostart $VM_ID`.

### Learn to Configure Exinda

To configure Exinda from the console you can attach to the VM console using the below command:

```
sudo sudo virsh console $VM_ID
```

Alternatively, you can use the web browser and access the Exinda configuration through the web GUI.

## 0.8.22 Creating an initial configuration using the Basic Wizard

The initial configuration wizard steps you through configuring the appliance's interfaces, IP settings, HTTP proxy settings, basic system information, license information, and storage volume. It also provides the option to upgrade the firmware and create the initial set of traffic policies.

1. The Exinda appliance by default picks up an IP Address from DHCP. The IP address is available on the management interface.

### NOTE

If a DHCP address is not picked up, the Exinda defaults to the IP Address 172.14.1.57. Open a web browser and connect to the Web User Interface by typing `https://172.14.1.57` in the address field. To connect, configure the IP address of your PC to the same subnet as the Exinda Appliance. For example, set your IP Address to 172.14.1.58, netmask 255.255.255.0

2. From the web browser, access Command Line Interface from [Accessing the Command Line Interface](#).

a. Enter **show interface eth0** and note the IP address assigned to the interface.

3. Navigate to `https://<IP Address>`, where IP Address is the address determined in the previous step.

4. Login with **username=admin** and **password=exinda**.

5. Select **Configuration > Basic Install Wizard** to start the configuration wizard .

- **Basic Wizard Step 1 - Interfaces:** This screen lists all the system interfaces, as well as reports any problem with the interfaces. You can set interface speed and duplex settings from this screen.

Step 1: Interfaces

| Interface | Speed | Duplex | Link Status |
|-----------|-------|--------|-------------|
| eth1      | Auto  | Auto   | ✔           |
| eth2      | Auto  | Auto   | Unplugged   |
| eth10     | Auto  | Auto   | ✔           |
| eth11     | Auto  | Auto   | ✔           |

Back   Next

- **Basic Wizard Step 2 - IP Settings:** This screen allows you to configure basic network connectivity settings. You can either manually specify these settings or select **Autoconf** to automatically acquire these settings. The type of auto configuration selected depends on your network. For IPv4 networks select **DHCP**, for IPv6 use **SLAAC**.

Step 2: IP Settings ▼

**Static**
 **Autoconf**

IPv4:  DHCP    IPv6:  SLAAC

---

\* Address (eth1) 192.168.110.70/24  
fe80::224:e8ff:fe3d:caed/64

Default IPv4 Gateway 192.168.110.1

Default IPv6 Gateway

\* Host Name

Primary DNS 172.16.1.254

Secondary DNS

---

br10

[eth1] [eth2] [eth10] [eth11]

LAN WAN

SLOT 1

\* Required field

- Basic Wizard Step 3 - HTTP Proxy Settings:** To allow the appliance to access Exinda's HTTP server for firmware updates, license updates, and messages, specify an HTTP proxy. If you have SDP enabled, ensure your proxy supports HTTPS.

Step 3: HTTP Proxy Settings ▼

Specify a HTTP proxy so the appliance can access Exinda's HTTP server for firmware updates, license updates and messages. If you have SDP enabled, please ensure your proxy supports HTTPS.

HTTP(S) Proxy Address

HTTP(S) Proxy Port

HTTP(S) Proxy Authentication

HTTP(S) Proxy Username

HTTP(S) Proxy Password

Do not verify SSL certificates

- Basic Wizard Step 4 - System:** This screen allows to configure basic system settings.

Step 4: System ▼

Domain Name

SMTP Server Name

Time Zone

New admin Password

Confirm Password

- Basic Wizard Step 5 - Licensing:** This screen allows you to configure the system's license. When you enter the screen, the Exinda appliance attempts to contact the Exinda licensing server on the Internet. If the

appliance has Internet connectivity and a new or updated license can be found, it is displayed in the text-box at the bottom of the screen. You can add this license to the system by clicking the **Add License** button.

Step 5: Licensing

|                              |                           |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |
|------------------------------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|---|----------|---|------------|---|
| Bandwidth                    | 102400 kbps               | <table style="width: 100%; border-collapse: collapse;"><tr><td>Monitor</td><td style="text-align: right;">✔</td></tr><tr><td>Optimize</td><td style="text-align: right;">✔</td></tr><tr><td>Accelerate</td><td style="text-align: right;">✘</td></tr></table> | Monitor | ✔ | Optimize | ✔ | Accelerate | ✘ |
| Monitor                      | ✔                         |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |
| Optimize                     | ✔                         |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |
| Accelerate                   | ✘                         |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |
| Software Subscription Expiry | <b>Dec 31, 2016 (45d)</b> |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |
| License Expiry               | No license expiry date    |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |
| Host ID                      | 0010f305cd54              |                                                                                                                                                                                                                                                               |         |   |          |   |            |   |

License(s) Installed:  
LK2-EXINDA-45A0-023R-GBKA-L5W3-E8H5-J434-005L-115M-05N4-BP00-5P23-45Q0-5R1L-5T24-N5U1-L5V2-G086-GT40-CB58-5KNX-KK0H-CBAY-GT38-X00K

Looking for a license online ...

Connection completed successfully. No new license found.

Add License

Back Next

- **Basic Wizard Step 6 - Storage:** This screen displays the available disks that can be added to the volume group.

Step 6: Storage

**Do you want to add the following disks to volume group when this wizard is completed?**  
**Note that this will delete all existing data on the disk**

**Volume:** sdb  
**Model:** Virtual disk  
**Size:** 17.1 GB

Yes    No

Back Next

- **Basic Wizard Step 7 - Firmware:** This screen displays the status of the firmware running on the Exinda appliance. If the appliance has Internet connectivity, the system checks for any newer firmware that may have been released. If a newer firmware image is available, you are asked if you want to download and install it.

Step 7: Firmware

**Firmware is up to date - no new update found.**

Back Next

- **Basic Wizard Step 8 - Optimization:** The final screen allows you to configure default Optimizer policies.

**Note: This wizard will delete all existing Optimizer Policies.**

Step 8: Optimizer

**Step 1:** Do you want to start Optimization when this wizard is completed?  Yes  No

**Step 2:** Do you want to configure new Optimization Policies?  
*Selecting YES will overwrite any existing policies you have configured.*  Yes  No

**Step 3:** Do you want to enable Optimization?  
*Selecting YES will create policies that optimize and accelerate WAN applications. Note: You must have another Exinda appliance on the WAN for this to work.*  Yes  No

**Step 4:** Do you want to enable Control?  
*Selecting YES will apply traffic shaping.*  Yes  No

**Step 5:** Select the topology type WAN or WAN + Internet?  WAN  WAN + Internet



Internet traffic for this site is routed over the WAN, usually via another site.

**Step 6:** Enter inbound bandwidth (kbps)?  kbps  
*(MAX = 1024000)*

**Step 7:** Enter outbound bandwidth (kbps)?  kbps  
*(MAX = 1024000)*

#### NOTE

Settings on each step are automatically applied when clicking the 'Back' or 'Next' buttons.

### 0.8.23 Licensing information

Exinda Appliances are licensed using a subscription-based model. Licenses are renewed on a recurring basis and include all upgrades and releases available during the subscription period.

Subscription prices vary between Exinda editions (Plus, Professional and Premium) and appliance types (hardware or virtual). All Exinda editions subscriptions come Edge Cache, Real Time Monitoring, Interactive Analytics, Reporting and Recommendation Engine. All hardware editions come with a 3 year warranty and Acceleration.

For more information about pricing plans, contact your authorized Exinda Inc. reseller. To find a reseller in your area, visit <https://www.gfi.com/partners/find-a-partner>.

## IMPORTANT

- » Product activation requires an Internet connection and a [GFI login account](#).
- » A license subscription period starts at the time of purchase, not at the time of activation.
- » The GFI subscription service connects to your appliance over the Internet every 24 hours. New licenses are automatically applied.
- » Hardware appliances automatically fetch license keys from the GFI license server on first boot-up.
- » Virtual appliances are shipped unlicensed. On first-boot, they automatically generate a unique Host ID. GFI must be notified of this Host ID before a license can be activated. Virtual appliances must be licensed within 96 hours or they stop functioning.

If your appliance is unable to access the Internet and the license server, the license key can be applied manually by entering the license key.

### Managing a license using the GFI Accounts Area

Licenses and subscription information are stored in the [GFI Accounts Area](#). Access requires a GFI Account.

Log in to the [GFI Accounts Area](#) using the email address and password entered when your GFI Account was created.

If you don't know the account password, reset it by following the instructions in [How to reset a lost password](#). If you don't know the email address associated with your account, contact your reseller.

If you don't have a GFI Account, create one in one of the following ways:

- » Upgrade a licensed product.
- » Have your Exinda Inc. reseller set up the account.
- » Request a trial of a GFI product.. For more information, see [Requesting a product trial](#)..

For more information about the GFI Accounts area, see [GFI Accounts Area Help](#).

### To see what licenses are available

1. Go to **Configuration > System > Setup > License**.

- The effective license limits are shown. If a feature is included in your license, a green checkmark  appears beside the feature.
- For Edge Cache, if the license includes this feature, the number of connections included in the license is listed.
- If multiple licenses have been applied to your appliance, the features included are a composite of the multiple licenses. The appliance will use the license that provides the highest specification limits.

| Licensed | Host ID      | Model                            | SS Expiry           |
|----------|--------------|----------------------------------|---------------------|
| ✔        | f129ade8c50f | Exinda x800 (1000/2500Mbps - HP) | <b>Sep 30, 2012</b> |
|          |              | Max Bandwidth:                   | 1000 Mbps           |
|          |              | Optimizer:                       | ✔                   |
|          |              | Max AA Bandwidth:                | 2500 Mbps           |
|          |              | Max Connections:                 | 10240000            |
|          |              | Max Connection Rate:             | 10000 / sec         |
|          |              | Max AA Connections:              | 12000               |
|          |              | Max PDF Reports:                 | 100                 |
|          |              | Max SLA Objects:                 | 100                 |
|          |              | Max APS Objects:                 | 100                 |
|          |              | Max APM Objects:                 | 100                 |
|          |              | Max Policies:                    | 1000                |
|          |              | SSL Acceleration:                | ✔                   |
|          |              | Virtualization:                  | ✔                   |
|          |              | Max Edge Cache Connections:      | 10000               |

Screenshot 46: License

| Licensed            | License Status                                                                                                                                                     |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Host ID             | Unique identifier of each Exinda appliance software.                                                                                                               |
| Model               | Exinda appliance model.                                                                                                                                            |
| SS Expiry           | Expiry date of Exinda Software Subscription. After this date, you are no longer entitled to support and no new software updates can be installed on the appliance. |
| Max Bandwidth       | Maximum monitoring and QoS bandwidth.                                                                                                                              |
| Optimizer           | Entitlement for QoS and Acceleration.                                                                                                                              |
| Max AA bandwidth    | Maximum acceleration bandwidth (WAN side).                                                                                                                         |
| Max Connections     | Maximum concurrent connections through the appliance.                                                                                                              |
| Max Connection Rate | Maximum number of new connections per second. Exceeding this will cause the network problems as any more connections will get dropped at setup time.               |
| Max AA Connections  | Maximum number of concurrent connections that can be accelerated. Exceeding this limit will mean the any new connections are not accelerated.                      |
| Max PDF reports     | Maximum number of PDF reports that can be automatically generated and emailed.                                                                                     |
| Max SLA Objects     | Maximum Service Level Agreement objects.                                                                                                                           |
| Max APS Objects     | Maximum Application Performance Score objects.                                                                                                                     |
| Max Policies        | Maximum number of optimization policies. Regardless of Circuit and VC.                                                                                             |

| Licensed         | License Status                                                                       |
|------------------|--------------------------------------------------------------------------------------|
| SSL Acceleration | Entitlement to SSL Acceleration.                                                     |
| Virtualization   | Entitlement to Virtualization.                                                       |
| Edge Cache       | Entitlement to Edge Cache Acceleration and the maximum number of edge cache objects. |

2. Please contact your local Exinda representative if you wish to enable a feature.

To see the last time that the auto-license service checked for a new license

1. Go to **Configuration > System > Setup > License**.
2. The **Last Check** and **Last Update** date & time is shown in the Auto-License section at the top of the page.

| Auto License Service: <b>Running</b> <input type="button" value="Restart"/> <input type="button" value="Stop"/> <input type="button" value="Disable"/> |                                   |                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|----------------------------------------------|
| License Server                                                                                                                                         | Last Check                        | Last Update                                  |
| license.exinda.com                                                                                                                                     | 2009/11/20 16:02:08 (28m 42s ago) | 2009/09/30 09:51:36 (1d 12h 36m 26.704s ago) |

To force the auto-license service to check for a new license

1. Go to **Configuration > System > Setup > License**.
2. Near the bottom of the page, click the **Check for License Online** link.
3. The system will either provided feedback that no license was found or will present you with a list of licenses for you to select from and apply.
4. Click the **Add Licenses** button to apply the newly found license.

To stop, restart, or disable the auto-license service

1. Go to **Configuration > System > Setup > License**.
2. At the top of the page in the Auto-License section, click the **Restart** button, **Stop** button, or **Disable** button as desired.

| Auto License Service: <b>Running</b> <input type="button" value="Restart"/> <input type="button" value="Stop"/> <input type="button" value="Disable"/> |                                   |                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|----------------------------------------------|
| License Server                                                                                                                                         | Last Check                        | Last Update                                  |
| license.exinda.com                                                                                                                                     | 2009/11/20 16:02:08 (28m 42s ago) | 2009/09/30 09:51:36 (1d 12h 36m 26.704s ago) |

To manually add a license key

1. Go to **Configuration > System > Setup > License**.
2. At the bottom page, type your license key. To add multiple licenses, separate each license with a comma (,).
3. Click the **Add Licenses** button.

To see previously applied licenses

1. Go to **Configuration > System > Setup > License**.
2. Below the current license status table, you can see all the licenses that have been applied to your appliance.

- Notice the Valid columns and Active columns to understand which licenses are currently applied.
- The appliance will use the license that provides the highest specification limits.
- License keys can also be removed from the system by clicking 'Remove'. Before removing, ensure that you keep a copy of the license key.

|                          | License Key                                                                                                                                               | Feature                             | Valid                               | Active                              |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| <input type="checkbox"/> | LK2-EXINDA-45A0-048C-W93E-45W3-F4N5-J3L0-05L1-15M3-L005-N4BP-005P-29C5-Q31E-V5R1-C5T2-3Q5U-24N5-V2C0-5Y11-4X11-6011-86GT-6W4Y-H2B8-TNCA-RCDE-1TC7-006J-JY | EXINDA                              | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
|                          | Tied to hex host ID:                                                                                                                                      | 002219d48dc4                        |                                     |                                     |
|                          | SS Expiry Date:                                                                                                                                           | 2012/06/19                          |                                     |                                     |
|                          | Max Bandwidth:                                                                                                                                            | 20480                               |                                     |                                     |
|                          | Optimizer Enabled:                                                                                                                                        | <input checked="" type="checkbox"/> |                                     |                                     |
|                          | Max AA Bandwidth:                                                                                                                                         | 20480                               |                                     |                                     |
|                          | Max Connections:                                                                                                                                          | 384000                              |                                     |                                     |
|                          | Max Connection Rate:                                                                                                                                      | 300                                 |                                     |                                     |
|                          | Max AA Connections:                                                                                                                                       | 1500                                |                                     |                                     |
|                          | Max PDF Reports:                                                                                                                                          | 12                                  |                                     |                                     |
|                          | Max SLA Objects:                                                                                                                                          | 120                                 |                                     |                                     |
|                          | Max APS Objects:                                                                                                                                          | 150                                 |                                     |                                     |
|                          | Max Policies:                                                                                                                                             | 384                                 |                                     |                                     |
|                          | SSL Acceleration:                                                                                                                                         | <input checked="" type="checkbox"/> |                                     |                                     |
|                          | Virtualization:                                                                                                                                           | <input checked="" type="checkbox"/> |                                     |                                     |
|                          | Edge Cache Acceleration:                                                                                                                                  | <input checked="" type="checkbox"/> |                                     |                                     |

#### To re-enable a virtual appliance that has shut down due to not connecting

Contact Exinda TAC to re-enable your virtual appliance.

#### To generate a virtual appliance trial license

In this step you will go to the Exinda website to create your trial license. Your trial license will be emailed to you after you complete this step.

1. Navigate to [https://license.exinda.com/virtual\\_licensing](https://license.exinda.com/virtual_licensing).
2. Enter in the Host-ID for the virtual appliance.
3. Select the License Type for the virtual appliance.
4. Type your email address, and which Hypervisor Type on which you have installed the virtual appliance.
5. Click **Create**. Your trial license is emailed to the address provided.
6. Once you have received your license key, copy the license key into the Exinda GUI.
  - a. In a browser, navigate to the IP address assigned to your Exinda Virtual machine.
  - b. To view the status of your license, select **Configuration > System > Setup** and switch to the **License** tab.
  - c. If your Host-ID has been previously entered into the system, click **Check for License Online**. If this is the first time you are licensing the virtual appliance, and your Host ID was recently created, paste the license key provided in the email..

## NOTE

You must be connected to the Exinda License Server at all times for the virtual appliance to work.

d. Click **Add License**.

7. Confirm your system has been licensed by refreshing the page.
8. Save any changes, and restart the virtual appliance.

### To purchase a virtual appliance license

In this step you will purchase a license from a web form. You must have your Host ID information for each EXN-V appliance and the PO number. As with the trial license process, you must have installed the EXN-V and captured the Host ID information to complete this process.

1. In a browser, navigate to the address of your Exinda Virtual Appliance.
2. Log into your Exinda VM. The default user name is `admin`, and the password is `exinda`.
3. On the **Dashboard > System** tab, find the Host ID that the hypervisor created for this virtual machine. You must have your purchase order number that details the type of license and number of licenses you have purchased. You will need the following:
  - Host ID
  - Hypervisor Type
  - License Level – this will be based on a bandwidth license

You will need this information for each EXN-V virtual appliance.

4. To purchase a license, navigate to [https://license.exinda.com/virtual\\_purchase/](https://license.exinda.com/virtual_purchase/).
5. Complete the Virtual Appliance Purchase form as required.

## 0.9 Managing multiple appliances with the Exinda Management Center

The Exinda Management Center (EMC) provides complete management insight and configuration control of your Exinda Network Orchestrator appliances from one central console. All applications, devices, users, and activities across all network locations are managed from a central location giving IT Administrators the ability to manage network policies and manage appliance configuration across the entire organization.

You can configure Exinda Appliances and monitor network usage directly from the appliance. However, once you have more than a few appliances to manage, it can become difficult to manage them individually and maintain standard configurations when needed.

The Exinda Management Center solves the management gap by enabling policy configuration on multiple appliances. When used in conjunction with SDP (Service Delivery Point), you can also have aggregated reporting of your network traffic across your appliances and/or reporting of the individual appliances all within a single report.

If deploying an on-premises instance of EMC, you can deploy it as a multi-tenant solution, where each estate is enrolled under a separate tenancy account. Objects and data cannot be shared across tenancies. For a single estate use of EMC, a single tenant is added to the system. Within a tenant, multiple appliance groups can be added to help organize the appliances and to ensure that the correct configuration is sent to the sets of appliances. These appliance groups can be nested in other appliance groups.

If using the Exinda-hosted service, only a single tenant appears in the system.

## 0.9.1 Getting started with EMC

Get started with Exinda Management Center using the topics below.

### 1 Integrate your Appliance with EMC

Make sure all your appliances are configured to call EMC.

1. On the Exinda Web UI of the appliance, go to the **Configuration > System > Setup > SDP**.
2. Enable the SDP Client option and key-in `mc.exinda.com` as the SDP Server.

### 2 Log in to EMC

1. Navigate to [mc.exinda.com](http://mc.exinda.com) and key-in your credentials.
2. Read and accept the license agreement.

### 3 Perform basic EMC configuration

After configuring the Exinda Management Center to communicate with the Exinda Network Orchestrator appliances, perform a basic configuration to be ready for creating policies and sending it to your appliance groups. For more information, refer to [Configuring the EMC](#) (page 155).

### 4 Import or perform manual appliance configuration

You may want to retain your appliance network objects and policies, so you have the option to import them into a common library for future use and integrate them into the global configuration for your group. For more information, refer to [Importing appliance configuration](#) (page 159).

In case the appliance is new, use the Library to add circuits, virtual circuits, policy sets and policies. These settings are all re-useable and are available as a selection when defining the optimizer policy for each appliance group. For each group, you can then create an optimizer policy tree to combine circuits, virtual circuits, policy sets and policy rules. For more information, refer to [Optimizer Policy Tree](#) (page 170).

The Optimizer Policy Tree is applied based on a hierarchical structure of Appliance groups, where child groups can inherit policy from parent groups. It is therefore important to start by defining the group names you will need. For more information, refer to [Configuring an appliance manually](#) (page 164).

The Optimizer Policy Tree defines actions to be taken on different types of traffic going through the appliance. The tree is processed in a top-to-bottom order, so the policies on traffic are applied accordingly.

### 5 Push Configuration

Once all the desired changes have been made to the group, you can simply push the configuration to the appliances within an Appliance Group. For more information, refer to [Sending configuration changes to the appliances](#) (page 176).

## 0.9.2 Deploying Exinda Management Center as a Virtual Machine

The Exinda Management Center product is available as a Software-as-a-Service, and as a virtual machine.

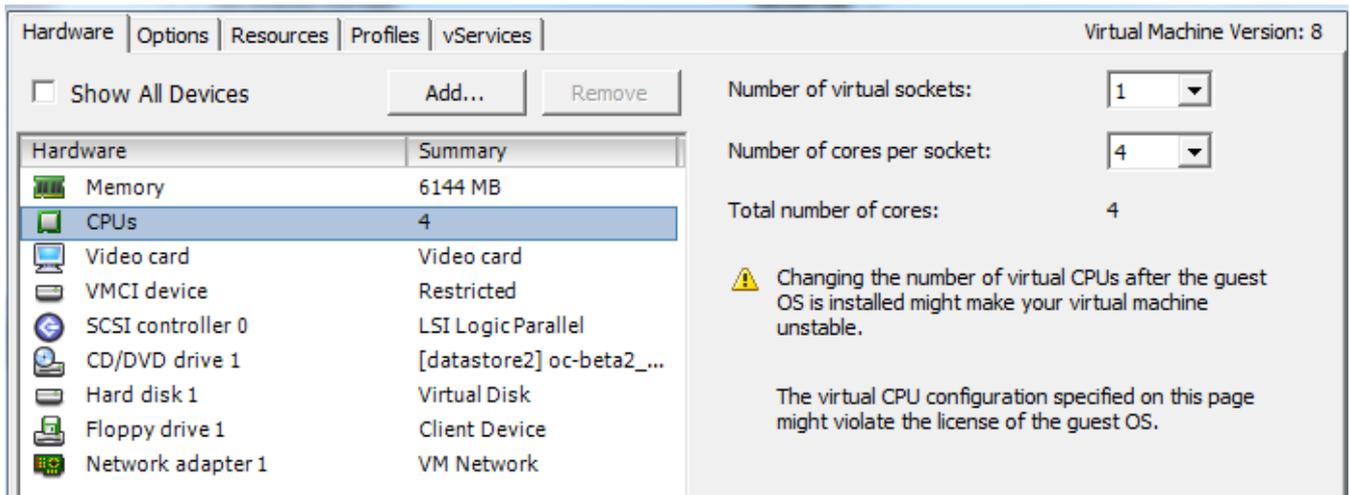
This topic describes how to deploy the Exinda Management Center as a virtual machine and how to start the VMware virtual machine.

---

### Install Requirements

The Exinda Management Center virtual machine needs the following resources:

- » Minimum of ESX 5.5 is required
- » 40 GB of disk storage is required
- » Within the VM, four CPU cores are required



- » The VMware OVA file is pre-configured with 6 GB of RAM
- » A NIC needs to be available on VMware
- » The virtual machine must be able to connect to the Internet and reach [license.exinda.com](http://license.exinda.com).

### Deploying on VMware ESX 5.5

The Exinda Management Center image is available as an OVA for VMware ESX 5.5 and higher.

1. Liaise with your local Exinda representative to obtain the download files.
2. Open the VMware vSphere client.
3. Select **File > Deploy OVF Template**.
4. Copy the URL of the latest release of the Exinda VMware Virtual Appliance from Exinda.com, and paste it into the **Deploy from...** field. Click **Next**.
5. Confirm the OVA template details are correct, and click **Next**.
6. Review and accept the End User License Agreement (EULA). Click **Next**.
7. Specify a name for the virtual machine.
8. If prompted, choose the location to deploy the virtual machine.
9. Click **Next**.
10. Choose the format for storing the virtual disks for the virtual machine. Exinda recommends **Thick Provisioning**, which is the default.

### Starting the VMware Virtual Machine

When you are ready to start the virtual machine for the first time:

1. Power on the virtual machine. The virtual machine boots and then displays a login prompt on the VMware console.
2. Log in with the default username `admin` and password `exinda`.
3. Get or set the IP address of the virtual machine.
  - a. If the first NIC is connected to a network that provides addresses using DHCP, the virtual machine should already have an IP address. On the virtual machine summary screen, VMware tools should display the IP address

the virtual machine has obtained. Or at the prompt, type : `show ip`

b. To set the IP, type :

```
enable
configure <network_card>
ip address <ip_address>[/<mask>]
ip default-gateway <ip_address>
```

3. Access the UI by entering `https://<ip-address>` in your web browser.

4. Login with the default username (`admin@example.com`) and password (`exinda`).

### 0.9.3 How EMC fits into the appliance feedback loop

Using Exinda Management Center, you can:

1. Monitor traffic on Exinda Appliances.
2. Use EMC to configure [Policy Rules](#) and apply them to multiple appliances.
3. Use Exinda Web UI of each appliance or EMC to create [Alerts](#) and [Application Performance Scores](#) for all appliances.
4. Wait for the appliances to send monitoring data notifications to EMC.
5. After receiving notifications, investigate them by monitoring your traffic on the appliances.
6. Configure and tune the [Optimizer Policy Tree](#) using Exinda Management Center and apply policy changes broadly across appliances as needed.
7. Repeat from step 4.

### 0.9.4 Exinda Management Center Concepts

You may be familiar with managing the Exinda Network Orchestrator appliances and may find it relatively easy working on EMC, but you must also know the new concepts introduced in the Exinda Management Center.

- » **Policy Sets:** A set of policies that can be re-used in multiple virtual circuits and various appliance groups.
- » **Circuit Type:** A named object that acts as the tie between circuits and appliance bridges. This allows different bridges on different appliances to be bound to the same circuit definition.
- » **Library:** Any item that is used and re-used in various policy trees is a Library item.
- » **Network Object Location:** The location of a network object (internal or external) is determined by the system comparing the IP addresses in the network object to those in the local network object.
- » **Service Level Agreements:** A set of Library objects that allow you monitor the availability of specified IP addresses.
- » **VLANs:** A set of Library objects that allow you to separate hosts on their functional attributes rather than their physical location.
- » **Dynamic Virtual Circuits:** A set of policies that allow you to enforce fair sharing among hosts or to limit the number of hosts on the circuit.

### 0.9.5 Best Practices

Here are the best practices when configuring the Exinda Management Center.

## Deployment

Regardless of the number of Exinda Appliances you are working with, it is very important to test EMC configuration on a single appliance before pushing the configuration onto other appliances.

Pull one appliance into the subgroup and push the configuration to the group. To ensure all the configuration and customizations are working as you expect, log on to the appliance and check the Optimizer Policy Tree.

If there is a problem with the policy, it is easier to back out a single appliance rather than an entire appliance group. After you confirm the configuration, add the rest of the appliances to the appliance group and then push the configuration to the group again.

## Appliance Group Inheritance

When working with device subgroups, remember to plan for group inheritance. Implement common configuration at the parent group level because all subgroups can inherit settings from the parent group.

Plan to support a common Optimizer Policy Tree that can also provide distinct network objects or applications per group, as appliance groups cannot inherit applications and network objects from their parent groups.

## Circuit Size

When configuring circuits and bridges, remember that you might need multiple circuits with different sizes to monitor various circuit types.

For example, if each link has a different Internet speed, you will have to create a different circuit for each link. Use the Library to create the different circuits for the different links, which are then re-usable when duplicating policy trees with only minor changes in link speeds.

## What to Configure in the EMC vs. the appliance

The following table outlines the configuration items you can only manage in the EMC and the items you can only update on individual appliances.

### NOTE

If you make changes to VLANs or protocols at the group level, you should send CLI commands from within EMC only.

| Configuration Item | Configure on EMC | Configure on Appliance |
|--------------------|------------------|------------------------|
| Optimizer          | ✔                |                        |
| Network Objects    | ✔                |                        |
| Users and Groups   |                  | ✔                      |
| VLANs              | ✔                |                        |
| Protocols          |                  | ✔                      |
| Applications       | ✔                |                        |
| Application Groups | ✔                |                        |
| Schedules          | ✔                |                        |
| Adaptive Response  |                  | ✔                      |

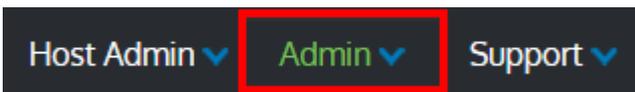
| Configuration Item | Configure on EMC | Configure on Appliance |
|--------------------|------------------|------------------------|
| Service Levels     |                  |                        |
| HTML Response      |                  |                        |

## 0.9.6 Configuring the EMC

To configure the Exinda Management Center to communicate with the Exinda Network Orchestrator appliances, follow the workflow below. After you complete the steps, you are ready to create policy and send it to your appliance groups.

### Step 1: Identify the SDP Location on the Exinda Management Center (if forwarding data to an SDP server).

Configure the location of your SDP so that data from the appliances is forwarded to this SDP.



At the top right of the interface, click **Admin > SDP Location** and specify the location of your SDP. The "SDP Location" configuration pop-up opens.

## SDP Location

**Location**

Enter the URL location of your SDP. The appliances will call into Exinda Management Center and the traffic data will be forwarded to SDP for reporting purposes.

SDP Host

Security  HTTPS

If using SDP on premise with HTTPS enabled, the SDP server must have a valid certificate installed for proper integration with Exinda Management Center.

All tenants will be associated with the same SDP.

### Step 2: Configure Administrator Email Settings

The mail server is used to send emails when a user needs to use the Forgot Password functionality.

At the top right of the interface, click **Admin > SMTP Server Settings** and specify the location of your Mail Server settings.

## Settings

▼ Email Server Settings

Configure SMTP settings to allow Exinda Management Center to send emails for Forgot Password and notifications

SMTP Host Name

SMTP Server Port

From Address

➤ Email Server Authentication

➤ Public Hostname

### Step 3: Configure your appliances to communicate with Exinda Management Center

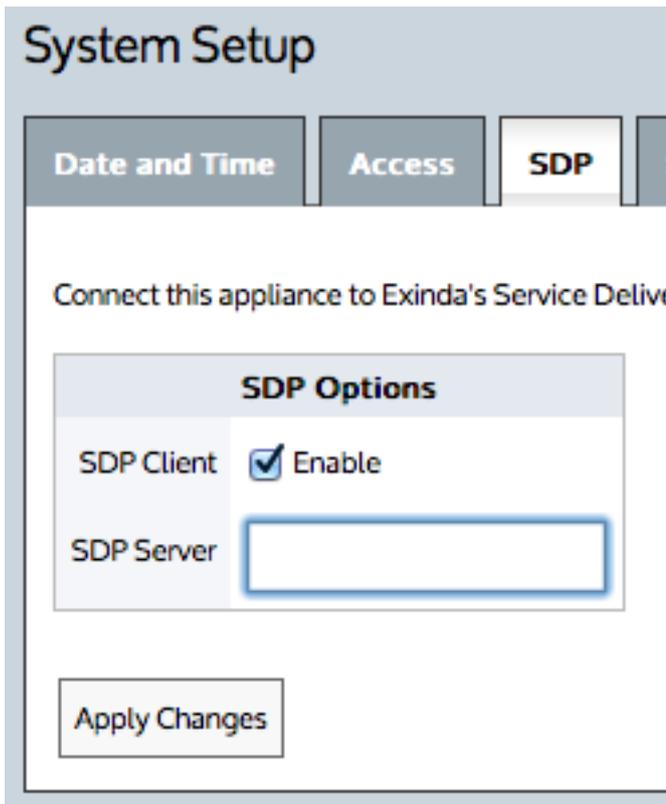
Exinda Network Orchestrator appliances need to know the location of the Exinda Management Center before they can communicate with it.

To allow an Exinda Appliance to communicate with the EMC, in the Exinda Web UI use the **SDP Options** to enable and set the location of the EMC system. The data collected by the Exinda Appliance is then forwarded to SDP for report generation. If you have an account in SDP, you can view and create reports there.

#### IMPORTANT

- » Only Exinda Appliances running firmware version 7.0.2 Update 1 or later support proper communication with the Exinda Management Center.
- » After the appliances have been configured to communicate with the EMC, sending configuration or performing firmware updates will not work from SDP. To restore this functionality in SDP, you need to reconfigure your appliances to communicate with SDP.

This step requires configuration on the Network Orchestrator appliances.



Screenshot 47: SDP settings

To configure the appliances to communicate with Exinda Management Center:

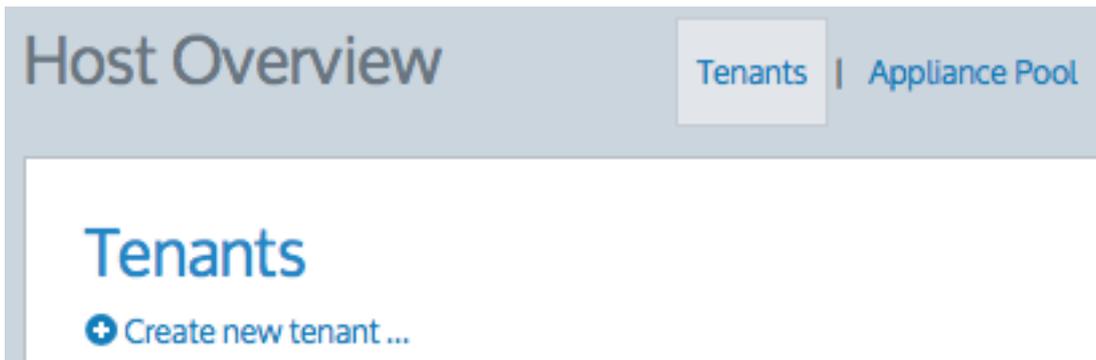
1. In the Exinda Web UI, go to **Configuration > System > Setup > SDP**.
2. Set the **SDP Server** to the location of your EMC.
  - If using the Exinda-hosted EMC, type `oc.exinda.com`.
  - If using your on-premises EMC, type the location of your EMC instance.
2. Select the **SDP Client Enable** checkbox.
3. Click **Apply Changes**.

#### Step 4: Add Tenants

##### NOTE

This screen appears only if using an on-premises instance of EMC.

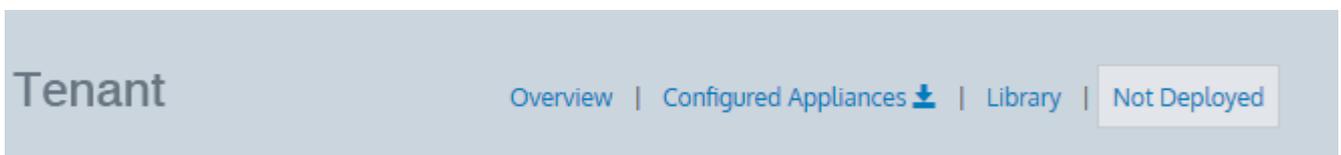
In the EMC web UI, click **Create new tenant** at the top of the tenant tree and specify the name of your tenant.



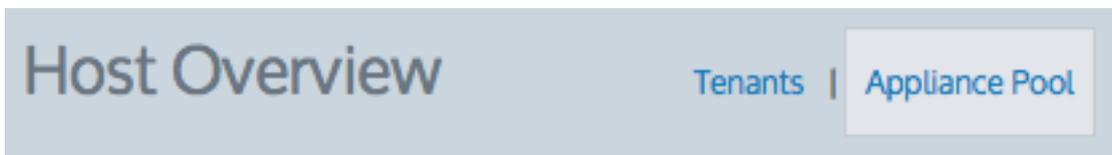
After you have created the tenant, click the tenant in the list to start managing the appliances in the tenancy.

#### Step 5: Add Appliances to Tenants

Wait for your appliance(s) to call in. If using the Exinda-hosted service or an on-premises instance with a single tenant, the appliance appears in the Not Deployed group.

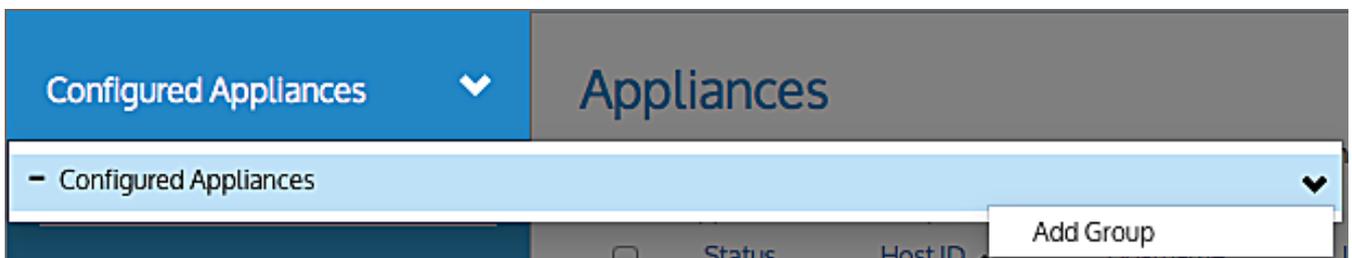


If you are using a multi-tenant on-premises instance, the appliance will appear in the **Appliance Pool**. Move the appliance from the **Appliance Pool** to **Unallocated** under the appropriate tenancy.



#### Step 6: Create Appliance Groups within a Tenant (optional)

Create an appliance group hierarchy under Configured Appliances. Appliances can be added to these groups. All appliances under the same group will receive the same configuration. Groups can be created hierarchically. Go to the Configured Appliances area. Click the drop down caret on the blue menu heading. Click the drop down caret on the desired appliance group and select **Add Group**. For more information, refer to [Managing Exinda Appliances with EMC](#) (page 419).



Move the appliances from the Not Deployed group to the Configured Appliances group (or one of the appliance groups that you created). Only appliances in a configured appliances group can be configured by Exinda Management Center. Select an appliance in the **Not Deployed** list and click **Move Appliances**. For more information, refer to [Managing Exinda Appliances with EMC](#) (page 419).

# Appliances

Move appliances into groups to manage and monitor appliances similarly.

[Move Appliances](#)



Status



Online

## 0.9.7 Importing appliance configuration

Although not compulsory, it is highly recommended that you import the available appliance configuration into the tenant library to reduce redundancy and avoid re-programming the same configuration in the EMC. The configuration import is not required if the appliance is new and does not contain any configuration.

### NOTE

Importing configuration is optional, but if an appliance has already been in use, its configuration can be applied globally across all other appliances instead of configuring the appliances individually.

Follow these steps to import configuration into a tenant:

When importing the configuration from an appliance, you work your way through a wizard that allows you to select the configuration items you need to import. There are eight classes of configuration that you can import, with each offering the configuration items that already exist on the appliance.

1. Go to **Configured Appliances > Appliances**.
2. Select an appliance and click **Import Configuration**.

The screenshot shows the 'Appliances' page in the EMC Exinda Network Orchestrator. The page title is 'Appliances' and the subtitle is 'Move appliances into groups to manage and monitor appliances similarly.' The page has a navigation bar with 'Overview', 'Configured Appliances', 'Library', and 'Not Deployed'. The 'Configured Appliances' menu is expanded, showing 'Appliances', 'Optimizer Policy Tree', and 'Bridge/Circuit Type Mapping'. The 'Appliances' section has three buttons: 'Move Appliances', 'Import Configuration' (highlighted with a red box), and 'Upgrade firmware'. Below the buttons is a table with columns for 'Status', 'Host ID', and 'Hostname'. The table has one row with a checked checkbox, a green circle, the text 'Online', the Host ID '0024e83dcaed', and the Hostname 'exinda-Riz-122'.

| Status                              | Host ID      | Hostname       |
|-------------------------------------|--------------|----------------|
| <input checked="" type="checkbox"/> | 0024e83dcaed | exinda-Riz-122 |

### Importing network objects

To import your network objects:

1. Click **Import Network Objects**, or click **Next** to skip this step.

Tenant Overview | **Configured Appliances** | Library | Not Deployed

**Configured Appliances** [Download] [Dropdown]

- Appliances
- Optimizer Policy Tree
- Bridge/Circuit Type Mapping
- Applications
- Application Groups

## Import Configuration

Network objects from the configuration file can be imported into the library.

The system will warn when a network object is not available for import or is already in the library.

### Step 1 - Network Objects

**Import Network Objects** Next Close

**NOTE**

If any network objects already exist in the library, a green checkmark appears in front of it. Once used, you will not be able to reuse it.

2. Click **Add Selected Network Objects to the Library**.

Tenant Overview | **Configured Appliances** | Library | Not Deployed

**Configured Appliances** [Download] [Dropdown]

- Appliances
- Optimizer Policy Tree
- Bridge/Circuit Type Mapping ⚠
- Applications
- Application Groups
- Network Objects
- Local Network Objects
- Application Performance Scores
- Service Level Agreements
- Configuration via CLI

## Import Configuration

Network objects from the configuration file can be imported into the library.

The system will warn when a network object is not available for import or is already in the library.

### Step 1 - Network Objects

| <input type="checkbox"/> | Name              |
|--------------------------|-------------------|
| <input type="checkbox"/> | www               |
| <input type="checkbox"/> | ✓ Rizwan-Sche-245 |
| <input type="checkbox"/> | ✓ Rizwan          |

**Add Selected Network Objects to the Library** Next Close

**Importing applications**

To import any applications that exist in the configuration of the appliance:

1. Click **Import Applications**, or click **Next** to skip this step.
2. Use check-boxes to select the applications you need to import.
3. Click **Add Selected Application to Library**.
4. Click **Next**.

#### **Importing schedules**

To import any schedules that exist in the configuration of the appliance:

1. Click **Import Schedules**, or click **Next** to skip this step.
2. Use check-boxes to select the schedules you need to import.
3. Click **Add Selected Schedules to Library**.
4. Click **Next**.

#### **Importing VLANs**

To import any VLANs that exist in the configuration of the appliance:

1. Click **Import VLANs**, or click **Next** to skip this step.
2. Use check-boxes to select the VLANs you need to import.
3. Click **Add Selected VLANs to Library**.
4. Click **Next**.

#### **Importing circuits**

To import any circuits that exist in the configuration of the appliance:

1. Click **Import Circuits**, or click **Next** to skip this step.
2. Use check-boxes to select the circuits you need to import.
3. Click **Add Selected Circuits to Library**.
4. Click **Next**.

#### **Importing virtual circuits**

To import any virtual circuits that exist in the configuration of the appliance:

1. Click **Import Virtual Circuits**, or click **Next** to skip this step.
2. Use check-boxes to select the virtual circuits you need to import.
3. Click **Add Selected Virtual Circuits to Library**.
4. Click **Next**.

#### **Importing policies**

To import any policies that exist in the configuration of the appliance:

1. Click **Import Policies**, or click **Next** to skip this step.

Tenant Overview | **Configured Appliances** | Library | Not Deployed

**Configured Appliances**  

Appliances

Optimizer Policy Tree

Bridge/Circuit Type Mapping

Applications

Application Groups

## Import Configuration

Policies from the configuration file can be imported into the library.

The system will warn when a policy is not available for import or already exists in the library.

Step 2 - Policies

[Previous](#) **[Import Policies](#)** [Close](#)

2. Use check-boxes to select the virtual circuits you need to import.
3. Click **Add Selected Policies to Library**.
4. Click **Next**.
5. Click **Import Policies** to import the current policies from the appliance.

Tenant Overview | **Configured Appliances** | Library | Not Deployed

**Configured Appliances**  

Appliances

Optimizer Policy Tree

Bridge/Circuit Type Mapping

Applications

Application Groups

## Import Configuration

Policies from the configuration file can be imported into the library.

The system will warn when a policy is not available for import or already exists in the library.

Step 2 - Policies

[Previous](#) **[Import Policies](#)** [Close](#)

### NOTE

EMC does not import policies that already exist in library or policies linked to a network object that does not exist in the library. Hover over the error icon to see the related error message.

6. Click **Add Selected Policies to the Library** to successfully add selected policies into the library:

Configured Appliances  

---

Appliances

---

Optimizer Policy Tree

Bridge/Circuit Type Mapping 

---

Applications

Application Groups

Network Objects

Local Network Objects

Application Performance Scores

Service Level Agreements

---

Configuration via CLI

## Import Configuration

Policies from the configuration file can be imported into the library. The system will s

The system will warn when a policy is not available for import or already exists in the

### Step 2 - Policies

| <input type="checkbox"/> | Name                                                                                             |
|--------------------------|--------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> |  A (Optimize: ) |
| <input type="checkbox"/> |  B (Optimize: ) |
| <input type="checkbox"/> | Everything (Optimize: 128 kbps-2048 kbps, Priority 1, Accelerated)                               |
| <input type="checkbox"/> | FTP (Optimize: 2000 kbps-8500 kbps, Priority 1, Accelerated)                                     |
| <input type="checkbox"/> | HTTP (Optimize: 3000 kbps-8000 kbps, Priority 1, Accelerated)                                    |
| <input type="checkbox"/> | RizDesktop (Optimize: 8000 kbps-100%, Priority 1, Accelerated)                                   |

### Importing service level agreements

To import any service level agreements that exist in the configuration of the appliance:

1. Click **Import Service Level Agreements**, or click **Next** to skip this step.
2. Use check-boxes to select the virtual circuits you need to import.

Tenant Overview | **Configured Appliances** | Library | Not Deployed

**Configured Appliances** ▼

Appliances

Optimizer Policy Tree

Bridge/Circuit Type Mapping ⚠

---

Applications

Application Groups

Network Objects

Local Network Objects

Application Performance Scores

Service Level Agreements

---

Configuration via CLI

## Import Configuration

Service level agreements from the configuration file can be imported into the library. The system will warn when a service level agreement already exists in the library, has the same definition.

The system will warn when a service level agreement already exists in the library, has the same definition.

### Step 8 - Service Level Agreements

| <input type="checkbox"/>            | Name                                                                                                            |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| <input checked="" type="checkbox"/> | Service Level Agreement name (Ping 10.10.9.55 with ping size of 1024 bytes, allowing for 5 ms of delay. Ping    |
| <input checked="" type="checkbox"/> | Service Level Agreement name 2 (Ping 10.10.9.195 with ping size of 1024 bytes, allowing for 5 ms of delay. Ping |
| <input checked="" type="checkbox"/> | XOXOX (Ping 10.10.9.55 with ping size of 1024 bytes, allowing for 5 ms of delay. Ping                           |
| <input checked="" type="checkbox"/> | YOYOYO (Ping 10.10.9.195 with ping size of 1024 bytes, allowing for 30 ms of delay. Ping                        |

Previous
Add Selected Service Level Agreements to the Library
Close

3. Click **Add Selected Service Level Agreements to Library**.

4. Click **Close**.

For each of these classes, you can select from the configuration items that exist, or you can skip to the next class. The process of importing each of the configuration classes is the same.

3. In the wizard, click the **Import <configuration class>** button. For the configuration class, if there are existing configuration items, they appear in a grid

4. If configuration items do not exist in the tenant library, select their check-boxes and click **Add Selected <configuration items> to the Library**.

#### NOTE

If configuration items already exist in the tenant library, a green check mark appears before the configuration item name. You cannot use the same configuration item again.

4. Click **Next** to move on to the next configuration class.

5. Repeat steps 3 and 4 for each configuration class.

6. At the end of the wizard, click **Close**.

## 0.9.8 Configuring an appliance manually

If your appliance is new and requires configuration, you can follow the steps below after the [basic configuration](#).

### Step 1: Configure the local network object for each appliance

The system uses the local network object to determine the location of all other network objects. The location of network objects is used to determine whether hosts and users are internal or external to the LAN behind your Network Orchestrator appliance. In the Configured Appliances area, select Local Network Objects in the menu. Click the appliance

where you want to configure the local network object. For more information, refer to [Configuring local network objects in the EMC](#) (page 436).

### Step 2: Configure your network objects (Optional)

Network objects are used for identifying the traffic affected by the policy as well as for monitoring traffic. Consider the following:

- » If you want to use a network object in the creation of a virtual circuit or policy, create a network object in the library. If you imported your network objects from your appliance, you may not need to create any network objects.
- » If you want to send a network object to the appliance for monitoring purposes (not for policy creation), then create a network object in the configured appliance group, which also adds the network object to the library. Or create a network object in the library then add it to the configured appliances group.
- » If you use a network object in the definition of a virtual circuit or policy for a given appliance group, then the network object is automatically added to the appliance group configuration. For more information, refer to [Creating network objects in the EMC](#) (page 433).

### Step 3: Configure custom applications (Optional)

Custom applications can be used for both identifying traffic affected by the policy, and for monitoring traffic. Consider the following:

- » If you want to use a custom application in the creation of your virtual circuit or policy, create a custom application in the library. When you create or edit the Virtual Circuits or policies, your custom application becomes available for selection.
- » If you want to send a custom application to the appliance for monitoring purposes, you need to add the application to a monitored application group.
- » If you use a custom application in the definition of a virtual circuit or policy for a given appliance group, then the custom application is automatically added to the appliance group configuration. For more information, refer to [Configuring applications with the EMC](#) (page 440).

### Step 4: Configure schedules (Optional)

Schedules can be used to specify when policies or virtual circuits takes effect. If you want to use a schedule in the creation of your virtual circuit or policy, first create a schedule in the library. When you create or edit the virtual circuits or policies, your schedule can then be selected and automatically added to the appliance group configuration. For more information, refer to [Configuring schedules in the EMC](#) (page 446).

### Step 5: Configure the Optimizer Policy Tree on the Configured Appliances group (or on one of the appliance groups that you created)

Each object element of the Policy Tree is stored in the library for use by other appliance groups or other areas within the Optimizer Policy Trees. For more information, refer to [Optimizer Policy Tree](#) (page 170).

#### IMPORTANT

When appliances are moved out of the Configured Appliances group to the Unallocated Appliances group or the Appliance Pool, the configuration that was applied using the edit forms is automatically removed from the appliances upon the next call into the Exinda Management Center.

### Step 6: Create named Circuit Types

The circuit type specifies the intended use of a circuit, for example, "Internet", "Voice", or "MPLS". Circuits and appliance

bridges are mapped to these circuit types. The circuit is sent to the appliances that have bridges mapped to the same circuit type as the circuit. Go to the library and create your required circuit types. For more information, refer to [Circuit types in the EMC](#) (page 423).

### Step 7: Map the appliance bridges to the Circuit Types

This determines the bridges to which the circuits are mapped on the appliances. Go to the configured appliances. Click Bridges/Circuit Type Mapping, and for each appliance specify the Circuit Type for each bridge. For more information, refer to [Working with bridge/circuit type mapping](#) (page 168).

### Step 8: Create a circuit

Circuits specify the physical connections to the WAN or Internet. Create a circuit in the Optimizer Policy tree for the configured appliance group. Or create a circuit in the library and add it to the Optimizer Policy tree for the configured appliance group. For more information, refer to [Circuits](#) (page 424).

### Step 9: Create a virtual circuit

Virtual circuits define what traffic is processed in a partition and how much bandwidth it is allowed. Create a virtual circuit in the Optimizer Policy tree for the configured appliance group. Or create a virtual circuit in the library and add it to the Optimizer Policy tree for the configured appliance group. For more information, refer to [Virtual Circuits](#) (page 426).

### Step 10: Add a policy set

Policy sets are groups of policies that will be added to the Virtual Circuits. There are pre-created policy sets in the library that correspond to the sets that result from running the wizard on the appliance.

Add a policy set to a configured appliance group Optimizer Policy tree virtual circuit. You can also create policy sets from the Optimizer Policy tree or in the library. For more information, refer to [Creating policy sets in the EMC](#) (page 432).

### Step 11: Add or edit a policy (Optional)

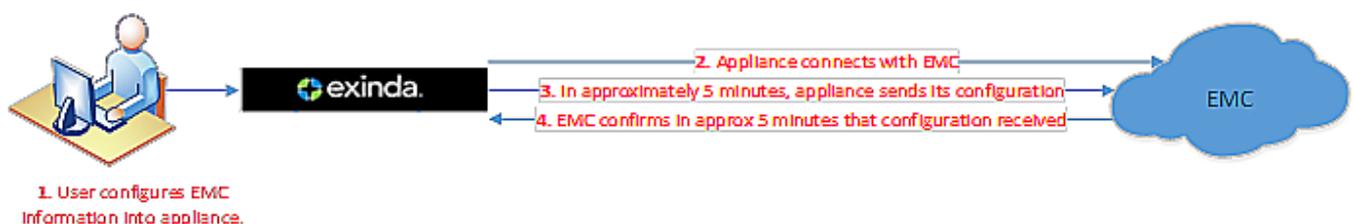
Policies are the rules that control the traffic. When adding or editing a policy set, you can add or edit a policy. For more information, refer to [Policies](#) (page 430).

### Step 12 Send the Configuration to the Appliance Group

Changes are sent to the appliances within an Appliance Group only when you choose to send the configuration. For more information, refer to [Sending configuration changes to the appliances](#) (page 176).

After these EMC settings are configured, the appliance calls in to the EMC for the first time. For more information, refer to [Configuring the EMC](#) (page 155). There can be a delay of up to five minutes while the appliance sends its current configuration to the EMC. Once the full configuration is received, EMC confirms the configuration, which may also take up to five minutes. The Bridge/Circuit Type mapping is not available on the EMC until the appliance has been online for about ten minutes. During this period, the EMC displays a message indicating that it is waiting for the appliance to call in.

Once you move the appliance to a group, and you push configuration to the group, the appliance then receives the updated configuration when it next calls in to the EMC.



Screenshot 48: First communication between the appliance and EMC

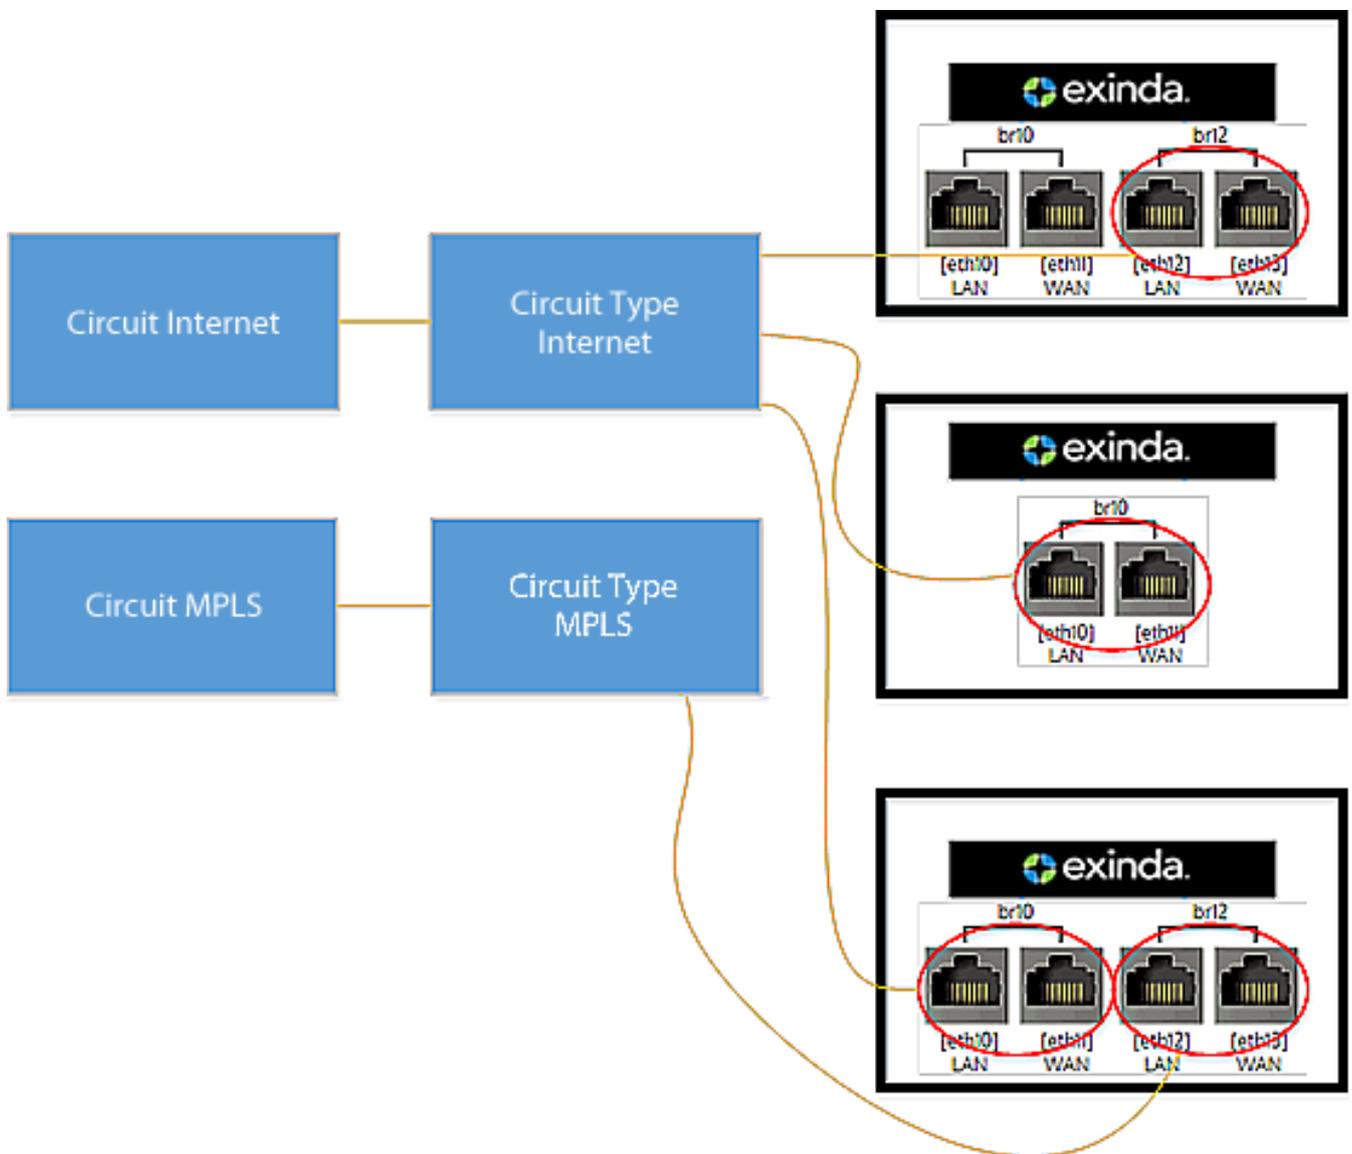
## 0.9.9 Configuring a bridge

It is important to understand the relationship between bridges, circuits, and circuit types before you start creating policy. Consider the following:

» **Circuit** – defines the physical connections to the WAN or Internet and the inbound and outbound bandwidth and the named circuit type. On the appliance, the circuit specifies which named bridge or bridges it is bound to.

» **Circuit Type** – an abstract concept that creates a virtual binding between the circuits and the appliance bridges. Circuit Types represent the intended use of a circuit. This allows you to configure a circuit for multiple appliances without requiring the bridges on the appliances to have the same name, such as br10. This is favorable where the number of bridges or names of bridges or the cabling of the bridges is not consistent across the appliances.

For an appliance to receive the Optimizer Policy Tree configuration rooted with a particular circuit, the bridge on the appliance must be mapped to the same Circuit Type as that Circuit. For example, if the circuit is bound to circuit type "Internet", and the appliance bridge(s) is mapped to 'Internet', then that circuit configuration is sent to that appliance bound to the specific bridges.



Screenshot 49: Diagram showing bridge configurations.

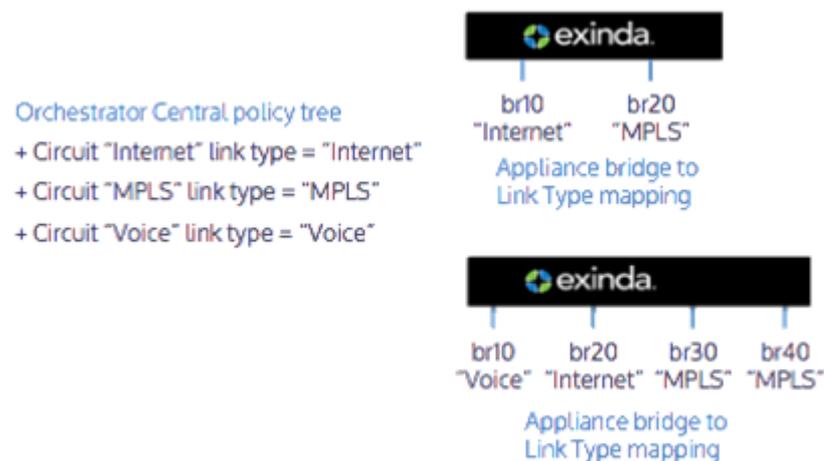
## Using library items to create a consistent definition and naming strategy

Current naming practices can make it complicated to track and understand the mappings between bridges, circuit types and circuits. Each bridge on an appliance is usually named with no relevancy to the purpose of the bridge. To simplify bridge configuration within the EMC, you should first focus on creating common Circuit Types and Circuits to create consistency throughout the appliances. By labeling circuits and circuit types within the library, and then mapping the library items to bridges, you can create an Optimizer Policy Tree that covers multiple scenarios.

For example, consider two different appliances with a different numbers of bridges and where they are cabled differently:

- » The first appliance has two bridges, br10 and br20, where br10 is mapped to the "Internet" Circuit Type and br20 is mapped to the "MPLS" Circuit Type.
- » The second appliance has four bridges, br10, br20, br30, and br40, where br10 is mapped to "Voice", br20 is mapped to "Internet", and br30 and br40 are mapped to "MPLS".

If you add the common Circuit Types used above to Library items, you can reuse the named Circuit Types when performing the bridge/circuit type mapping in the Optimizer Policy Tree.



Screenshot 50: Using naming strategies to simplify bridge configuration.

## Working with bridge/circuit type mapping

For an appliance to receive the Optimizer Policy Tree configuration rooted with a particular circuit, the bridge for an appliance must be mapped to the same Circuit Type as that Circuit. That is, if the circuit is bound to circuit type 'Internet' and the appliance bridge(s) is mapped to 'Internet', then the circuit configuration sent to that appliance is bound to the specific bridges.

The Bridge/Circuit Type Mapping list shows each appliance in the appliance group. You can edit the Bridge/Circuit Type Mapping for each appliance. A warning icon appears next to each appliance that does not have any of its bridges mapped. This warns that the appliance will not be sent any part of the Optimizer Policy Tree configuration. Also, the warning icon will be shown in the blue menu (on the left) next to the Bridge/Circuit Type Mapping menu item, if there are any appliances with a warning icon.

## Create and map the circuit types to the appliance's bridges

1. Go to **Configured Appliances > Bridge/Circuit Type Mapping**.
2. Click **Host ID** or **Bridge ID** of the appliance to bind the circuit type to the bridge.

## Bridge/Circuit Type Mapping

Map the circuit types to the appliance's bridges. Circuits are mapped to the circuit type. The circuit types allow the appliance bridges to be bound to circuits according to their circuit type.

| Host ID ▲    | Hostname       | IP Address  | Group  | ▲ | Bridge/Circuit Type Mapping |
|--------------|----------------|-------------|--------|---|-----------------------------|
| 0024e83dcaed | exinda-Riz-122 | 10.10.7.122 | Rizwan | ▲ | br10 ->                     |

3. You can either select the existing circuit types from the drop-down list or create a new circuit type.

## Appliance Bridge to Circuit Type Mapping

Map the circuit types to the appliance's bridges. Circuits are mapped to the circuit types. The circuit types allow the appliance bridges to be bound to circuits according to their circuit type.

▼ Appliance bridge to circuit type mapping

Host ID 0024e83dcaed

Hostname exinda-Riz-122

br10

[+ Create new circuit type in the library ...](#)

Screenshot 51: Appliance bridge to circuit type mapping.

It is recommended to provide the circuit type the same name as circuit, as it makes it easier to map them together.

### To edit an appliance's Bridge/Circuit Type Mapping

1. Click the **Host ID** or **Bridge/Circuit Type Mapping**.
2. For each bridge, select a circuit type from the drop-down list. If the desired named circuit type is not in the list, click **Create new circuit type in the library...** to create a new circuit type. Once created, it is then available in the drop-down list.
3. Click the **Update Mapping** button.

### Why does it say Pending in the Bridge/Circuit Type Mapping column?

Pending means that the Exinda Management Center has not received the list of bridges from the appliance yet. The first time the appliance calls in, the Exinda Management Center requests bridge information from the appliance. Note that in general, there should not be much time between the first communication and the second communication with the bridge information.

## NOTE

This may also occur if you are using an appliance with a firmware version prior to the 7.0.2 Update 1.

### What triggers a warning?

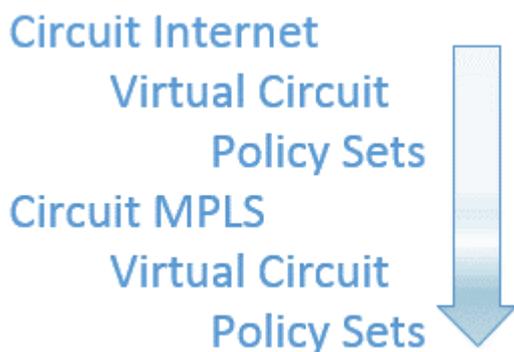
- » If none of the bridges are mapped to a circuit type, then a warning appears for the specific appliance.
- » If there is one or more appliances with a warning, then a warning appears next to the Bridge/Circuit Type Mapping menu item.

### What happens to the mappings when I move an appliance?

- » The bridge retains its circuit type mappings when moving between appliance groups.
- » If the appliance is moved to the **Not Deployed** area, then the circuit type mapping is removed.

## 0.9.10 Optimizer Policy Tree

All network behavior that you need to modify is specified by policies in the optimizer. This includes traffic shaping, prioritization, acceleration, and packet marking. These policies are arranged hierarchically in a tree so that you can assign different policy rules to different types of traffic on your network. The hierarchy consists of circuits, virtual circuits, policy sets, and policy rules. Note that policy sets are a concept within EMC only; they do not exist on Network Orchestrator appliances. The tree is processed in a top-to-bottom order, so the policies on traffic are applied accordingly.



### Policy sent to the appliance is dependent on Circuit Type

When the Optimizer Policy Tree is assigned to an appliance group, generally, all appliance groups that are nested under that group inherit the Optimizer Policy Tree. In which case, a message area above the Policy Tree indicates that the tree is inherited. If you do not want child appliance groups to inherit the tree, you can stop the inheritance.

All the appliances within a group will potentially receive that Policy Tree when the policy is sent. An appliance receives the Policy Tree configuration that corresponds to its bridge/circuit type mapping.

## EXAMPLE

Consider a Policy Tree with three circuits, Internet, MPLS, and Voice, where the circuits map to Circuit Types of the same name. For example, Internet maps to a circuit type named "Internet". Now consider an appliance (within this Policy Tree's appliance group) that has its two bridges mapped to only two of these circuits: br10 to Internet, and br20 to MPLS. When the configuration is sent to this appliance, the circuit "Internet" is mapped to the appliance's bridge br10 and the circuit "MPLS" is mapped to the appliance's bridge br20. The "Voice" circuit is not sent to the appliance as there is not a matching circuit type.

## Policy sets can be reused in multiple Virtual Circuits

Policy sets are a concept within the EMC only. You can create a named policy set then apply the policy set to multiple virtual circuits. All virtual circuits using the policy set then have exactly the same policy.

## Everything references a library item

Within a tenant, the EMC treats everything as a library item so that the configuration components can be reused. When you modify a library item, everywhere it is used is also affected. For example, when you create and reuse a virtual circuit, whenever that virtual circuit is changed, all instances of its use are also changed.

## Required objects will automatically be queued to be sent

When policy rules or virtual circuits use objects in their definitions, such as network objects or schedules, then those objects are automatically added to the configuration that must be sent to the appliances.

All other uses of the Policy Tree and its components are the same as on the appliance itself.

## Where do I find the Optimizer Policy tree?

Go to **Configured Appliances > Optimizer Policy Tree**.

The screenshot shows a web interface for a tenant. At the top, there is a navigation bar with the following items: "Tenant", "Overview", "Configured Appliances" (with a dropdown arrow), "Library", and "Not Deployed". Below this, a sidebar menu is visible with the following items: "Configured Appliances" (with a dropdown arrow), "Appliances", "Optimizer Policy Tree" (which is highlighted), and "Bridge/Circuit Type Mapping". The main content area displays the "Optimizer Policy Tree" title, followed by the text "The Optimizer policy tree defines what actions are taken on different types" and a link "Create new circuit ...".

## To add a circuit to the Policy Tree

1. Click **Create new circuit...**
2. Name the circuit, program the desired bandwidth and bind the circuit to the circuit type.

## Circuit

Define physical connections to the WAN/Internet.

You can bind different circuits to each bridge or you can treat all bridges as one combined circuit. Typically, one circuit would be created for each physical link. Ensure each bridge has a unique policy and will be monitored in a catch-all circuit.

Note that the order of the circuits is important.

|                                                             |        |
|-------------------------------------------------------------|--------|
| > Name                                                      |        |
| ▼ Bandwidth                                                 |        |
| Identify the inbound and outbound bandwidth of the circuit. |        |
| Inbound Bandwidth <input type="text"/>                      | kbps ▼ |
| Outbound Bandwidth <input type="text"/>                     | kbps ▼ |
| > Bind to Circuit Type                                      |        |

3. If a circuit exists in a library then you can add it from there. Click **Add Circuit from Library...** and select a desired circuit.

Tenant Overview | Configured Appliances ↓ | Library | Not Deployed

**Configured Appliances** ↓

- Appliances
- Optimizer Policy Tree
- Bridge/Circuit Type Mapping

### Optimizer Policy Tree

The Optimizer policy tree defines what actions are taken on different

Create new circuit ... **Add Circuit from Library ...**

## Optimizer Policy Tree

The Optimizer policy tree defines what actions are taken on different

Create new circuit ... | Select a circuit ▼

Internet

### To add a virtual circuit to the Policy Tree

Virtual circuits logically partition the circuit. A virtual circuit defines the traffic that is processed in the partition and the bandwidth it consumes. Each virtual circuit has its own set of policies.

1. Click **Create new virtual circuit...**

Tenant Overview | **Configured Appliances** ↓ | Library | Not Deployed

Configured Appliances ↓

Appliances

**Optimizer Policy Tree**

Bridge/Circuit Type Mapping

## Optimizer Policy Tree

The Optimizer policy tree defines what actions are taken on different

- Internet - Circuit (1024 kbps on circuit type 'Internet')

Create new virtual circuit ...

---

Create new circuit ...

2. Define the virtual circuit and click **Create and Add**.

## Virtual Circuit

Define how to logically partition the circuit. The virtual circuit defines what traffic will be processed in this partition and how much bandwidth it is allowed. Each virtual circuit will have its own set of policy rules.

▼ Name: Local<-->All

The name will be used to identify this virtual circuit when applying to policy trees. Keeping the Local Site name generic will be better for use across different appliances and different scenarios.

Auto-suggest the name

Name

Local Site

---

➤ Filter: Bi-directional to/from All

---

➤ Bandwidth

---

➤ Dynamic Virtual Circuit: Disabled

---

➤ Schedule: Always

---

Create

Cancel

Similar to circuits, virtual circuits can be added from the library, if present.

### To add a policy set to the Policy Tree

Policies define what actions are to be taken on different types of traffic.

1. There are two options:

a. Click **Create new policy set...** to create your own set of policies.

## Optimizer Policy Tree

The Optimizer policy tree defines what actions are taken on different types of traffic.

- Internet - Circuit (1024 kbps on circuit type 'Internet')

---

- ■ Local<-50%>->All - Virtual Circuit (50% in/out matching 'All')

**Create new policy set ...** Add Policy Set from Library ...

Create new virtual circuit ...

---

Create new circuit ...

b. Click **Add Policy Set from Library...** to select a pre-defined policy set template for a different type of traffic.

## Optimizer Policy Tree

The Optimizer policy tree defines what actions are taken on different types of traffic.

- Internet - Circuit (1024 kbps on circuit type 'Internet')

---

- ■ Local<-50%>->All - Virtual Circuit (50% in/out matching 'All')

Create new policy set ... | Select a policy set ▼

Create new virtual circuit ...

---

Create new circuit ...

- Internet inbound
- Internet outbound
- Monitor Only
- WAN inbound
- WAN inbound (Service Provider)
- WAN outbound
- WAN outbound (Service Provider)
- WAN outbound (with acceleration)

2. In this example, the Internet outbound policy set is selected and it automatically populates all the policies within this set into the virtual circuit:

# Optimizer Policy Tree

The Optimizer policy tree defines what actions are taken on different types of traffic. Each element in the

- Internet - Circuit (1024 kbps on circuit type 'Internet')

---

- Local<-50%->All - Virtual Circuit (50% in/out matching 'All')
  - Internet inbound - Policy Set
    - ⚙️ P2P - Choke 1%-3% (Optimize: 1%-3%, Priority 10)
    - ⚙️ Streaming - Limit Low 2%-10% (Optimize: 2%-10%, Priority 10)
    - ⚙️ Software Updates - Limit Med 3%-50% (Optimize: 3%-50%, Priority 9)
    - ⚙️ Voice - Guarantee Critical 15%-100% (Optimize: 15%-100%, Priority 1)
    - ⚙️ Thin Client - Guarantee High 10%-100% (Optimize: 10%-100%, Priority 3)
    - ⚙️ Files - Guarantee Med 8%-100% (Optimize: 8%-100%, Priority 5)
    - ⚙️ Web - Guarantee High 10%-100% (Optimize: 10%-100%, Priority 3)
    - ⚙️ Mail - Guarantee Med 8%-100% (Optimize: 8%-100%, Priority 5)
    - ⚙️ Unified Communications - Guarantee Med 8%-100% (Optimize: 8%-100%, Priority 3)
    - ⚙️ ALL - Guarantee Low 5%-100% (Optimize: 5%-100%, Priority 7)

---

- [Create new virtual circuit ...](#)

---

- [Create new circuit ...](#)

## To add a policy rule to the Policy Tree

1. Click the policy set to which you would like to add the policy.
2. In the policy set form, create a policy or add a policy from the library.
3. Click **Update in Library**.

## To reorder Virtual Circuits

In the Optimizer Policy Tree, drag and drop the virtual circuit to its new location.

## To reorder a policy rule

1. Click the policy set that contains the policy that you would like to reorder.
2. In the policy set form, drag and drop the policy to its new location.
3. Click **Update in Library**.

## To remove elements from the Policy Tree

For each element that you want to remove, click the **x** icon at the far right.

### NOTE

You cannot delete elements from the Policy Tree if you are looking at a sub appliance group that inherits the Policy Tree. You must edit the Policy Tree in the appliance group that defined the tree.

## To disinherit an appliance group from a Policy Tree

1. When an appliance group inherits an Optimizer Policy Tree, there is a banner across the top indicating that it is inherited.
2. Click **Stop inheriting Policy Tree**. The tree is removed and you can start building up another Policy Tree.

### 0.9.11 Sending configuration changes to the appliances

Changes that you make in the Exinda Management Center are not sent to the appliances until you choose to send them. To do so, click the **Send Configuration** icon  next to the appliance group. This sends the configuration to all the appliances in the group and any nested appliance groups when the next time each appliance calls into the system.



When sending the configuration, you have the option to restart the optimizer on the appliances or save the configuration on the appliances. The next time the appliances call in, they receive the configuration, restart the optimizer, and save configuration as instructed.

If multiple appliance groups exist, then you can push the configuration individually for an appliances group by clicking on the main **Configured Appliances**  icon. This pushes the configuration to all the appliances groups and appliances within them. Pushing the configuration restarts the Optimizer and saves the configuration on the appliance.

#### When sending the configuration:

- Restart the Optimizer on the appliances
- Save the running configuration on the appliances

**Send Configuration** **Cancel**

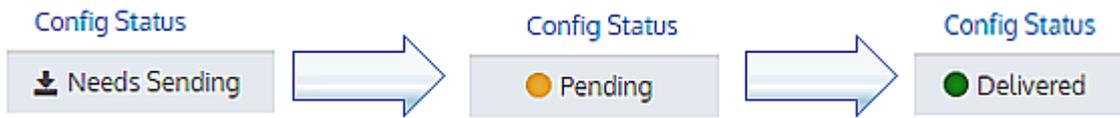
The first user-initiated push of the configuration to the appliance(s) deletes the following configuration on the appliance:

- » The definitions for local and private net network objects.
- » The Circuits in the Optimizer (and hence the entire optimizer tree).
- » The **After Work** and **Work Hours** schedules.

#### NOTE

The configuration is sent only to the appliances belonging to that appliance group, not to all the groups within a tenant.

When the configuration is being edited, the configuration status is **Needs Sending**. When the configuration is pushed, the status changes to **Pending**, and when the appliance receives the configuration, the status changes to **Delivered**.



Screenshot 52: The life cycle of configuration status

# Using

This topic focuses on the day-to-day use of your Exinda Appliance like defining traffic policies, setting alerts, monitoring performance, monitoring traffic, and understanding solutions and recommendations.

---

---

## 0.10 Defining a network environment

One of the first things you do after connecting an Exinda to your network, is define how the Exinda sees your network and its components.

As an analogy, imagine walking around your office or data center and placing sticky notes on all the servers, cables and racks to identify them and note what they do. That's essentially what you do when you define objects in Exinda.

There are a variety of object types available, representing almost every physical, virtual and logical network component in your environment.

### 0.10.1 Adding network objects

Network objects represent hosts on a network and can include subnets, single hosts, or groups of both. Once defined, a network object may be used throughout the Exinda Appliance for monitoring and identifying which traffic should be processed in the policy engine.

Network objects are in the configurations of other objects, such as applications, adaptive response rules, application performance score objects, and application performance metric objects.

Network objects are also used to determine which traffic is considered inbound to your network and which traffic is outbound.

The location of a network object determines the direction of traffic. If one end of the conversation is defined in an external network object and the other is defined in an internal network object, then traffic from an external network object to an internal network object is considered inbound traffic.

Conversely, traffic from an internal network object to an external network object is considered outbound traffic.

You can indicate whether you want to report on the traffic relative to the network object, that is chart the traffic in and out of a given network object.

#### Adding network objects in the Exinda Web UI

By checking the Subnet Report checkbox, the data for the network object will be shown on the subnet monitor page. This setting only affects the display of the data. The data will be collected regardless of this setting.

Some network objects are automatically created by the appliance: **ALL**, **private net** and **local**

- » **All**— Represents all traffic on the network. When used in Optimizer Policies, it matches all traffic. This network object is not editable and cannot be deleted.
- » **private net**— Represents all possible non-routable, private IP addresses.
- » **local**— Created when an IP address is assigned to one or more bridge interfaces. The object contains the IP address and subnet mask of each bridge interface.

Add New Network Object

Screenshot 53: Adding a new network object.

Where to configure it

Go to **Configuration > Objects > Network Object > Network Objects**.

To create a new network object

1. Specify a name for the network object.
2. Select the location of the network object - internal, external, or inherit. Packets are matched to a network object, and the closest subnet within that network object determines the location. See examples below.

- **Internal**— All subnets and hosts defined by the network object will be considered on the LAN side of the appliance.
- **External**— All subnets and hosts defined by the network object will be considered to be on the WAN side of the appliance.
- **Inherit**— The locations of the subnets and hosts defined by the network object is determined or inherited by closest match to other network objects.
  - If all the subnets in this network object are contained in other network objects that are internal, then the location of this network object will inherit the internal location.
  - Similarly, if all the subnets in this network objects are contained in other network objects that are external, then the location of this network object will inherit the external location.
  - If some subnets in this network object are contained in other internal network objects and some are contained in other external network objects, then the location of this network object will be mixed.
  - If no network objects match, then the location defaults to external.

**NOTE**

When creating network objects that have location set to "inherit", you can use the CLI command `show network-object <name>` to show the location.

3. Select whether the traffic for this network object should be shown on the Subnet reports. See For more information, refer to [Monitoring subnets](#) (page 291)..
4. Specify the network IP address and netmask length of the subnet. IPv4 and IPv6 addresses are accepted. Although only four lines for IP addresses are displayed for a new object, add more IP addresses by saving the network object and click **Edit** to be presented with an extra 4 lines.
5. Click **Add new Network Object**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.

## Examples of network object definitions

### EXAMPLE – Network object defining two internal proxy servers

Create a network object that defines two internal proxy servers, 192.168.1.10 and 192.168.1.11:

```
Name: Web Proxies  
Location: Internal  
Subnets: 192.168.1.10 /32  
Subnets: 192.168.1.11 /32
```

### EXAMPLE – Head office defining a network object for a remote branch

Create a network object that defines the Head Office location, that has a subnet 10.0.100.0/24, where this Exinda appliance is NOT deployed:

```
Name: Head Office  
Location: External  
Subnets: 10.0.100.0 /24
```

### EXAMPLE – Network object defining an internal IPv6 server

Create a network object that defines the internal IPv6 server at 2001:db8::1234:5678

```
Name: FileServer6  
Location: Internal  
Subnets: 2001:db8::1234:5678 /128
```

### EXAMPLE - Network object with inherited location

Define three network objects as follows:

```
Name: HQ Subnets: 10.0.0.0/8 Location: External  
Name: Office-A Subnets: 10.0.1.0/24 Location: Internal  
Name: Server-1 Subnets: 10.0.1.200/32 Location: Inherit
```

Subnets are matched by decreasing netmask length. The Server-1 network object 10.0.1.200 will be internal, as it most closely matches the Office-A Network Object which is internal. Since the Server-1 Network Object contains a single subnet that can be matched to Office-A, its location is shown as internal.

What about internal-to-internal or external-to-external traffic?

When the **Ignore Internal-to-Internal** option is set on the Monitoring configuration page, all traffic between network objects marked as internal is ignored and passed through the Exinda appliance unaffected. See For more information, refer to [Monitoring Configuration](#) (page 621)..

How to know whether a network object with location 'inherit' resolved to internal or external location?

You can use the CLI command to see what location the network object resolved to:

```
show network-object <name>
```

Creating Network Objects based on FQDN?

It is possible to configure network objects using fully-qualified-domain-names instead of IP addresses. Should it later become necessary to change network settings on application servers, the Exinda appliance can then automatically detect the change through DNS.

To configure a network object based on a fully qualified domain name, use the following commands:

```
>en
```

```
#conf t
```

```
(config) # network-object <NAME> fqdn <fully qualified domain name>
```

An Exinda appliance must be configured with a DNS server if it is to perform name resolution using FQDN. Each record retrieved has a life cycle equal to the TTL (Time to live) defined for such a record. When the TTL is exceeded, Exinda automatically refreshes the record to verify that the IP address has not changed. When appliance reboots occur, or changes to DNS configuration, interface configurations, or link states on any interface, this causes an automatic refresh of the network object. Should you need to perform a refresh, you can use the following command:

```
(config) # network-object <NAME> refresh
```

When the TTL is lower than 5 minutes, Exinda waits the full five minutes before attempting a refresh in order to avoid DNS flooding.

#### **CAUTION**

Please be aware that if using a cluster of Exinda Appliances, the resolution for a given FQDN-Network Object, among the different appliances, might be unexpected if any of the following conditions are met:

- » If the cluster members are using different DNS servers
- » If the name being resolved returns a different set of addresses for each resolve (i.e. 'load balancing')

## Adding network objects in the EMC

A network object can be created in the Library for later use in other components and appliance groups, or it can be created directly in the Optimizer Policy Tree, which also saves it to the Library. You can also import network objects into the Library from an appliance.

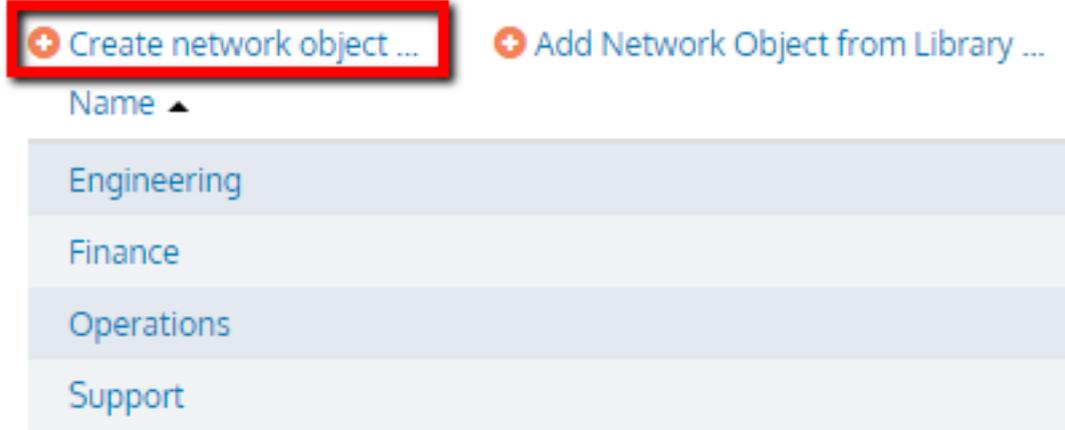
### **To create and apply a network object**

To create and apply a network object:

1. Go to **Library > Network Objects**.
2. Click **Create new network object in the library....**

# Network Objects

Define network objects to represent subsets of your network, which can



[+ Create network object ...](#) [+ Add Network Object from Library ...](#)

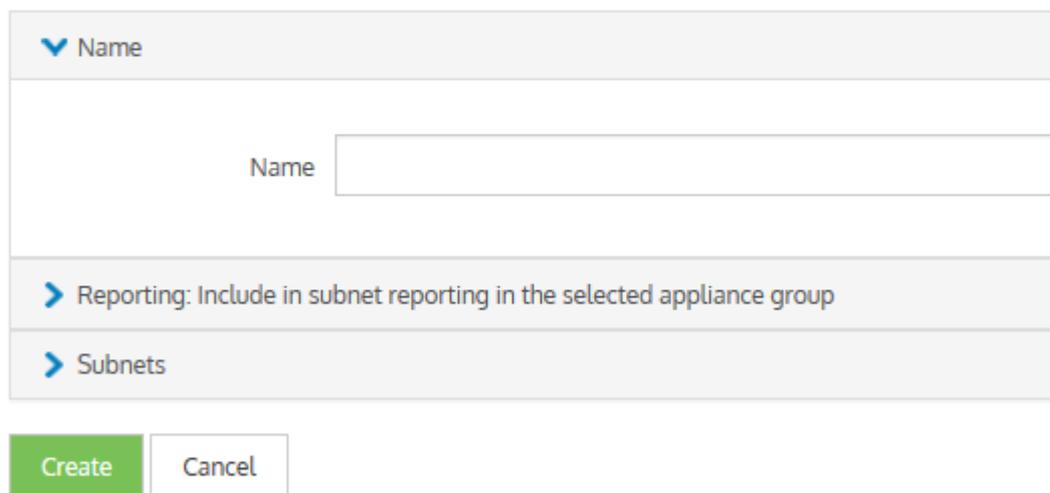
Name ▲

- Engineering
- Finance
- Operations
- Support

3. Key-in a **Name** for the object.

## Network Object

Define network objects to represent subsets of your network, which can include multiple subnets :



▼ Name

Name

▶ Reporting: Include in subnet reporting in the selected appliance group

▶ Subnets

4. In the **Subnets** section, define subnets to include in this network object by **IP Network Address** and **Mask Length** or by **Fully Qualified Domain Name (FQDN)**.

## Network Object

Define network objects to represent subsets of your network, which can include multiple s

> Name

> Reporting: Include in subnet reporting in the selected appliance group

▼ Subnets

Define which subnets to include in this network object.

IP Network Address / Mask Length

/  ✕

+ Add Another Subnet

Create Cancel

5. Click **Create**. Repeat from step 1, if you want to create another network object.
6. Go to **Configured Appliances > Network Objects**.
7. Click **Add network object from Library...**

## Network Objects

Define network objects to represent subsets of your network, which can include multiple

+ Create network object ... + Add network object from library ...

Name ▲

|                     |
|---------------------|
| Network Object      |
| Test Network Object |

Click here

8. Select one or more network objects to add and click **Add Network Object to 'Configured Appliances'**.
9. To define whether to include monitor information of network objects in subnet reporting, edit the required network objects and check **Include in subnet reporting** option, within the **Reporting** section.

# Network Objects

Define network objects to represent subsets of your network, which can include multiple subnets and multiple hosts. The network objects can then be used to monitor traffic or to configure traffic policy.

[+ Create network object ...](#)

[+ Add Network Object from Library ...](#)

| Name ▲      | IP Network Address | Monitoring               |
|-------------|--------------------|--------------------------|
| Engineering | 10.10.1.0/24       | ✓                        |
| Finance     | 10.11.0.0/16       | ✓                        |
| Marketing   | 10.50.1.0/24       | <input type="checkbox"/> |
| Operations  | 10.30.0.0/16       | ✓                        |
| Support     | 10.10.10.5/32      | ✓                        |

## To use a network object in a policy rule definition

You can apply a network object filter to a policy to filter by subnet (communicating with other specific hosts or subnets). Use the instructions above to create the Network Object in the library, and then refer to [Policy Library](#) for further instructions. The Network Objects saved as Library items appear in the **Filter** section, under the **Source** list.

## To use a network object in a virtual circuit definition

When defining virtual circuits to partition a circuit, you can apply Network Object library items to filter the traffic by subnet. Use the instructions above to create the Network Object in the library, and then refer to [Virtual Circuits](#) for further instructions.

## To use a network object in an Application

When defining applications to classify traffic, you can apply Network Object library items to classify traffic based on a combination of Network Object, TCP Port, UDP Port, DSCP, and Protocols. Use the instructions above to create the Network Object in the library, and then refer to [Applications](#) for further instructions.

## Configuring local network objects in the EMC

Local network objects define which part of the network is considered as the local area network relative to an appliance.

Local network objects take their definition from an IP network address and mask length to identify the range of IP addresses that exist in the LAN behind the appliance. You can define a local network object for each appliance, so that each appliance can differentiate between traffic that is external and internal to the LAN on which it operates. The local network objects cannot be deleted.

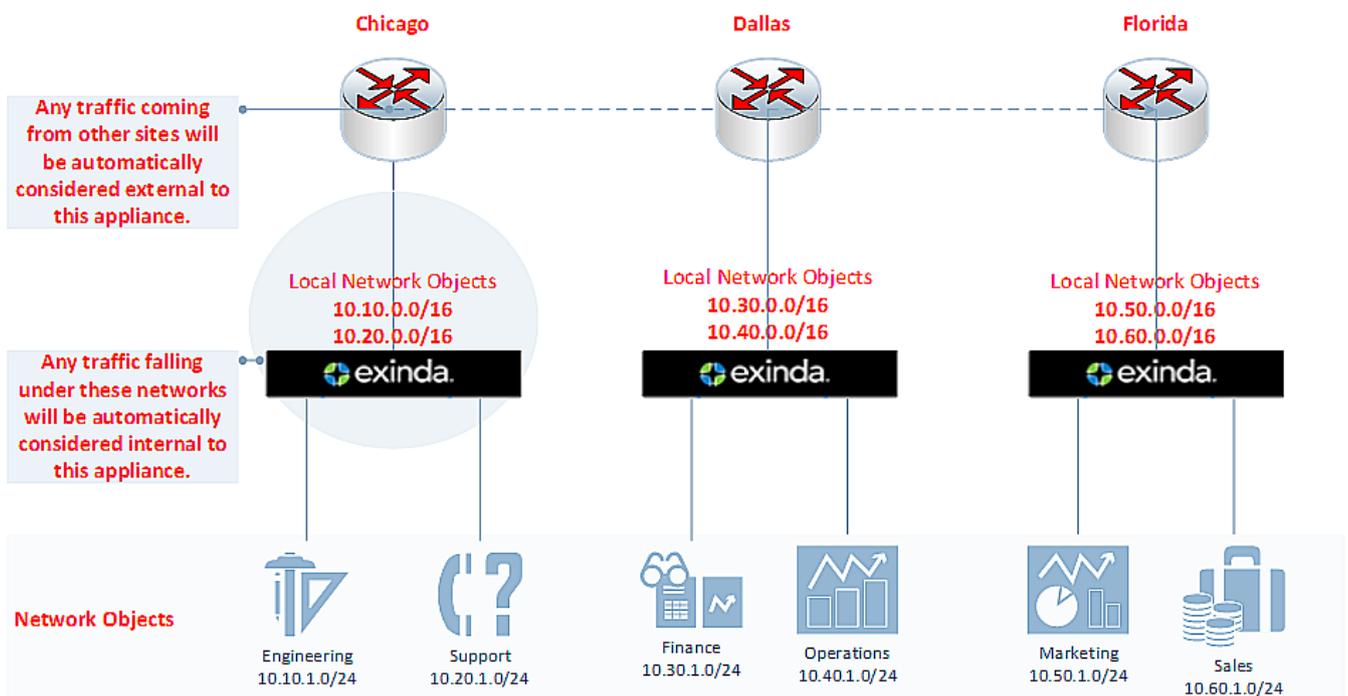
For reporting purposes, hosts and users are defined as internal or external by comparing the IP address with the network objects.

## EXAMPLE

Considering that there are three sites, Chicago, Boston, and Dallas where each site has an appliance, and each appliance has a local network object. Do the following:

- » Create network objects to represent the Chicago site, the Boston site, and the Dallas site.
- » On the Chicago appliance, when compared to the appliance's local network object, the Chicago network object gets set as internal, and the Boston and Dallas network objects are set as external.
- » On the Boston appliance, the Boston network object gets set as internal, and the Chicago and Dallas network objects are set as external.

In the EMC configuration, local network objects are appliance specific, so appliances cannot share these objects.



Screenshot 54: Diagram depicting separated local network objects for each appliance.

To configure the local network object:

1. Go to **Configured Appliances > Local Network Objects**.
2. Click **Local to Appliance (Host ID)** of the appliance to configure the local network object.

## Local Network Objects by Subnet

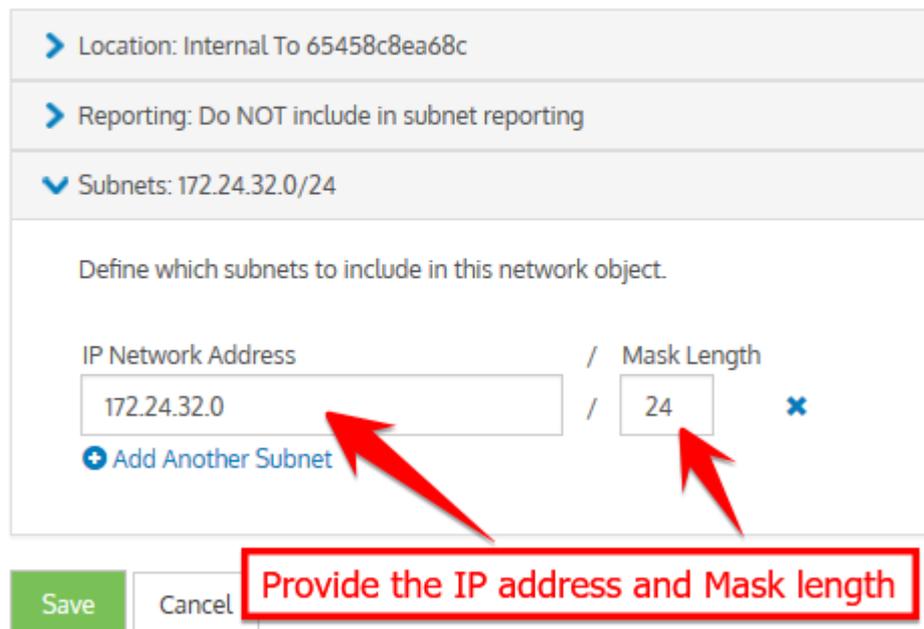
Local network objects represent the subnet local to each Exinda appliance in the system to-local traffic. The local network objects cannot be deleted.



3. Key-in the **IP Network Address** and **Mask Length**.

## Local Network Objects by Subnet

Local network objects represent the subnet local to each Exinda appliance in the system to-local traffic. The local network objects cannot be deleted.

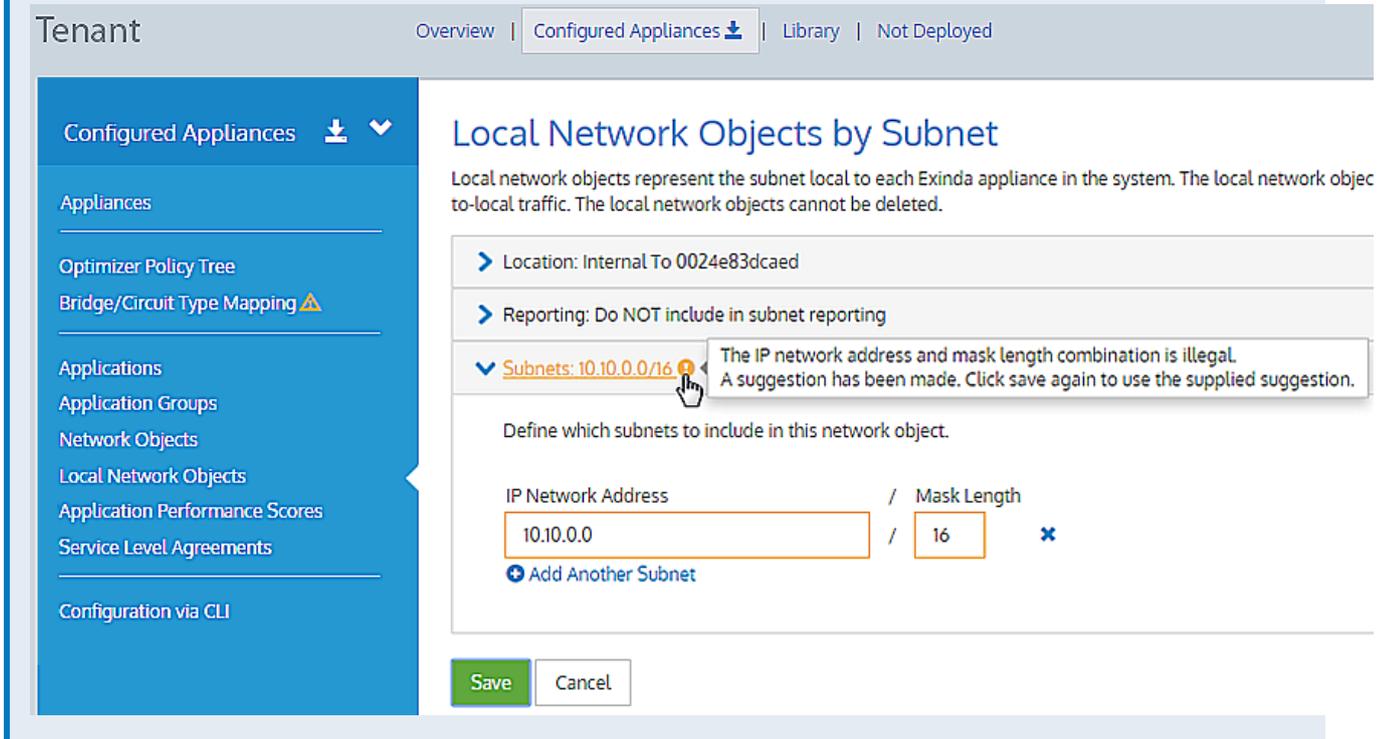


Screenshot 55: Defining local network object by subnet.

4. Click **Save**.

## NOTE

If the IP address and mask length you define are illegal, the EMC returns a warning and provides a suggestion for correcting the issue.



The screenshot shows the 'Tenant' management interface. The breadcrumb trail is 'Overview | Configured Appliances | Library | Not Deployed'. The left sidebar contains a 'Configured Appliances' menu with sub-items: Appliances, Optimizer Policy Tree, Bridge/Circuit Type Mapping, Applications, Application Groups, Network Objects, Local Network Objects, Application Performance Scores, Service Level Agreements, and Configuration via CLI. The main content area is titled 'Local Network Objects by Subnet'. Below the title is a description: 'Local network objects represent the subnet local to each Exinda appliance in the system. The local network object to-local traffic. The local network objects cannot be deleted.' The configuration area includes: 'Location: Internal To 0024e83dcaed', 'Reporting: Do NOT include in subnet reporting', and a 'Subnets: 10.10.0.0/16' entry with a warning icon. A tooltip points to the warning icon with the text: 'The IP network address and mask length combination is illegal. A suggestion has been made. Click save again to use the supplied suggestion.' Below this is a form to 'Define which subnets to include in this network object.' with fields for 'IP Network Address' (10.10.0.0) and 'Mask Length' (16), and an 'Add Another Subnet' button. At the bottom are 'Save' and 'Cancel' buttons.

Screenshot 56: Defining local network object by subnet.

## Importing network objects

If you have existing appliances that are already configured with network objects, in the Exinda Management Center you can import these network objects into the library (one at a time) from the Not Deployed list. When importing network objects, the importer indicates if the network object already exists in the library (or is included in another network object), or if it conflicts with another network object in the library. By importing your network objects, you can more quickly start building a library of objects to use. The system lets you know if the imported network object is already in the library or if the imported network object definition is a subset of a network object in the library. This allows you to quickly ensure consistency across multiple appliances.

To import network objects:

1. From the **Not Deployed** appliance list in your tenant, select the appliance from which you want to import the network objects.
2. Click **Import Configuration**.

# Import Configuration

Network objects from the configuration file can be imported into the library. The system will warn when a network object is not available for import or already has the same definition.

The system will warn when a network object is not available for import or already has the same definition.

## Step 1 - Network Objects

| <input type="checkbox"/>   | Name              | IP Network Address                     |
|----------------------------|-------------------|----------------------------------------|
| <input type="checkbox"/>   | Exinda Appliances | 1.2.3.5/32<br>1.2.3.6/32<br>1.2.3.4/32 |
| <input type="checkbox"/> ✓ | Toronto           | 10.21.0.0/24                           |
| <input type="checkbox"/>   | Data Center       | 10.10.10.0/24                          |

Screenshot 57: Importing configuration

3. Click **Import Network Objects**. A list of network objects, with their IP addresses, appears. If a network object has the same name as a network object in the library, the following indicate the status:

- ✓ – if the IP addresses are the same, the name has a green check mark before it.
- ⊆ – if all of the object IP addresses are contained within the one in the library, the name has a green subset symbol before it.
- ⚠ – if the imported network object has an IP address that is not in the network object with the same name in the library, the name has a warning icon before it. In this case, you have to resolve this conflict manually by modifying the network object in the Library.

4. Select the network objects that you would like to import.

5. Click **Add Selected Network Objects to the Library**. The network objects are imported into the library.

### NOTE

While importing network objects, neither the location (internal or external) of the network object, nor the reporting flag are imported.

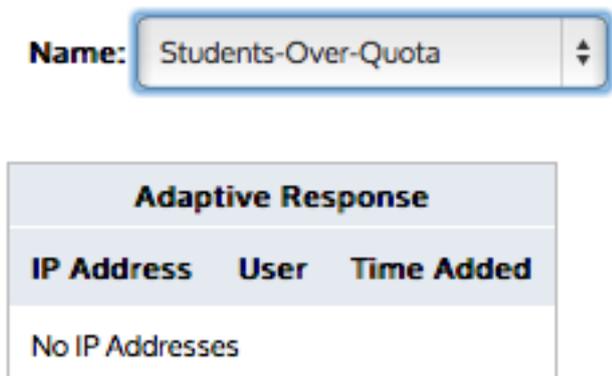
## 0.10.2 Working with dynamically created network objects

Dynamic network objects are network objects that are automatically updated and maintained by the Exinda Appliance. They can be used anywhere static network objects are used, however, they cannot be manually modified.

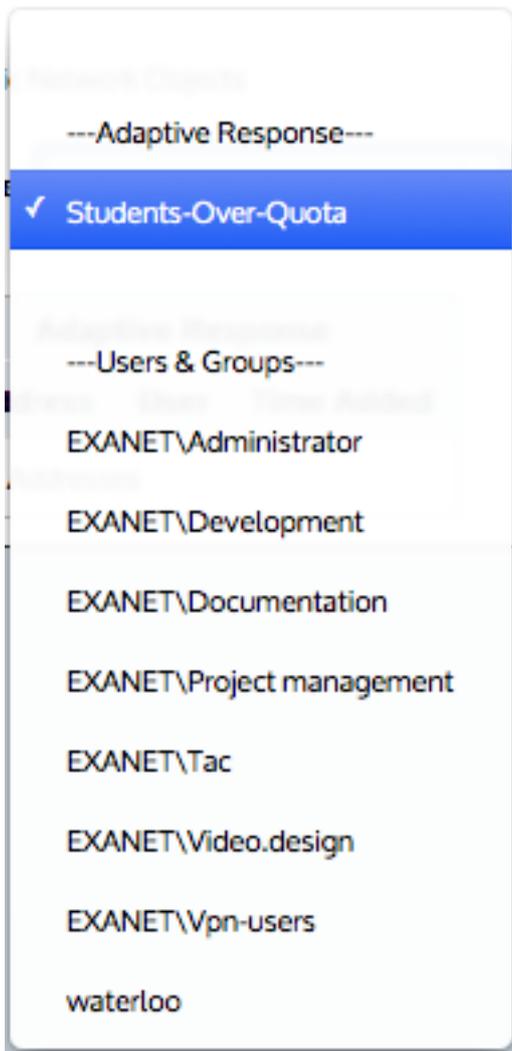
You can view the contents of a dynamic network object by selecting it from the drop-down at the top of the page. It displays the IP addresses, usernames (if applicable) and the date/time the specific address was dynamically added.

There are two types of dynamic network objects.

- » Adaptive Response Dynamic Network Objects - When adaptive response rules are created, a corresponding dynamic network object is automatically created. This dynamic network object is populated by the hosts that have exceeded their adaptive response quota. For more information, refer to [Configuring adaptive response limits](#) (page 217).
- » Active Directory Dynamic Network Objects - When Active Directory users or groups are defined, a corresponding dynamic network object is automatically created. This dynamic network object is populated by the hosts that make up that particular Active Directory user or group. For further information, see the [Integrate with Active Directory](#) and [Users and Groups](#).



Screenshot 58: Select the desired dynamic network object to show the IP addresses/User



Screenshot 59: Select a configured adaptive response network object or a user network object or user group network object

To view it, go to **Configuration > Objects > Users & Groups**.

### 0.10.3 Working with users and groups as objects

Users and groups objects are used to define pre-populated users and groups such that they can be used for monitoring and optimization.

There are two ways the Exinda Appliance can learn about user and group information:

1. **Active Directory**: The Exinda Appliance can receive user and group information using the Exinda Active Directory Service, installed on Active Directory Servers.
2. **Static Users and Groups**: Static users and group information can be only entered using the CLI "networkuser" command.

Once the appliance has learned about users and groups, you can use the users and groups pages to define which users and groups to expose as **Dynamic Network Objects**, for use in monitoring and optimization.

- » To define users as Dynamic Network Objects, see [Create Network User Objects](#).
- » To define groups as Dynamic Network Objects, see [Create Network Group Objects](#).

## Defining network user objects

The Network Users page displays a pre-populated list of users (and their associated IP addresses) from either the Exinda AD Connector, or from static users entered using the CLI. Select which individual users you want to define as dynamic network objects.

Once a user is defined as a dynamic network object, the object can be used in the Optimizer policies.

| <input type="checkbox"/> | User (Domain)           | IP         | Network Object |
|--------------------------|-------------------------|------------|----------------|
| <input type="checkbox"/> | Dev_user_1 (HEADOFFICE) | 172.1.1.6  | ✘              |
| <input type="checkbox"/> | Dev_user_2 (BRANCH1)    | 172.1.1.19 | ✘              |
| <input type="checkbox"/> | Dev_user_3 (BRANCH2)    | 172.1.1.13 | ✘              |
| <input type="checkbox"/> | Dev_user_4 (BRANCH2)    | 172.1.1.14 | ✘              |
| <input type="checkbox"/> | Dev_user_5 (BRANCH2)    | 172.1.1.15 | ✘              |
| <input type="checkbox"/> | Dev_user_6 (BRANCH1)    | 172.1.1.16 | ✘              |
| <input type="checkbox"/> | Qa_user_7 (BRANCH1)     | 172.1.1.9  | ✘              |
| <input type="checkbox"/> | Qa_user_8 (BRANCH1)     | 172.1.1.10 | ✘              |
| <input type="checkbox"/> | Qa_user_9 (BRANCH1)     | 172.1.1.11 | ✘              |

Screenshot 60: A list of network users displayed on the Network Users page.

## Defining and removing users as dynamic network objects

Use the following instructions to define users as dynamic network objects and to stop identifying them as necessary. The instructions focus on dealing with one user at a time, but you can define or remove many users by selecting multiple checkboxes.

1. Go to **Configuration > Objects > Users & Groups > Network Users**.
2. Find the user in the listing.

### TIP

If you have many users, use the links at the top of the page to help find the user.

3. Select the checkbox for the user.
4. At the bottom of the page, click **Add Network Object**. The Network Status icon for the user changes to , indicating it is a network object.

### To stop identifying (remove) a user as a dynamic network object

1. Go to **Configuration > Objects > Users & Groups > Network Users**.
2. Find the user in the listing.

### TIP

If you have many users, use the links at the top of the page to help find the user.

3. Select the checkbox for the user.
4. At the bottom of the page, click **Remove Network Object**. The Network Status icon for the user changes to , indicating it is no longer a network object.

## Configuring network user group objects

The Network Groups page displays a pre-populated list of groups from either the Exinda AD Connector, or from static groups entered using the CLI. This page allows you to select which groups you want to define as dynamic network objects. Once a group is defined as a dynamic network object, it can be used in the Optimizer policies.

### Related topic

[Defining and removing user groups as dynamic network objects](#)

[Defining and removing user groups as dynamic network objects](#)

### To define a group as a dynamic network object

Use the following instructions to define a user group object.

1. Go to **Configuration > Objects > Users & Groups > Network Groups**
2. Find the group in the list, and click **Edit**.
3. To map all users within the selected network group to the network object, select **Map to Network Object**.
4. Select **Ignore Domain** to exclude the domain prefix.
5. Click **Apply**. The Network Status icon for the group changes to , which indicates it is now a network object. If the dynamic network object is created from multiple groups, the groups are combined into a single entry and each domain is identified after the group name.

### To remove a group as a dynamic network object

1. Go to **Configuration > Objects > Users & Groups > Network Groups**
2. Locate the group in the list, and click **Edit**.
3. Clear the **Map to Network Object** checkbox.
4. Click **Apply**. The Network Status icon for the user group changes to , which indicates it is no longer a network object.

#### NOTE

If the dynamic network object was created from multiple groups, each group is again listed individually in the list.

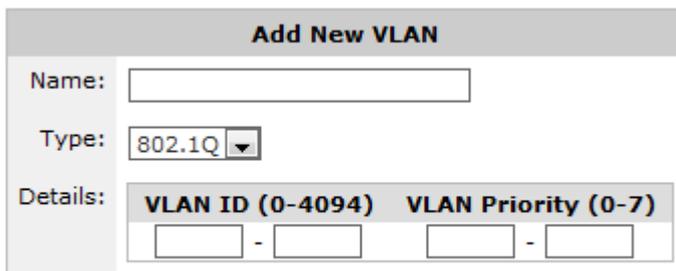
## 0.10.4 Configuring VLAN objects

Virtual LAN (VLAN) objects are used to logically separate hosts (or groups of hosts) on a functional basis rather than on a physical basis. Once VLAN Objects are defined, they can be used in Optimizer policies to filter traffic.

### Configuring VLAN objects in the Exinda Web UI

By default, the Exinda Appliance has a single VLAN defined called "ALL", which matches all traffic (regardless if that traffic is part of a VLAN or not). Additional VLAN Objects can easily be added.

All the defined VLAN objects are shown in the table. Each VLAN object can be edited or deleted by clicking the appropriate button in the table. The **ALL** VLAN object is protected and cannot be edited or deleted.



**Add New VLAN**

Screenshot 61: Adding a new VLAN.

To add a new VLAN object:

1. Go to **Configuration > Objects > VLANs**.
2. Enter a meaningful name for the VLAN object.
3. Specify the type of VLAN to define. Currently only 802.1Q VLANs are available.
4. Specify the range of VLAN IDs to define. To define all VLAN IDs, leave this field blank or type 0 - 4094. A single VLAN ID can be defined by entering the same value in both fields.
5. Specify the VLAN Priority range to define. To define all VLAN Priorities, leave this field blank or type 0 - 7. A single VLAN Priority can be defined by entering the same value in both fields.
6. Click the **Add New VLAN** button. The VLAN will be added to the list of VLANs in the table.

#### EXAMPLE

Consider VoIP traffic that has a VLAN ID of 10. Create a VLAN object with this ID. This object can then be used to prioritize VoIP traffic using the Optimizer.

Name: VoIP

Type: 802.1Q

VLAN ID: 10 - 10

VLAN Priority: 0 - 7 (or leave this field blank)

The VLAN priority is a field in the 802.1Q header that networking devices use for their own QoS-purpose. In order for the VLAN Object to work properly, the VLAN Priority field indicated in the object must match the Priority set in the 802.1Q header. If the priority set for a specific VLAN is unknown, Exinda recommends to configure the VLAN priority inside the VLAN object as "0 - 7", which covers all possible scenarios. For most networking vendors, if the VLAN priority was not specified, the packets are tagged with Priority 0 by default.

## Configuring VLAN objects in the EMC

Virtual LAN (VLAN) Objects are used to logically separate hosts (or groups of hosts) on a functional basis rather than on a physical basis. Once VLAN Objects are defined, they can be used in Optimizer policies to filter traffic. By default, the Exinda Appliance has a single VLAN defined called **ALL**, which matches all traffic (regardless if that traffic is part of a VLAN or not). The **ALL** VLAN object is protected and cannot be edited or deleted. Additional VLAN Objects can easily be added.

### To add a new VLAN object

1. Go to **Library > VLANs**.

2. Click **Create new VLAN in the library...**
3. In the **Name** field, key-in a meaningful name for the VLAN.

## VLAN

Define 802.1Q VLANs to logically group network nodes. These VLANs can then be used in optimizer policies to filter traffic.

▼ Name: VLAN name

Name

4. Expand the **Definition** section:

- a. In the VLAN ID **Start** and **End** fields, key-in the range of VLAN IDs that must appear in the VLAN.

### NOTE

The absolute range is 0–4094. This would equate to all VLAN IDs being in the VLAN. Leaving both fields blank would give the same result. To define a lesser range, type a range somewhere within the absolute range. To isolate one VLAN, type its ID value in both the **Start** and **End** fields.

- b. In the VLAN Priority **Start** and **End** fields, key-in the range of values for this VLAN.

### NOTE

You can define priorities within a maximum range of 0–7. This would equate to all priorities being assigned to the VLAN. Leaving both fields blank would give the same result. To define a lesser range, type a range somewhere between the maximum range. To define just one priority, type the same value in both the Start and End fields.

# VLAN

Define 802.1Q VLANs to logically group network nodes. These VLANs can then be used in optimizer policies to filter traffic.

> Name: VLAN name

▼ Definition

VLAN ID (0 - 4094)

Start:  End:

VLAN Priority (0 - 7)

Start:  End:

4. Click **Create**.

## Deleting a VLAN

### NOTE

You can only delete those VLAN objects that are not currently in use. If a  icon appears in the row of a VLAN object that you wish to delete, it means that the object cannot be deleted because it is in use. If you still want to delete it, you first need to unlink it from the appliances.

Also the ALL VLAN object is protected and cannot be deleted.

**The "All" VLAN is available by default and cannot be deleted**

1. Go to **Library > VLANs**.
2. In the list of VLAN library objects, find the VLAN you need to delete.

## VLAN Library

Define 802.1Q VLANs to logically group network nodes. These VLANs can then be used in optimizer policies to filter traffic.

+ Create new VLAN in the library ...

| Name ▲            | VLAN ID   | VLAN Priority |        |
|-------------------|-----------|---------------|--------|
| All               | 0 - 4094  | 0 - 7         |        |
| Database          | 0 - 25    | 2 - 4         |        |
| Social Networking | 100 - 500 | 0             | Delete |

3. In the VLAN object row, click

Are you sure you want to delete this VLAN from the library?

Note: This operation cannot be undone.

Don't ask when deleting VLANs again.

4. Confirm and click **Delete**.

### 0.10.5 Adding protocol objects

Protocol objects are used to define IPv4 protocol numbers that can then be used to define application objects. By default, the appliance factory setting includes all major Internet Protocol (IPv4) related protocols, including ICMP (Internet Control Message Protocol), TCP (Transmission Control Protocol) and UDP (User Datagram Protocol). Additional IPv4 protocols can easily be added by simply specifying IPv4 protocol number.

#### NOTE

Protocol numbers are unique and can only be defined once.

All the defined protocol objects are shown in the table. Each protocol object can be edited or deleted by clicking the appropriate button in the table. Some protocols are protected and cannot be edited or deleted.

Add New Protocol

Name:

Number:

Screenshot 62: Adding a new protocol.

To configure protocol objects:

1. Go to **Configuration > Objects > Protocols**.
2. Enter a meaningful name for the protocol.
3. In the **Number** field, specify the IPv4 protocol number.
4. Click the **Add New Protocol** button. The protocol will be added to the list of protocols in the table.

#### EXAMPLE

Consider where SCTP (Stream Control Transport Protocol) is undefined by default and need to be defined.

Name: SCTP

Number: 132

### 0.10.6 Adding application objects

Application objects are used to classify traffic on the network and are made up of layer 7 signatures or TCP/UDP port numbers and port ranges. Application classification can be used to monitor traffic or to create application-specific policy. There are many predefined applications on the appliance. You can add any applications that are not already in the list.

Applications can be created from various combinations of L7 signatures, TCP/UDP port numbers or ranges, and network object. The following are valid combinations.

- » Applications based on L7 signatures. For example, you can create an application for a particular website by selecting http, host, and entering the domain of the website.
- » Applications based on L7 signature and TCP/UDP port numbers or ranges, which are OR'd together. For example, you could define HTTP based on TCP port 80 OR 'http' L7 signature.
- » Applications based on network object and TCP/UDP port numbers or ranges, which are AND'd together. For example, you could define an application based on a particular port number on a particular server (specified by network object).
- » Applications based on only network object. For example, you could define an application based on a particular application server (specified by network object).
- » Applications based on only TCP/UDP port number or ranges. For example, you could create an application based on a particular port.

Network objects cannot be used in conjunction with a layer 7 signature.

**Add New Application**

|                         |                                                                                                                                                                                                                                                                                                                       |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name:</b>            | <input style="width: 80%;" type="text" value="Exinda Website"/>                                                                                                                                                                                                                                                       |
| <b>Network Object:</b>  | <input style="width: 80%;" type="text" value=""/>                                                                                                                                                                                                                                                                     |
| <b>L7 Signature:</b>    | <input style="width: 25%;" type="text" value="http ---&gt;"/> <input style="width: 15%; border: none;" type="button" value="▼"/> <input style="width: 25%;" type="text" value="host"/> <input style="width: 15%; border: none;" type="button" value="▼"/> <input style="width: 20%;" type="text" value="exinda.com"/> |
| <b>Ports/Protocols:</b> | <input style="width: 25%; border: none;" type="button" value="TCP Port/Range"/> <input style="width: 50%; border: none;" type="button" value="▼"/> <input style="width: 25%;" type="text" value="eg. 80,8080,3127-3128"/>                                                                                             |

[Show a List of Common Port Numbers](#)

**Add New Application**

Screenshot 63: Adding a new application object.

**NOTE**

When creating applications based on ports, any given port number can only be defined once for TCP and once for UDP. The same port number can be defined for TCP and UDP. For example, if you define an application object with a port range TCP 500-510, you cannot then define another application object on TCP port 505. However, you can define another application object with UDP port 505.

You can define duplicate ports/port ranges if a network object is also specified.

Many of the L7 signatures have sub-type classifications, which makes layer 7 visibility much more granular. For instance, for reporting on specific web applications, most vendors can only report on port 80 traffic. Exinda allows a deeper look into Layer 7 applications. For example, by comparison:

- » Layer 4 reporting tools report on web applications as: port 80 or HTTP
- » Layer 7 reporting tools report on web applications as: Yahoo or Skype
- » Layer 7 with sub-type classification report on web applications as: Yahoo video, Yahoo voice, or Yahoo webchat.

This allows you to monitor at a much more granular level.

## Adding application objects in the Exinda Web UI

**Add New Application**

Name:

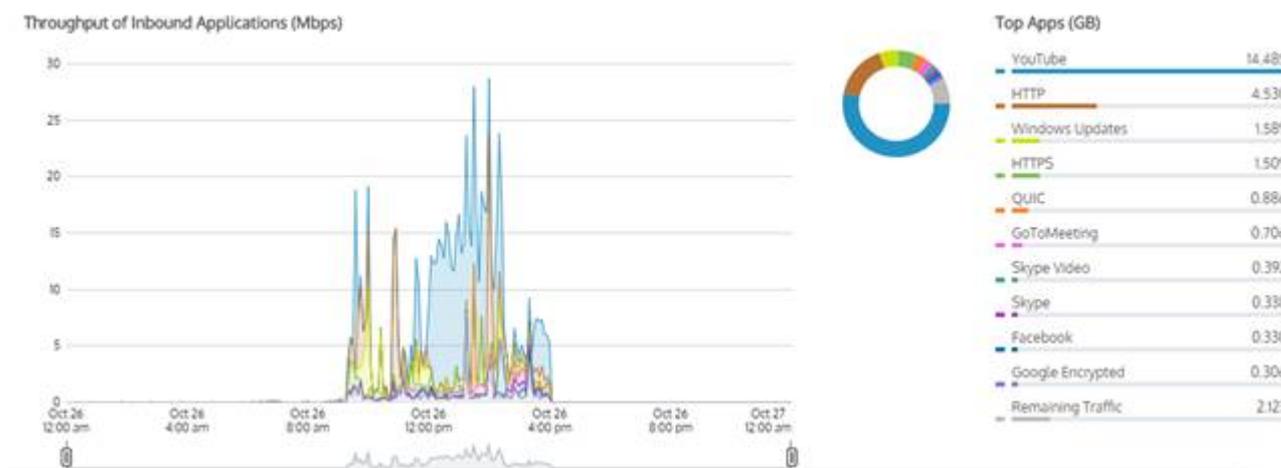
Network Object:

L7 Signatures:

Ports/Protocols:

[Show a List of Common](#)

- file-transfer
- unknown
- video
- voice
- webchat



Screenshot 64: Chart displaying bandwidth throughput for applications.

### Related Topics

- » [Predefined Applications and Supported L7 Signatures](#)
- » [What L7 signature options are there?](#)
- » [Adding a new application](#)
- » [Example: How to create a custom application based on the HTTPS protocol](#)

## Adding application objects in the EMC

The Exinda Management Center provides a comprehensive set of built-in Applications for you to use. These cannot be edited, but you can define Custom Applications. You can add the custom application as part of an ignore policy at the bottom of the optimizer tree that is sent to the appliances.

### NOTE

if you use a custom application in the definition of a virtual circuit or policy for a given appliance group, then the custom application is automatically added to the appliance-group configuration.

Applications can be defined in the Library directly, and then later assigned to an Optimizer Policy Tree. Applications can also be created in an Optimizer Policy Tree, in which case they are also saved to the Configuration Library for use elsewhere. All applications within a tenant must have unique names. When an application is modified, all uses of it are also modified.

## Where do I find Applications?

Application library items can be found in **Library > Applications**. You can define custom applications for each appliance group. Go to the desired appliance group in the **Optimizer Policy Tree**.

## How do I view built-in Applications?

Built-in application library items can be found in **Library > Applications > Built-in**. You can view built-in applications, but not edit.

## To create a Custom Application in the Library

1. Go to **Library > Applications**.
2. Click **Create new application...**
3. In the **Name** section, key-in a name for the application. The name must be unique in the tenant.
4. In the **Definition** section, choose either the Layer 7 Signature, or select a combination of Network Object, Ports, DSCP and Protocols.
5. Click the **Create** button. The custom application will be added to the library list.

## To create a new Application directly in the Optimizer Policy Tree

1. Go to **Configured Appliances > (desired appliance group) > Applications**.
2. Click **Create new application...**
3. In the **Name** section, key-in a name for the application. The name must be unique in the tenant.
4. In the **Definition** section, choose either the Layer 7 Signature, or select a combination of Network Object, Ports, DSCP and Protocols.
5. Click the **Create** button. The custom Application is added to the Optimizer Policy tree for this current Appliance Group and is also saved to the Library.

## To add an Application from the library to an appliance group

1. Go to **Configured Appliances > (desired appliance group) > Applications**.
2. Click **Add Application from Library**.
3. Select one or more applications to apply to the current appliance group, and then click **application to <group\_name>**.

## To use a Custom Application in a policy definition

When creating a policy, you can use a custom Application to filter traffic to or from this application. The custom Applications appear in the Application list within the **Filter** section. For more information, refer to [Policies](#) (page 430).

## 0.10.7 Adding and updating application group objects

To properly classify applications on your network it is important to understanding what is happening and for controlling or protecting a particular type of traffic.

The Exinda Appliance comes with a long list of predefined applications used to classify your network traffic. If, however, you want to create your own application, you can create new applications based on L7 signatures, TCP/UDP port numbers and port ranges, or network objects.

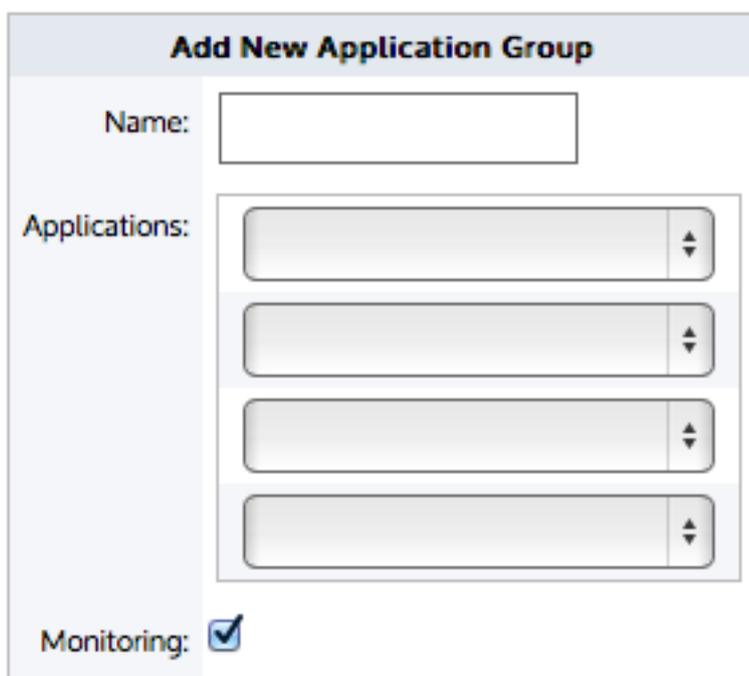
You may also want to monitor, control, or protect your traffic by grouping a set of applications. For instance, controlling social networking applications as a group in most cases provides adequate granularity. The Exinda Appliance comes with a default set of application groups. You can add new applications to these groups, or create new groups, or delete existing groups.

There are several predefined application groups, such as Mail, P2P, Voice, etc. You can edit existing application groups or create new ones.

#### NOTE

A given application can exist in multiple application groups. However, monitored groups must not contain applications which are already a member of another group being monitored. Any given application can only be monitored within a single application group.

### Adding application group objects in the Exinda Web UI



The screenshot shows a web form titled "Add New Application Group". It includes a "Name:" field, four "Applications:" dropdown menus, and a "Monitoring:" checkbox which is checked.

Add New Application Group

Screenshot 65: Adding a new application group.

#### To add a new application group

1. Go to **Configuration > Objects > Applications > Application Groups**.
2. In the **Add New Application Group** area type a name for the new group.
3. Select the applications that belong in the new group. By default, there are four drop-downs available to add Application Objects. If you need to add more, save the application group object, then select the **Edit** button next to the newly created application group. You will be presented with four additional drop-downs to add more applications.
4. If you want this application group to be monitored in the Application Group report, select the **Monitoring** checkbox.
5. Click **Add New Application Group**.

## To update an application group

1. Go to **Configuration > Objects > Applications > Application Groups**.
2. Locate the group from which to add or delete applications and click **Edit**.
3. Select a new application from a blank drop-down list. Or to remove an application, open the drop-down list with the application to remove and select the blank row at the top.
4. Click **Apply Changes**.

## What application groups are predefined?

For more information, refer to [Predefined Application Groups](#) (page 795).

## Adding application groups in the EMC

The Exinda Management Center provides a comprehensive set of built-in Application Groups for you to use, but you can also define Custom Application Groups.

Although an application can be a member of multiple application groups, to prevent conflict it can only be a member of the application group that is currently monitoring traffic. For example, Skype cannot be added to both the Voice group and the Messaging group because EMC can gather data from only one monitored application group for reporting.

When an Application Group is created in the library, it is applied tenant-wide and is therefore available in every appliance group within an Appliance Group section. If an application group is set for monitoring and/or being used in the optimizer tree, then this application group is pushed to the respective appliances within the push configuration.

If a custom application is added to an application group, where the group does not exist in the configuration of the appliances, the application is first added to the appliances and then the group is imported.

In addition, the following limitations should be noted:

- » If one application within a group is not supported by an appliance, then that application definition will not be sent to that specific appliance.
- » If you try to add an application to an appliance with a firmware version does not support the application, the EMC displays an error for the appliance and the application is not imported. However, the appliance does import the application group along with other settings.

### Where do I find Application Groups?

Application Group library items can be found in **Library > Application Group**. A lock icon in the Monitoring column indicates that an Application Group is in use and cannot, therefore, be deleted.

### To use an Application Group in a policy rule definition

When creating a policy rule, you can use an Application Group to filter traffic to or from the applications within the group. The Application Groups appear in the Application list within the **Filter** section. Refer to [For more information, refer to Policies](#) (page 430). for details.

### How do I configure built-in Application Groups?

You can modify application groups either within the configured appliances or from the library.

1. Click the desired application group name to edit.

# Application Groups

Application groups can be used when defining a policy and to monitor groups of applications. An application can be a member of multiple application groups, but can only be a member of one application group that is monitored.

| Name ▲            | Applications |
|-------------------|--------------|
| Database Services | MS-SQL       |
|                   | MySQL        |
|                   | Oracle       |
|                   | PostgreSQL   |
|                   | TDS          |

2. You can configure monitoring status and add or remove applications within this group.

## Application Group

Application groups can be used when defining a policy and to monitor groups of applications. An application can be a member of multiple application groups, but can only be a member of one application group that is monitored.

> Name: Database Services

> Reporting: Include in application group reporting

> Applications:

> In Use: 6

[Update in Library](#) [Cancel](#)

3. You could also view which policies are currently using this application group under **In Use**.

# Application Group

Application groups can be used when defining a policy and to monitor groups of applications. An application can be a member of multiple application groups, but can only be a member of one application group.

> Name: Database Services

> Reporting: Include in application group reporting

> Applications:

▼ In Use: 6

| Type of use | Name                                            |
|-------------|-------------------------------------------------|
| Policy      | Database - Guarantee High 10%-100% - Accelerate |
| Policy      | Database - Guarantee Low 5%-100%                |
| Policy      | Database - Guarantee Med                        |
| Policy      | Database - Guarantee Med 8%-100%                |
| Policy      | Database - Guarantee Med 8%-100% - Accelerate   |
| Policy      | Database - Limit High 4%-70%                    |

[Update in Library](#) [Cancel](#)

How do I create a custom application group?

1. Go to **Library > Application Groups** and click **Create new application group in the library...**

## Application Group Library

Application groups can be used when defining a policy :  
An application can be a member of multiple application

**+ Create new application group in the library ...**

- Name ▲

- Database Services

2. Provide a name, configure monitoring, if you want this group to be monitored, and add the applications to be part of this group.

# Application Group

Application groups can be used when defining a policy and to monitor  
An application can be a member of multiple application groups, but ca

|                 |
|-----------------|
| > Name:         |
| > Reporting:    |
| > Applications: |

Create Cancel

## How do I know which application groups are enabled for monitoring?

By default, all the built-in application groups are enabled for monitoring. On the main **Application Groups** page, you can view the specific groups that are set for monitoring.

# Application Groups

Application groups can be used when defining a policy and to monitor groups of applications.  
An application can be a member of multiple application groups, but can only be a member of one application group that is monitored.

| Name ▲            | Applications | Monitoring                          |
|-------------------|--------------|-------------------------------------|
| Database Services |              | <input checked="" type="checkbox"/> |
|                   | MS-SQL       |                                     |
|                   | MySQL        |                                     |
|                   | Oracle       |                                     |
|                   | PostgreSQL   |                                     |
|                   | TDS          |                                     |

| Applications | Monitoring                          |
|--------------|-------------------------------------|
|              | <input checked="" type="checkbox"/> |
| MS-SQL       |                                     |
| MySQL        |                                     |
| Oracle       |                                     |
| PostgreSQL   |                                     |
| TDS          |                                     |

You can always change the monitoring configuration by clicking on the application group name and changing it.

## Predefined Applications and Application Groups

This section provides information about the predefined applications and application groups that your Exinda Appliance (s) is pre-configured to manage.

To see a list of predefined applications and application groups navigate to [http://go.gfi.com/?pageid=exos\\_help#cshid=pre\\_apps](http://go.gfi.com/?pageid=exos_help#cshid=pre_apps)

### 0.10.8 Configuring anonymous proxy detection and monitoring

Anonymous proxies are typically used to circumvent security policies, allowing users to access prohibited recreational, adult or other non-business sites by tunneling this traffic over a regular or encrypted HTTP session. Anonymous Proxies also provide anonymity; users accessing websites through an Anonymous Proxy cannot easily be traced back to their original IP.

Exinda Appliances have built-in support for anonymous proxy detection. The Exinda Appliance receives daily updates from [www.exinda.com](http://www.exinda.com) containing updated anonymous proxy definitions, much like anti-virus applications receive daily threat updates.

The anonymous proxy application is a special application object that is used to detect anonymous proxy websites and services. However, the anonymous proxy service is disabled by default.

If the anonymous proxy service is enabled, the Exinda appliance fetches a list of anonymous proxy definitions from the Exinda web servers on a daily basis.

An application object called 'Anonymous Proxy' is automatically created. The Anonymous Proxy application tracks all traffic sent through one of the anonymous proxies in the list. This application object is displayed in the monitoring reports like any other application object and can also be used in the Optimizer policies.

#### NOTE

- » Anonymous Proxy classification only occurs if the Anonymous Proxy ASAM module is enabled on the **Configuration > System > Setup > Monitoring** page.
- » In order to receive daily Anonymous Proxy definition updates, the Exinda appliance must be able to contact the Exinda web servers and the appliance must also have a valid software subscription.

| Anonymous Proxy Options |                                            |
|-------------------------|--------------------------------------------|
| Auto Update Service     | <input checked="" type="checkbox"/> Enable |

Apply changes

| Settings    |                                           |
|-------------|-------------------------------------------|
| URL         | http://updates.exinda.com/aplist/alist.gz |
| Last Check  | 2014/07/18 20:19:02                       |
| Last Update | 2014/07/18 20:19:02                       |
| Status      | Ok                                        |

The **renumerate** button refreshes the Anonymous Proxy list immediately

Renumerate

Screenshot 66: The form to enable the Anonymous Proxy service to keep of list of anonymous proxy sites.

| ASAM            |                                            |
|-----------------|--------------------------------------------|
| Anonymous Proxy | <input checked="" type="checkbox"/> Enable |

Screenshot 67: The form to enable/disable the Anonymous Proxy ASAM required for classification.

Where to configure it

- » To enable the anonymous proxy service, go to **Configuration > Objects > Applications > Anonymous Proxy**.
- » To enable the anonymous proxy traffic classification, go to **Configuration > System > Setup > Monitoring**.

To enable the anonymous proxy traffic classification

1. Check the Auto Update Service **Enable** checkbox. The appliance will communicate with the Exinda web servers daily and fetch any new anonymous proxy definitions.
2. Ensure that the Anonymous Proxy ASAM module is enabled by going to the **Configuration > System > Setup > Monitoring** page and ensuring the **Anonymous Proxy** checkbox is checked in the **ASAM** section. The Anonymous Proxy ASAM is **on** by default. The appliance will classify traffic by matching the traffic against the anonymous proxy list.

To see when the appliance last updated the anonymous proxy definitions

1. Look at the **Settings** section.

2. The **Last Check** field indicates the last time that the appliance checked the Exinda service for new anonymous proxy definitions.
3. The **Last Update** field indicates the last time new anonymous proxy definitions were found and updated.

To force a check of the anonymous proxy definitions

Click the **Renumerate** button. The appliance will check the Exinda web servers immediately to check for new anonymous proxy information.

To disable the anonymous proxy traffic classification

1. Uncheck the Auto Update Service **Disable** checkbox.
2. Disable the Anonymous Proxy ASAM by going to the **Configuration > System > Setup > Monitoring** page, unchecking the **Anonymous Proxy** checkbox in the **ASAM** section, and clicking the **Apply Changes** button. Disabling the ASAM will clear the existing anonymous proxy definitions.

### 0.10.9 Configuring service level agreement objects

The Service Level Agreement (SLA) objects are used to monitor the availability of a particular IP site. By creating a SLA object, you indicate which IP site to monitor. The Exinda appliance will send one ICMP ping every 10 seconds to the IP address. You can specify the ping packet size to use. You can also specify when an alert will be triggered by specifying the ping latency threshold and the duration that the ping latency threshold was exceeded. An alert is triggered when the latency of the SLA site exceeds the latency threshold for longer than the specified duration.

Configuring service level agreement objects in the Exinda Web UI

| Add New SLA Site                                            |                                     |
|-------------------------------------------------------------|-------------------------------------|
| Name:                                                       | <input type="text"/>                |
| Destination IP:                                             | <input type="text"/>                |
| Latency Threshold (ms):                                     | <input type="text" value="500"/>    |
| Ping Size:                                                  | <input type="text" value="64"/>     |
| Duration:<br>(Duration for which the threshold is exceeded) | <input type="text" value="1 hour"/> |
| Enable:                                                     | <input type="checkbox"/>            |

**Add New SLA Site**

**Cancel**

Screenshot 68: Adding an SLA site.

To access this configuration, go to **Configuration > Objects > Service Levels > Service Level Agreements**.

To create a Service Level Agreement (SLA) object:

1. Click the **Add New SLA Object** button.
2. Type a name for the SLA object.
3. Type a IP address in the **Destination IP** field that will be pinged.
4. Type the **Latency Threshold**(in ms), such that you want to be notified if this threshold is consistently exceeded. The default is 500 milliseconds.

5. Enter the ping packet size (in bytes) to use in the **Ping Size** field. The default is 64 bytes.
6. Select the duration, that is the amount of time that the latency threshold needs to be exceeded before the alert is sent. The options are:
  - 30 seconds
  - 60 seconds
  - 5 minutes
  - 30 minutes
  - 1 hour (Default)
  - 0 - (Disable), which disables the alert.
7. Select the **Enable** checkbox to enable the SLA object to starting pinging the IP site.
8. Click the **Apply Changes** button. The object is added to the list of configured SLA objects.

#### NOTE

Ensure that the Send Email alert is enabled for this on the Configuration > System > Setup > Alerts page.

Valid SMTP and email settings are required for email alerts. To configure, see For more information, refer to [SNMP configuration](#) (page 547). and For more information, refer to [Email configuration](#) (page 544)..

## Configuring service level agreement objects in the EMC

The Service Level Agreement (SLA) library objects are used to monitor the availability of particular IP addresses. By creating an SLA object, you identify the IP address to monitor. The Exinda Management Center then sends one ICMP ping every 10 seconds to the IP address. You can specify the ping packet size to use. You can also specify when an alert is triggered by defining the ping latency threshold and the duration by which the ping latency threshold was exceeded. An alert is triggered when the latency of the SLA site exceeds the latency threshold for longer than the specified duration.

### Where do I find Service Level Agreements?

The Service Level Agreements library items can be found in:

**Library > Service Level Agreements.**

Or,

**Configured Appliances > Service Level Agreements.**

### Setting up a Service Level Agreement in the Configuration Library

1. Go to **Library > Service Level Agreements.**
2. Click **Create new service level agreement in library...**
3. In the **Name** field, key-in a meaningful name for the new Service Level Agreement.

▼ Name

Name

4. Click **Service Agreement** and do each of the following:

- In the **Destination** field, key-in the IP address of the server whose availability you need to monitor.
- In the **Latency Threshold (ms)** field, key-in a value for the response time.
- In the **Ping Size (bytes)** field, key-in a packet size, for example, 1024.
- Select the **Enable Ping** checkbox.

▼ Service Agreement: Ping 192.168.0.25 with ping size of 1024 bytes, allowing for 25 ms of delay. Ping is enabled.

Destination IP

Latency Threshold (ms)

Ping Size (bytes)

Enable Ping

4. Click **Notification**, and select a delay period from the drop-down list. The options are:

- 0 – Disabled – this disables the alert
- 30 seconds
- 60 seconds
- 5 minutes
- 30 minutes
- 1 hour – the Default setting

**TIP**

The default delay is 1 hour. If this setting fits your needs, you do not need to change anything here.

▼ Notification: Notify if threshold is met for 1 hour.

An alert will be sent out after the threshold has been surpassed for the specified amount of time.

Notification

1 hour



7. Click **Create**.

#### Setting up a Service Level Agreement directly in the Appliances

1. Go to **(desired appliance group) > Service Level Agreements**.
2. Click **Create new service level agreement...**
3. In the **Name** field, key-in a meaningful name for the new Service Level Agreement.

▼ Name

Name

Service Level Agreement name



4. Click **Service Agreement** and do each of the following:

- In the **Destination** field, key-in the IP address of the server whose availability you need to monitor.
- In the **Latency Threshold (ms)** field, key-in a value for the response time.
- In the **Ping Size (bytes)** field, key-in a packet size, for example, 1024.
- Select the **Enable Ping** checkbox.

▼ Service Agreement: Ping 192.168.0.25 with ping size of 1024 bytes, allowing for 25 ms of delay. Ping is enabled.

Destination IP 192.168.0.25

Latency Threshold (ms) 25

Ping Size (bytes) 1024

Enable Ping

4. Click **Notification**, and select a delay period from the drop-down list. The options are:

- 0 – Disabled – this disables the alert
- 30 seconds
- 60 seconds
- 5 minutes
- 30 minutes
- 1 hour – the Default setting

**TIP**

The default delay is 1 hour. If this setting fits your needs, you do not need to change anything here.

Notification: Notify if threshold is met for 1 hour.

An alert will be sent out after the threshold has been surpassed for the specified amount of time.

Notification

7. Click **Create**. The Service Level Agreement is added to the appliances and is also saved to the configuration library.

**Modifying a Service Level Agreement**

You can modify and update all the properties of the Service Level Agreement except its name.

**Deleting a Service Level Agreement**

**NOTE**

You can only delete those SLA items that are not currently in use. If a  icon appears in the row of a SLA item you wish to delete, it means that the item cannot be deleted because it is in use. If you still want to delete it, you first need to unlink it from the appliances.

1. Go to **Library > Service Level Agreements**.
2. In the list of SLA library objects, locate the SLA item you need to delete.

## Service Level Agreement

Service level agreements (SLA) provide a method for evaluating how well a server is running.

[+ Create new service level agreement in the library ...](#)

| Name ▲                       | Destination IP | Threshold | Ping Size (bytes) | Notification | Enabled |
|------------------------------|----------------|-----------|-------------------|--------------|---------|
| Service Level Agreement name | 10.10.9.55     | 5 ms      | 1024              | 5 minutes    | ✓       |

Delete

3. In the SLA item row, click .

Are you sure you want to delete this service level agreement from the library?

Note: This operation cannot be undone.

Don't ask when deleting service level agreements again.

**Delete**      Don't Delete

4. Confirm and click **Delete**.

### 0.10.10 Creating an HTML response object

In particular circumstances, you may want to present a specific message to your network users through a webpage. Consider informing your network users that they have exceeded their usage quota. An HTML Response object can be created so that the Exinda appliance will host the webpage. This HTML Response object can then be used in an Optimizer policy to present the webpage.

To create an HTML Response object all you need is a name and the html for the webpage.

**Add New HTML Response**

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name:</b>          | Quota Exceeded                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>HTML Response:</b> | <pre>&lt;html&gt; &lt;head&gt; &lt;style&gt; hr {color:sienna;} p { margin-left:20px; font-family:Verdana; } &lt;/style&gt; &lt;/head&gt; &lt;body style="background:#e3e9ef"&gt; &lt;div style="padding:30px;margin-left:auto;margin-right:auto;width:80%;background:white"&gt; &lt;p style="color:#79c157;font-size:20pt;"&gt;OOPS! YOU'VE REACHED &lt;/br&gt; &lt;span style="color:#3da58e;font-size:20pt;"&gt;YOUR DAILY DATA LIMIT&lt;/span&gt;&lt;/p&gt; &lt;hr&gt; &lt;p style="color:#515151;"&gt;Thank you for visiting the Springfield Mall. Please come again.&lt;/p&gt; &lt;/div&gt; &lt;/body&gt; &lt;/html&gt;</pre> |

**Add Response**

Screenshot 69: Creating a HTML Response object

Use the following instructions to create an HTML Response object.

1. Go to **Configuration > Objects > HTML Response**.
2. In the **Name** field, type a suitable name for the HTML Response object.
3. In the **HTML Response** text box, type the HTML code for the response.

**TIP**

Any HTML with embedded CSS or Javascript can be used.

4. Click the **Add Response** button. The object is added to the list of configured HTML response objects. You can also preview the webpage by pressing the Preview button in the list of HTML Response objects.

The object is now ready to be used within a policy.

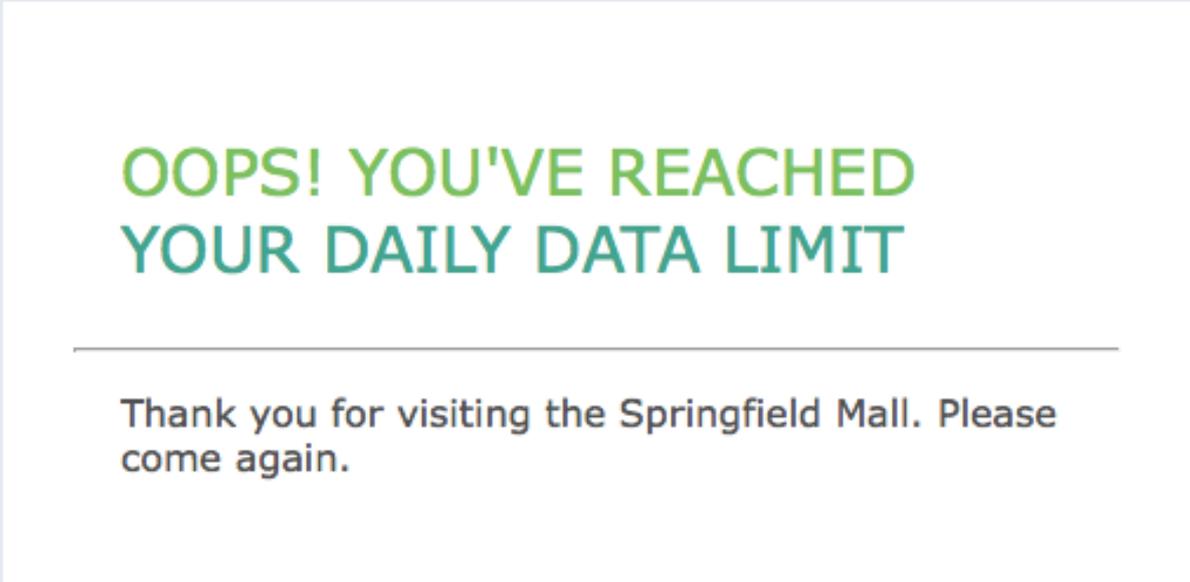
## EXAMPLE

Create a webpage that states that the user has reached their daily data limit.

Enter the following as the HTML Response:

```
<html>
  <head>
    <style>
      hr {color:sienna;}
      p {margin-left:20px;font-family:Verdana;}
    </style>
  </head>
  <body style="background:#e3e9ef">
    <div style="padding:30px;margin-left:auto;margin-right:auto;width:80%;background:white">
      <p style="color:#79c157;font-size:20pt;">OOPS! YOU'VE REACHED </br>
      <span style="color:#3da58e;font-size:20pt;">YOUR DAILY DATA
      LIMIT</span></p>
      <hr>
      <p style="color:#515151;">Thank you for visiting the Springfield Mall.
      Please come again.</p>
    </div>
  </body>
</html>
```

If you click **Preview**, a new browser window opens displaying the message:



OOPS! YOU'VE REACHED  
YOUR DAILY DATA LIMIT

---

Thank you for visiting the Springfield Mall. Please  
come again.

### 0.10.11 Configuring schedule objects

The Exinda Appliance allows you to automate your network optimization policies for different times of the day and different days of the week. For example, you may wish to lock down your network at night to improve security, while still allowing automated backup services and email to function.

## Configuring schedule objects in the Exinda Web UI

By default, there are three schedule objects defined. Each schedule can be edited or deleted by clicking the appropriate button in the table. The **ALWAYS** schedule is protected and cannot be edited or deleted.

| Name               | From Day | To Day   | From Time | To Time | Edit                                | Delete                                |
|--------------------|----------|----------|-----------|---------|-------------------------------------|---------------------------------------|
| <b>After Hours</b> |          |          |           |         | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
|                    | Monday   | Friday   | 0:00      | 08:00   |                                     | <input type="button" value="Delete"/> |
|                    | Monday   | Friday   | 18:00     | 24:00   |                                     | <input type="button" value="Delete"/> |
|                    | Saturday | Saturday | 0:00      | 24:00   |                                     | <input type="button" value="Delete"/> |
|                    | Sunday   | Sunday   | 0:00      | 24:00   |                                     | <input type="button" value="Delete"/> |
| <b>ALWAYS</b>      |          |          |           |         |                                     |                                       |
|                    | Sunday   | Saturday | 0:00      | 24:00   |                                     |                                       |
| <b>Work Hours</b>  |          |          |           |         | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
|                    | Monday   | Friday   | 08:00     | 18:00   |                                     | <input type="button" value="Delete"/> |

Screenshot 70: Predefined schedules

Additional schedules can easily be added.

**Add New Schedule**

Name:

Times:

| From Day                                 | To Day                                   | From Time                                | To Time                                  |
|------------------------------------------|------------------------------------------|------------------------------------------|------------------------------------------|
| <input style="width: 80%;" type="text"/> |
| <input style="width: 80%;" type="text"/> |
| <input style="width: 80%;" type="text"/> |
| <input style="width: 80%;" type="text"/> |

Screenshot 71: Add a new schedule by specifying one or more date ranges and time ranges.

### NOTE

- » A single Schedule Object cannot specify different times that overlap. Time must be unique within the same Schedule Object.
- » The 'Start Time' and 'End Time' values must fall within the same day. In order to specify a time from one day to the next, use two lines: one from 'Start Time' to 24:00 on the first day, and the other from 00:00 to 'End Time' on the next day.

To add a new schedule:

1. Enter a meaningful name for the schedule.
2. For each time period, in the **From Day** field and the **From Time** field, specify a day of the week and time of day that marks the beginning of the time period.

3. In the **To Day** field and the **To Time** field, specify a day of the week and time of day that marks the ending of the time period.
4. Add multiple time periods if desired.
5. Click the **Add New Schedule** button.

#### NOTE

By default, there are four time period drop-down sets available to add schedule times. If you need to add more, save the schedule, then select the **Edit** button next to the newly created schedule. You will be presented with four additional sets of drop-downs to add more schedule times.

## Configuring schedule objects in the EMC

To create a Schedule in the Library

1. Go to **Library > Schedules**.
2. Click **Create new schedule in the library**.
3. In the **Name** section, key-in a name of the schedule item. The name must be unique in the tenant.
4. In the **Times** section, select the days of the week and the **Start Time** and **End Time** for this schedule.
  - To layer the time ranges, click **Add another time range**. For example, if you want to apply a schedule for Monday through to Friday from 9:00 to 17:00, but you need a different start and end time for weekends, you can add another range for Saturday and Sunday.
4. Click **Create**. The Schedule item is added to the Schedules Library category and is then available when defining [Policies](#) and [Virtual Circuits](#).

#### NOTE

By default, the **Always** schedule library item is pre-configured. You cannot change or edit it.

## 0.10.12 Configuring adaptive response limits

Adaptive response limits allow administrators to specify rules based on combinations of data transfer amounts and/or the amount of time on the network, and then establish priorities so that different policies can be applied both before and after the limit is reached. This allows you to provide data transfer caps, throttle data transfer after a limit has been reached, or to throttle only particular types of traffic after the limit has been reached.

#### VERSION INFO

In 7.0.1 and above releases in the 7.0 firmware product line, the option to set Adaptive Response Limits is based on the amount of time a user is using the network.

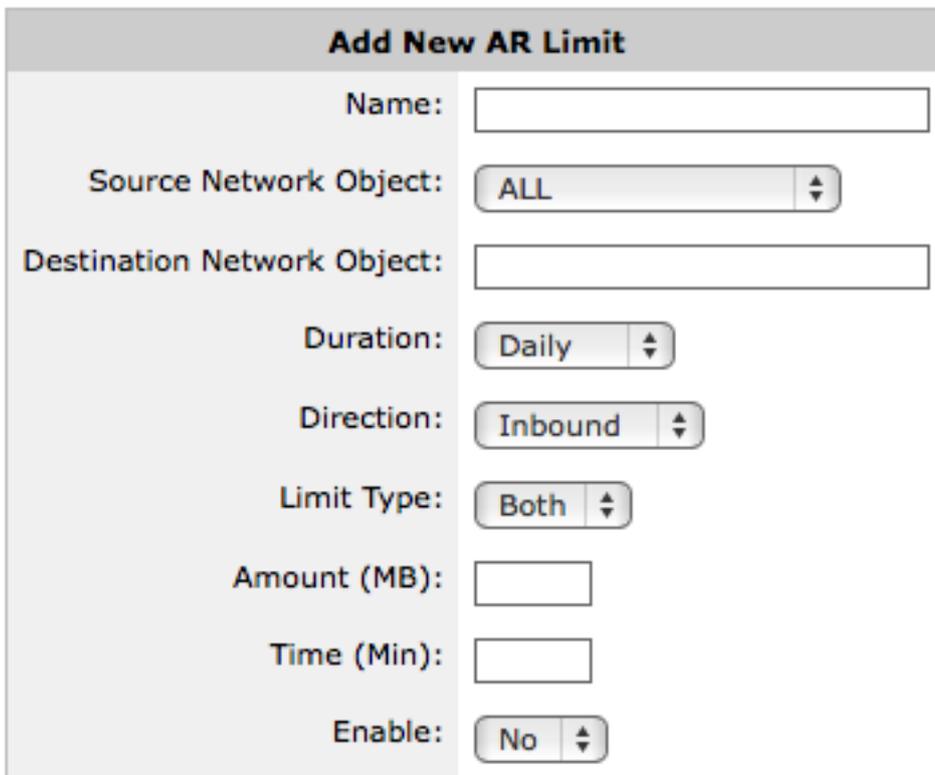
To implement such policies, the following steps are required:

1. Create a network object that defines what traffic is to be monitored.

The source network object can either be a static network object, which includes one or more subnets, or the source network object can be a dynamic network object mapped from an Active Directory group.

2. Create an adaptive response limit object. The adaptive response limit object allows administrators to specify the traffic to monitor using a network object, and what data limit or time limit should be applied and for what period. The appliance then dynamically creates a new network object that keeps track of the IP addresses that have exceeded their limit.

When a time limit is specified, the time is tracked in increments of 5 minutes and starts counting down from the first flow for a given user.



| Add New AR Limit            |                                      |
|-----------------------------|--------------------------------------|
| Name:                       | <input type="text"/>                 |
| Source Network Object:      | <input type="text" value="ALL"/>     |
| Destination Network Object: | <input type="text"/>                 |
| Duration:                   | <input type="text" value="Daily"/>   |
| Direction:                  | <input type="text" value="Inbound"/> |
| Limit Type:                 | <input type="text" value="Both"/>    |
| Amount (MB):                | <input type="text"/>                 |
| Time (Min):                 | <input type="text"/>                 |
| Enable:                     | <input type="text" value="No"/>      |

**Add New Limit**

Screenshot 72: Adding a new AR limit.

3. Create policies intended for the traffic matching the over-the-limit network object and policies intended for the traffic matching the source network object.

Ensure that the over-the-limit policy filters the traffic using the over-the-limit network object, and that the policy appears in the policy tree before the policies intended for the users who have not exceed their quota. Traffic attempts to match the policy tree nodes in a top-down order. Since IP addresses that have exceeded their quota will match either the destination or source network object, you need those that exceed their quota to be matched against the destination node first.

When creating the adaptive response limit object, you can create exceptions such that certain IP addresses, specified by one or more network objects, can be excluded from the limit rules. By editing the adaptive response limit object, you can specify which IP addresses are exempt from the rule. The exception network object can be internal or external. By creating an exception for an internal network object, those IP addresses will not have the limit applied to them. By creating an exception for an external network object, those IP addresses in the source network object will be excluded when they are conversing with an IP in the external exception network object.

You can also create an alert that sends an email to the Exinda Appliance email recipients when a specified quota threshold has been exceeded. The email will contain all the IP addresses that have exceeded the threshold. Note that the **Info Emails** checkbox must be checked for each recipient on the email configuration page.

**Add New Internal Exception Object**

Internal Network Object

**Add New External Exception Object**

External Network Object

**Add Alert Percent**

Percent

%

| Network Object | Exception | Action                                |
|----------------|-----------|---------------------------------------|
| IT subnet      | Internal  | <input type="button" value="Delete"/> |

| Alert | Action                                |
|-------|---------------------------------------|
| 50%   | <input type="button" value="Delete"/> |

To see examples, read [Quota Enforcement](#) in the Common Use Cases section.

Where do I find this configuration?

Go to **Configuration > Objects > Adaptive Response**.

To create an adaptive response limit object

1. Enter a name for the Adaptive Response Limit object.
2. From the **Source Network Object** drop-down list, select the desired network object. The source network object can be a static network object specified as subnets, or it can be a dynamic network object based on Active Directory users or user groups.
3. In the **Destination Network Object** field, type a name for the network object that will be dynamically created to store the IP addresses of the users that exceeded the limit.
4. From the **Duration** drop-down list, select the time period that the limit applies to. For example, if you want to apply a 10MB cap per day, select daily. The options are daily, weekly, monthly.
5. From the **Direction** drop-down list, select which traffic direction is counted towards the limit. The options are inbound, outbound, both.
6. From the **Limit Type** drop-down list, select whether you want the limit to be based on data volume, the amount of time on the network, or both (meaning whichever occurs first).
7. In the **Amount** field, specify the data volume limit, beyond which the IPs will be added to the destination network object. Specified in MB.
8. In the Time field, specify the time on the network limit, beyond which the IPs will be added to the destination network object. Specified in minutes.
9. If you are ready to make this adaptive response limit rule active, select **Yes** from the **Enable** drop-down list.
10. Click the **Add New Limit** button. The object is added to the list of configured adaptive response limit objects.

To add an alert notification when users reach a particular threshold

As an example, consider an educational institution that has a group of students who have IP addresses in the subnet 192.168.0.0/16. Each student is allowed 10GB data transfer (uploads and downloads) per month. The resident assistants are excepted from the rule as well as students who are trying to access a particular university server.

1. Edit an existing adaptive response limit object.
2. In the **Add Alert Percent** field, type the percentage of threshold that should trigger an email. The percentage can exceed 100.
3. Click the appropriate **Add Alert** button.

#### NOTE

Valid SMTP and email settings are required for email alerts.

The **Info Emails** checkbox needs to be checked for each email recipient on the Configuration > System > Network > Email tab.

To create user exceptions to the adaptive response limit rules

1. Edit an existing adaptive response limit object.
2. In the **Add New Internal Exception Object** or **Add New External Exception Object** section, select the network object that specifies the IPs that you would like excluded from the adaptive response limit rule. Selecting an internal network object will exclude the associated IP addresses from the limit rule. Selecting an external network object will exclude the hosts having conversations with particular external hosts from the limit rule. For example, consider an educational institution that has a group of students who have IP addresses in the subnet 192.168.0.0/16. Each student is allowed 10GB data transfer (uploads and downloads) per month. The resident assistants are excepted from the rule as well as students who are trying to access a particular university server.
3. Click the appropriate **Add Exception** button. Upon reaching the threshold, an email is sent that has content similar to:

Description: Adaptive Limit Threshold(s) exceeded for limit 'Student-Quota-Rule'

| IP Address | Current User | Limit | Usage    | Threshold | Timestamp (Canada/Eastern) |
|------------|--------------|-------|----------|-----------|----------------------------|
| 10.2.6.150 |              | 1M    | 6990.28k | 50%       | 2014/02/25 14:34:37        |

To create a network object

For more information, refer to [Adding network objects](#) (page 178).

To see which hosts have exceeded their quota

Consider an educational institution that has a group of students who have IP addresses in the subnet 192.168.0.0/16. Each student is allowed 10GB data transfer (uploads and downloads) per month.

Go to **Configuration > Objects > Network > Dynamic** and select the appropriate dynamic network object.

Name: Students-Over-Quota

| Adaptive Response |                     |                     |
|-------------------|---------------------|---------------------|
| IP Address        | User                | Time Added          |
| 10.2.6.32/32      | EXANET\Sam.wilson   | 2014/02/24 19:52:53 |
| 10.2.6.64/32      | EXANET\Rory.grey    | 2014/02/24 19:52:53 |
| 10.2.6.67/32      | -                   | 2014/02/24 19:52:53 |
| 10.2.6.137/32     | -                   | 2014/02/24 19:52:53 |
| 10.2.6.145/32     | EXANET\David.watson | 2014/02/24 19:52:53 |

When does the quota reset?

The Duration setting indicates the frequency at which the quota is reset.

- » Daily Duration resets at midnight
- » Weekly Duration resets at midnight Saturday night/Sunday morning
- » Monthly Duration resets at the 1st of the month.

If the quota is not reached for a given day, then the quota is reset. For example, consider Daily Duration, where the quota is reset at midnight. If the time limit was set to 90 minutes and the user started using the network at 11:30pm, then the system would allow 2 hours of continuous use, which includes 30 minutes until midnight, then another 90 minutes when the new day started at midnight.

### Creating adaptive response rules with CLI

Adaptive Response rules can be created using the CLI (in configure terminal mode):

```
adaptive limit <limit-name> network-object source <src> destination <dst>
adaptive limit <limit-name> amount <N (mb)>
adaptive limit <limit-name> duration <daily|weekly|monthly>
adaptive limit <limit-name> direction <inbound|outbound|both>
adaptive limit <limit-name> enable
```

#### EXAMPLE

Create an Adaptive Response rule which adds IP addresses from the static Students Network Object to the Dynamic Network Object Students-Over-Quota, once 200 MB has been downloaded per day.

```
adaptive limit Students-AR network-object source Students destination Students-Over-Quota
adaptive limit Students-AR amount 200
adaptive limit Students-AR duration daily
adaptive limit Students-AR direction inbound
adaptive limit Students-AR enable
```

See the following topics for more information:

---

---

## Adding a dynamic network object to Optimizer with CLI

The aim of this step is create a virtual circuit which references a dynamic network object created above.

Assuming we have created a Virtual Circuit named "Wan Inbound Choke" with reduced bandwidth, we can now reference the Dynamic Network Object created above using the following CLI command.

```
(config) # circuit default vcircuit "WAN Inbound Choke" destination Students-Over-Quota
```

## Disabling an adaptive response rule

To disable an adaptive response rule, run the following command. No IPs will belong to the destination network object, so any Optimizer virtual circuits or policies using the destination network object will effectively do nothing.

```
(config) # no adaptive limit Students-AR enable
```

## Excluding hosts or subnets from the quota

It is possible to configure Adaptive Response rules to exclude both internal or external hosts and subnets from the data transfer limits. This configuration option is available using the following CLI commands:

```
adaptive limit <limit-name> except network-object {internal|external} <network object>
```

The following examples illustrate how to exclude IP addresses or subnets from the Adaptive Response quota. The first example excludes an internal IP address that exists on the LAN-side of the Exinda Appliance. The second example excludes an entire subnet that exists on the WAN-side of the Exinda appliance.

### EXAMPLE

Create an Adaptive Response rule which adds IP addresses from the static Students Network Object to the Dynamic Network Object Students-Over-Quota once 200 MB has been downloaded per day, except for the IP address 192.168.0.50.

```
network-object IgnoreUser subnet 192.168.0.50 /32
network-object IgnoreUser location internal
adaptive limit Students-AR network-object source Students destination Students-Over-Quota
adaptive limit Students-AR amount 200
adaptive limit Students-AR duration daily
adaptive limit Students-AR direction inbound
adaptive limit Students-AR enable
adaptive limit Students-AR except network-object internal IgnoreUser
```

## EXAMPLE

Create an Adaptive Response rule which adds IP addresses from the static Students Network Object to the Dynamic Network Object Students-Over-Quota once 200 MB has been downloaded per day except for the DMZ subnet 203.122.212.128 /27.

```
network-object IgnoreDMZ subnet 203.122.212.128 /27
network-object IgnoreDMZ location external
adaptive limit Students-AR network-object source Students destination Students-Over-Quota
adaptive limit Students-AR amount 200
adaptive limit Students-AR duration daily
adaptive limit Students-AR direction inbound
adaptive limit Students-AR enable
adaptive limit Students-AR except network-object external IgnoreDMZ
```

## Other useful adaptive response CLI commands

The following command may be used to show Adaptive Response rules:

```
show adaptive limit <limit-name>
```

Adaptive Response evaluates rules every 5 minutes by default. IP addresses are added to destination dynamic Network Objects when the amount of traffic for the specified direction and duration exceeds the specified amount. Network objects are cleared at the end of the duration (e.g. daily, weekly or monthly). The following command can be used to change the frequency at which the rules are evaluated:

```
adaptive update-time <seconds>
```

Use the following command to show network objects created by Adaptive Response:

```
show network-object <network object>
```

The following command will clear all IPs from all Adaptive Response destination network objects. The network objects will be repopulated when rules are next evaluated.

```
adaptive clear
```

## 0.10.13 Configuring application performance score objects

The application performance score (APS) object is used to assess how network users enjoy the network performance experience of business-critical applications. The score, ranging between 0 and 10, where 0 is poor and 10 is excellent, indicates whether the app is performing as well as expected or is performing poorly. By creating an APS object, you specify an application to monitor. Optionally, you can also specify a network object so that the application is only monitored when observed on that part of the network. You set thresholds on one or more network metrics. Later, traffic for that application is assessed against those thresholds to determine how well the application is performing.

The appropriate thresholds for an application is unique for each network environment. You can manually set the thresholds for the network metrics or you can have the system automatically create threshold values by having the system observe traffic to determine reasonable baseline values. The metrics include network delay, server delay, round trip time, jitter, and network loss. Note that you can manually set the network loss metric, however, it will not be automatically be calculated during the baseline analysis. You can use one or more of these metrics in your APS object. Most applications use transactional protocols. Applications like Citrix XenApp server or Microsoft Remote Desktop use non-transactional protocols that send information between the client and server at arbitrary times. With these types of applications, the standard method of calculating the network delays and server delays does not produce an accurate metric. If the application uses a non-transactional protocol, you must specify that when creating APS object.

For the baselining analysis, traffic is analyzed during the specified period, and a set of metric thresholds is generated. The threshold recommendations target an APS of 8.5. If the application reports an APS below 8.5, the application is performing worse than the baseline. If no traffic is observed during the baselining period, then the appliance will automatically start another the baseline analysis for the next larger time period. Email will be sent for each unsuccessful baseline analysis.

**NOTE**

It is a best practice to start the baseline analysis during a time period when you would expect traffic for the application is typical. This will ensure that the baseline values accurately reflect the typical usage of the application. This means that if network conditions changes, it is recommended that the thresholds are re-evaluated.

**NOTE**

APS is not supported for small-packet applications like Citrix and RDP. The metrics are normalized as if the application runs with larger packet sizes, leading to larger values.

You can also set alerts so that you will be notified when the score drops below a certain threshold value. There is an alert trigger delay setting which requires that the score remains below the alert threshold for a specified period of time before triggering the alert. This prevents brief temporary poor scores from appearing like an emergency.

**Add New APS Object**

APS Name:

Application:

Network Object - Internal:

Network Object - External:

Alert Enable:

Alert Threshold:

Alert Trigger Delay:

Auto Baseline

Auto Baseline Period:

Non-Transactional Protocol

Add New APS Object

Cancel

Screenshot 73: The form to add a new APS object.

When editing the APS object, you can modify the alert configuration, restart the baselining operation, and modify the threshold values. If you change the network object settings, it is recommended that you re-evaluate the metric thresholds and possibly re-start a baseline.

### Edit APS Object

APS Name:

Application:

Network Object - Internal:

Network Object - External:

Alert Enable:

Alert Threshold:

Alert Trigger Delay:

Non-Transactional Protocol

### Baseline Info

Status: Stopped

Average Packet Size (bytes): 757

Traffic Seen (KB): 6001

Start Date: Mon Feb 10 09:00:00 EST 2014

End Date: Mon Feb 10 10:00:00 EST 2014

Auto Baseline Period:

### Scoring Metrics

| Metric                            | Config                           | Baseline |
|-----------------------------------|----------------------------------|----------|
| Normalized Network Delay (ms/kb): | <input type="text" value="48"/>  | 48       |
| Normalized Server Delay (ms/kb):  | <input type="text" value="22"/>  | 22       |
| Network Delay (ms):               | <input type="text" value="176"/> | 176      |
| Server Delay (ms):                | <input type="text" value="146"/> | 146      |
| Network Jitter (ms):              | <input type="text" value="26"/>  | 26       |
| Round Trip Time (ms):             | <input type="text" value="143"/> | 143      |
| Network Loss (%):                 | <input type="text"/>             | -        |

Screenshot 74: Editing an APS object.

## Creating an application performance score object

Use the instructions that follow to create a new APS object. During this set up, you can set a scope for the monitoring process. The scores can focus on specific internal and/or external network objects, or on ALL in one or both categories.

### Before you begin...

- » If you need to enable alerts, ensure that you have set Email on the **Configuration > System > Setup > Alerts** page. For more information, see the Exinda Web UI help.

» You also need to set up SNMP on the **Configuration > System > Network > SNMP** page. for more information, see the Exinda Web UI help.

## Creating an application performance score object in the Exinda Web UI

### To create the object:

1. Go to **Configuration > Objects > Service Levels > Application Performance Score**.
2. Click the **Add New APS Object** button.
3. In the **APS Name** field, type a name for the score.
4. In the **Application** list, select the application traffic to monitor.
5. Open the **Network Object - Internal** drop-down and either select a specific network object or select ALL.
6. Open the **Network Object - External** drop-down and either select a specific network object or select ALL.

### NOTE

By specifying both an internal and external network object, only the application conversations between the specified network objects is tracked.

7. If you want to be alerted when the application performance score drops below a particular threshold, set the following alert settings:
  - a. Ensure the **Alert Enable** checkbox is selected.
  - b. In the **APS Threshold** field, set a threshold value between 0 and 10.
  - c. In the **Alert Trigger Delay** field, specify how many minutes that the APS score to be below the threshold before the notification is sent.

### EXAMPLE

If the alert threshold is set to 7.0 and the alert trigger delay is set to 5 minutes, then the alert needs to be below 7.0 for 5 minutes before the alert is triggered.

8. If you need baselining to start immediately, select the **Auto Baseline** checkbox and select the **Auto Baseline Period**.
9. If the application uses a non-transactional protocol for traffic between the client and server, such as Citrix XenApp Servers or Microsoft Remote Desktop, select the **Non-Transactional Protocol** checkbox.
10. Click **Add New APS Object**. The object is added to the list of configured APS objects

### Editing APS objects created in the Solution Center

You can modify the thresholds either manually or by re-starting the baseline analysis. You can also enable an alert, and change the internal and/or external network object setting to monitor the application for a subset of your network.

### IMPORTANT

Do not change the application being monitored, and do not change the name of the APS object.

## Related Topics

- » [Manually creating APS thresholds](#)
- » [Automatically calculating APS thresholds](#)
- » [How to know if baselining is in progress](#)
- » [How the Performance Metric thresholds are calculated](#)

## Creating Application Performance Score (APS)

Before you can start to monitor the performance of the important applications on your network, you need to identify those applications in the Configuration Library. The Configuration Library comes with definitions for a very large number of supported applications.

1. In the EMC interface, click **Library > Application Performance Scores**.
2. On the right side, select the **Create new application performance...** link to open the APS set up page.
3. Click **Name** to expand the section. Provide a meaningful name for the new performance score.

Tenant Overview | Configured Appliances | Library | Not Deployed

Library

- Circuits
- Circuit Types
- Virtual Circuits
- Policy Sets
- Policies
- Network Objects
- Applications
- Application Groups
- Schedules
- Application Performance Score
- Service Level Agreements
- VLANs

### Application Performance Score

Application Performance Scores provide a method for evaluating how well applications on your network are performing.

▼ Name: E-mail

Name

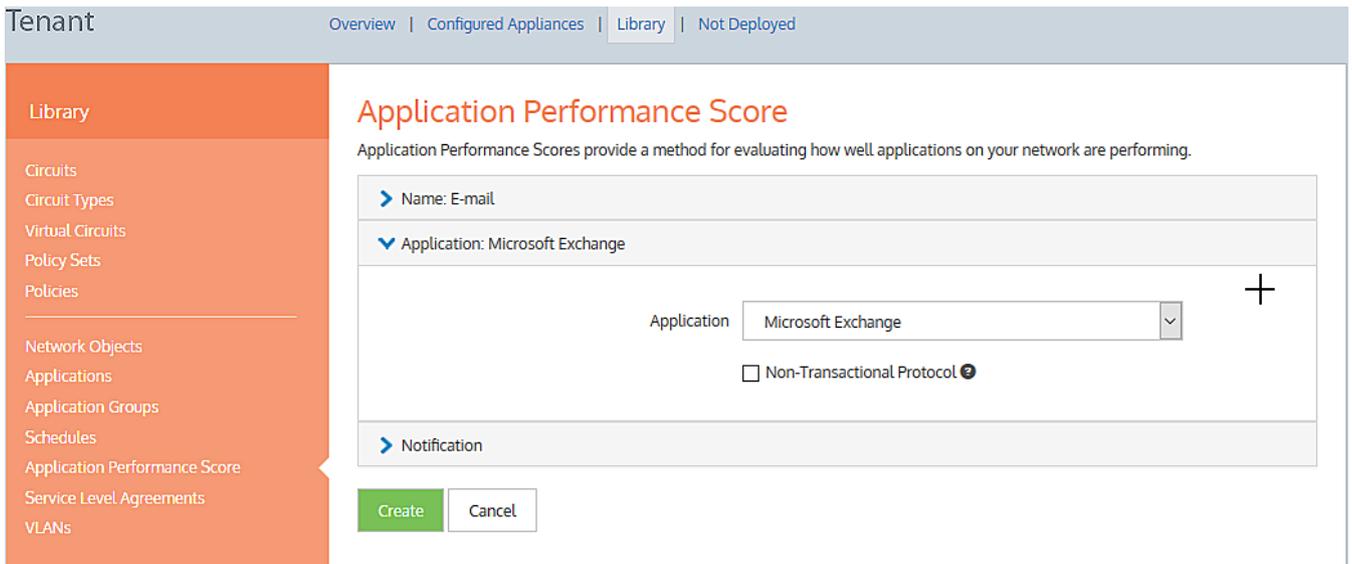
▶ Application

▶ Notification

Create Cancel

Screenshot 75: Setting up a performance score for email services.

4. Click **Application** section to expand.
5. Use the **Application** drop-down list to select the application you need.

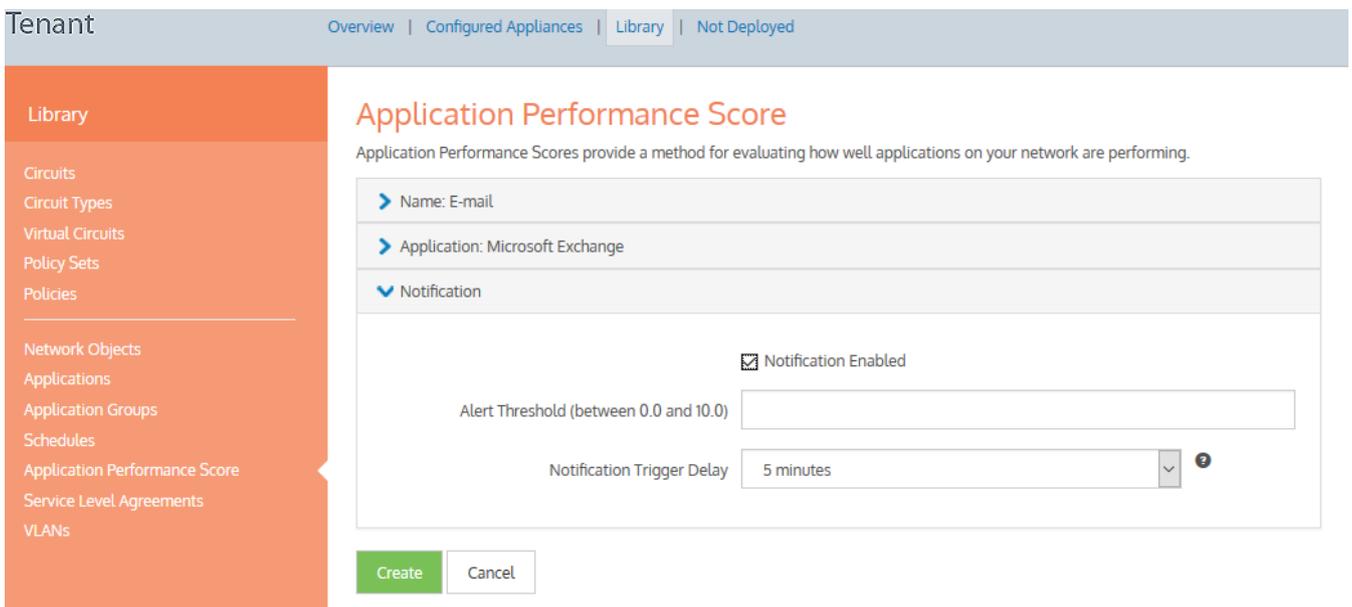


Screenshot 76: Selecting an application

#### IMPORTANT

Before proceeding, you should consider the type of protocol used by the application. If it uses a non-transactional protocol, you should select the checkbox. For more information, click the ? button.

6. Click **Notification** to expand the section.



Screenshot 77: Configuring notification settings

7. Configure the following options:

| Option               | Description                                                                                                                                                                                                                                                                   |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Notification Enabled | Select if you want to be notified when the Alert Threshold is exceeded.                                                                                                                                                                                                       |
| Alert Threshold      | Type an alert level between 0.0 and 10.0. This is a measure of how important is the service the application provides. For example, an application that provides real-time communications would require a higher alert level than one whose responsiveness is not as critical. |

| Option                     | Description                                                                                                                                                                                    |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Notification Trigger Delay | click the down-arrow and select an acceptable delay period. The notification is only triggered if the condition persists for the entire delay period. Click the ⓘ button for more information. |

8. Click **Create**. The screen refreshes to show an entry for the new performance score.

#### NOTE

You can also create APS directly from the **(desired appliance group) > Application Performance Scores** interface. In this case, the APS is applied to the appliance group, and it is also saved in the configuration library.

### Adding APS to the appliances in the EMC

After completing the creation of the APS item definitions in the Configuration Library, you can then apply them to the appliances. When applying the scores, this is a global application; all appliances in the same appliance group receive the same configuration. After applying the scores, you must then push the configuration to the appliances in order to get any notifications.

1. Go to **Configured Appliances > Application Performance Scores**.
2. Click **Add application performance from the library...** link.
3. On the **Add Application Performance Score from Library** pop-up, expand the **Application Performance Score** section.
4. Use the **Application Performance Score** drop-down list to select the APS you need to apply.

### Add Application Performance Score from Library

Screenshot 78: Selecting APS to apply

#### NOTE

If you have appliances that do not have the most recent firmware, you may notice the interface reference to the application name is highlighted in red, when you try to add an APS to the configured appliances. You can hover over ⓘ to view details about the issue on a tooltip.

5. Expand the **Network Object...** section.

## Add Application Performance Score from Library

Application Performance Score: E-mail

Network Object: Measure from All to All

Internal Network Object All

External Network Object All

Baseline: For next hour after pushing the configuration

Add Application Performance Score to 'Configured Appliances' Cancel

Screenshot 79: Selecting internal and external network objects

- Use the **Internal Network Object** drop-down list to select the object that fits your needs.

### TIP

The internal objects are the [Local Network Objects](#) defined in the Configuration Library. The default is **All**, but you can choose from any of the others that have been defined.

- Use the **External Network Object** drop-down list to select the previously Network Object that fits your needs.

When complete, the title of the Network Object section changes to summarize the settings you have made.

6. Expand the **Baseline** section and choose the duration of the baseline calculation.

### NOTE

To establish a baseline for the performance of an application, its performance in the network must be monitored for period of time. The Baseline Length you define is the initial monitoring period, but if the baselining fails, the initial monitoring is automatically extended. The calculation will start at the beginning of the hour the configuration is pushed.

For example, if it is pushed at 4:37, the baseline calculation will start at 4:00. If there is no data for that period, it will extend the duration to the next longest period.

## Add Application Performance Score from Library

> Application Performance Score: E-mail

> Network Object: Measure from All to All

▼ Baseline: For next hour after pushing the configuration

Baseline will be performed for the selected amount of time after the configuration is pushed to the appliance. If baseline does not succeed within the time period the next biggest time period is started.

Baseline Length

Add Application Performance Score to 'Configured Appliances' Cancel

Screenshot 80: Choose the duration of the baseline calculation

- When finished, click **Add Application Performance Score to 'Configured Appliance'**. The screen refreshes to show an entry for the new performance score.
- Repeat each of the preceding steps to apply all of the APS definitions you need.

### IMPORTANT

After applying all the APS definitions, you must push the configuration to all or any particular appliance group. For more information, refer to [Sending configuration changes to the appliances](#) (page 176).

### Removing APS from the Appliances

If any existing APS is no longer required, you can remove it from the appliance configuration. Removing an APS in this way does not remove it from the Configuration Library; the APS item remains there for future reuse. If you do want to remove it from the library, see [Maintaining APS in the Configuration Library](#). But you must remove an APS item from the appliance configuration before deleting it from the Configuration Library.

- Go to **Configured Appliances > Application Performance Scores**. The currently assigned APS items appear on the page.
- In the **Name** column, locate the APS item and use the **X** icon on the right to delete.

Screenshot 81: Confirming removal

- Confirm and click **Remove**.

Are you sure you want to remove this application performance score from the appliance group?

Note: It will still be available in the application performance score library.

Don't ask when removing application performance scores again.

**Remove**    Don't Remove

Maintaining APS in the EMC configuration library

Over time, it may become necessary to modify the Application Performance Score items stored in the Configuration Library. When modifying an APS item, you are modifying its use wherever it has been applied.

1. In the EMC interface, click **Library > Application Performance Scores**. A listing of the currently defined performance scores appears on the right.
2. In the **Name** column, click on the name of the APS item you need to modify.

## Application Performance Score (APS) Library

Application Performance Scores provide a method for evaluating how well applications on your network are performing.

[+ Create new application performance score in the library ...](#)

| Name                   | Application        | Notification |
|------------------------|--------------------|--------------|
| <a href="#">E-Mail</a> | Microsoft Exchange | ✓            |

3. Do the following, as needed:

## Application Performance Score

Application Performance Scores provide a method for evaluating how well applications on your network are performing.

▼ Name: E-Mail

Name

▶ Application: Microsoft Exchange

▶ Notification: Notify when the application's performance score goes below 6 for 5 minutes

▶ In Use: 1

- Modify the APS **Name**.
- Change the **Application**
- Update the **Notification** definition

4. Before completing the modification, expand the **In Use** section to see how the APS items is currently being used.

5. Click **Update in Library**.

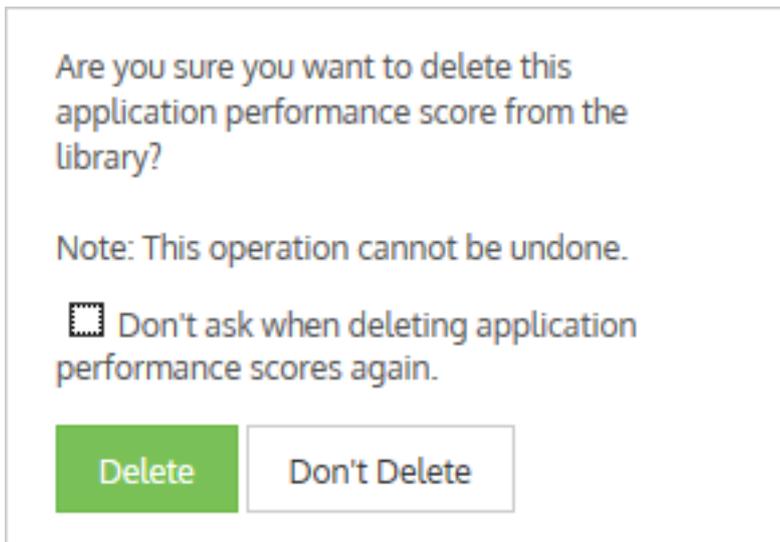
### Deleting APS from the Configuration Library

Over the time, it may become necessary to delete an Application Performance Score item stored in the Configuration Library.

1. Go to **Library > Application Performance Scores**.
2. Locate the APS item to delete from the list and click the  icon.

#### TIP

You can only delete those APS items that are not currently in use. If a  icon appears in the row of an APS item you wish to delete, it means that the item cannot be deleted because it is in use. If you still want to delete it, you first need to unlink it from the appliances. For more information, refer to [Maintaining APS in the EMC configuration library](#) (page 232).



Screenshot 82: Confirming deletion of APS from the Configuration Library

3. Confirm and click **Delete**.

### How performance metric thresholds are calculated

Network performance metrics are calculated based on the observed traffic. Each threshold is calculated to be 0.85 of a standard deviation above the average observation for that metric. This ensures that the calculated thresholds target is an APS of 9.0. If the application reports an APS below 9.0, the application is performing worse than the baseline.

### Configuring APS thresholds manually

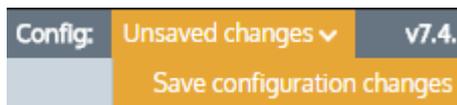
Metric thresholds can be set manually when initially creating the APS object or upon editing an APS object even if they were automatically determined by the baselining operation. For example, if the baselining operation set all of the thresholds and you really only care about round trip time, normalized server delay, and normalized network delay, then you can remove the threshold settings for the other metrics.

1. Go to **Configuration > Objects > Service Levels > Application Performance Score**.
2. On the **Add New APS Object** form, uncheck the **Auto Baseline** checkbox. Note if a baseline analysis is running, you'll need to press the **Stop Baseline** button. The threshold values are only editable there is not a baseline running. The metrics will appear on the screen. Or edit the APS object in the list, then on the **Edit APS Object** form, the **Scoring Metrics** appear at the bottom of the form.
3. Enter or modify the values for the metrics that you are interested in setting thresholds for. Note that any metric that does not have a threshold set will not be analyzed when calculating the APS score.
  - Network delay – the time taken for data to traverse the network (on the wire) in one direction from the client through the Exinda appliance to the server (or in the opposite direction) in ms
  - Server delay – the time taken for a server to respond to the request in ms
  - Normalized network delay – the time taken for data to traverse the network in one direction, where the delay is measured independent of the transaction size by assuming a normalized packet size of 1024 bytes
  - Normalized server delay – the time taken for a server to respond to the request, where the delay is measured independent of the transaction size by assuming a normalized packet size of 1024 bytes
  - Round-trip time – the time taken for a packet to travel from a client through the Exinda appliance to the server and back

- Jitter – the measure of variability of network delay, defined as one standard deviation of normalized network delay
- Inbound loss – the percentage of packet loss on inbound traffic
- Outbound loss – the percentage of packet loss on outbound traffic

#### 4. Click **Apply Changes**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.



### Configuring automatic APS threshold calculation

The baselining process can be started when initially creating the APS object or upon editing an APS object. At any time you can restart the baselining process if you would like the system to recalculate the thresholds.

1. Go to **Configuration > Objects > Service Levels > Application Performance Score**.
2. On the **Add New APS Object** form, ensure that the **Auto Baseline** checkbox is checked and set how long you want the system to observe traffic when calculating the thresholds by using the **Auto Baseline Period** drop-down list. Select the time period for the baseline based on how popular the application is. For example, if there is a lot of HTTP traffic on the network, the 1 hour period will be long enough to analyze traffic and create an accurate baseline. For an application that is not used very often, use the 1 week baseline period to ensure that enough traffic is analyzed to generate baseline recommendations.
3. Or on the **Edit APS Object** form, set the **Auto Baseline Period** and click **Start Baseline**.

#### NOTE

The Network Loss metric is not calculated during the baseline analysis.

#### NOTE

If no traffic matching this APS object is observed during the baseline period, the appliance restarts the baseline analysis for the next larger time period. For example, if no traffic observed during the one hour period, the traffic continues to be analyzed for one day. If no traffic is observed during the one day period, then the traffic is analyzed for a week. If the traffic is analyzed for one week and no traffic has been transferred, the auto baseline analysis stops. Each time the system unsuccessfully baselines the traffic (that is, when no traffic is observed during the auto baseline period), an email notification is sent to the users configured on the **Configuration > System > Network > Email** page.

#### Related Topics

- » [Creating an Application Performance Score object](#)
- » [Manually creating APS thresholds](#)
- » [How to know if baselining is in progress](#)
- » [How the Performance Metric thresholds are calculated](#)

## Checking if baselining is in progress

On the **Application Performance Score** configuration tab, the list of APS objects is shown. If the APS is currently baselining the application traffic, there will be a green checkmark in the **Auto Baseline** column.

Press the **Edit** button for the APS object. The **Baseline Info** section specifies the status (Running or Stopped) and the Start and End Date and time of the baseline period. Note that it also shows the average packet size and the amount of traffic seen.

### 0.10.14 Configuring an application performance metric object

The Application Performance Metric (APM) objects are used to monitor particular application performance metrics. By creating an APM object, you indicate which application to monitor. Optionally, you can also specify a network object so that the application is only monitored when observed on that part of the network. You set a threshold on a single network metric. Later, traffic for that application is assessed against that threshold to determine how well the application is performing. An alert is triggered when the threshold is exceeded for a given length of time.

The following metrics are available:

- » bytes lost
- » network delay
- » server delay
- » transaction delay
- » normalized network delay
- » normalized server delay
- » normalized transaction delay
- » round trip time
- » tcp connections aborted
- » tcp connections ignored
- » tcp connections refused
- » tcp connected started.

Add New APM Object

|                            |                                                                   |
|----------------------------|-------------------------------------------------------------------|
| APM Name:                  | <input style="width: 100%;" type="text"/>                         |
| Metric:                    | <input style="width: 100%;" type="text" value="round-trip-time"/> |
| Application:               | <input style="width: 100%;" type="text"/>                         |
| Network Object - Internal: | <input style="width: 100%;" type="text" value="private net"/>     |
| Network Object - External: | <input style="width: 100%;" type="text" value="ALL"/>             |
| APM Threshold:             | <input style="width: 100%;" type="text" value="0"/>               |
| Alert Trigger Delay:       | <input style="width: 100%;" type="text" value="5 minutes"/>       |
| Alert Enable:              | <input type="checkbox"/>                                          |

Screenshot 83: Adding a new APM object.

**NOTE**

APM values are not shown on any report; they are used solely to generate alerts.

Use the following instructions to create an APM object.

Before you begin...

- » If you need to enable alerts, ensure that you have set Email on the **Configuration > System > Setup > Alerts** page. For more information, see the Exinda Web UI help.
- » You also need to set up SNMP on the **Configuration > System > Network > SNMP** page. for more information, see the Exinda Web UI help.

To create an APM object

1. Go to **Configuration > Objects > Service Levels > Application Performance Metric**.
2. Click the **Add New APM Object** button.
3. Type a name for the APM object.
4. Select the metric that you need to monitor. The following metrics are available:
  - **bytes-lost** — Bytes lost due to retransmissions.
  - **network-delay** — The time taken for data to traverse the network.
  - **server-delay** — The time taken for a server to respond to a request.
  - **transaction-delay** — The total time for a transaction (network delay + server delay)

- **normalized-network-delay** — The time taken for data to traverse the network where the packet size is normalized to 1024 bytes.
- **normalized-server-delay** — The normalized measure of the time taken for a server to respond to a transaction request.
- **normalized-transaction-delay** — The normalized measure of the time taken for a client request to be sent to a server, and the server's reply to be received by the client.
- **round-trip-time** — The time taken for a packet to travel from a device, cross a network, and return.
- **tcp-connections-aborted** — The number TCP connections reset after the connection is established. (RST from client or server)
- **tcp-connections-ignored** — The number TCP connections that expire in the SYN-SENT state. No response is received from the server.
- **tcp-connections- refused** — The number TCP connections that are reset before the connection is established. (RST in SYN-SENT state)
- **tcp-connections-started** — The number of TCP connections initiated.

5. In the **Application** list, select the application traffic to monitor.
6. If you want to just monitor the application for a particular internal network object, specify the desired internal network object; otherwise select ALL.
7. If you want to just monitor the application for a particular external network object, specify the desired external network object; otherwise select ALL. By specifying both the internal and external network object, only the application conversations between the specified network objects will be tracked.
8. Select the **Alert Enable** checkbox.
9. In the **APM Threshold** field, type the threshold that will trigger an alert if the score drops below that value.
10. In the **Alert Trigger Delay** list, select how long the metric needs to remain below the threshold before the alert is sent. For example, if the alert is tracking the number of bytes lost, the threshold is set to 100, and the alert trigger delay is set to 5 minutes, then the number of bytes lost needs to be above 100 for 5 minutes before the alert is triggered.
11. Set the threshold for the APM metric. The units of the threshold is relative to the metric being measured. That is, delays and round trip time are measured in milliseconds, tcp connections and bytes lost are counts.
12. Click **Add New APM Object**. The object is added to the list of configured APM objects.

## 0.11 Monitoring your network

After installing and configuring your Exinda Appliance you can monitor your network, gaining full visibility into the applications users access, inbound traffic, outbound traffic and network throughput. Before customizing the Optimizer, best practices suggest letting the Exinda Appliance collect enough data to help you make informed decisions and policies.

---

|                                                      |     |
|------------------------------------------------------|-----|
| 0.11.1 Dashboards .....                              | 239 |
| 0.11.2 Monitoring network traffic in real time ..... | 244 |
| 0.11.3 Monitoring network interfaces .....           | 256 |
| 0.11.4 Monitoring network throughput .....           | 259 |
| 0.11.5 Monitoring service levels .....               | 261 |

---

|                                                              |     |
|--------------------------------------------------------------|-----|
| 0.11.6 Monitoring applications .....                         | 272 |
| 0.11.7 Monitoring network users .....                        | 282 |
| 0.11.8 Monitoring hosts traffic volume .....                 | 285 |
| 0.11.9 Monitoring network conversations .....                | 288 |
| 0.11.10 Monitoring subnets .....                             | 291 |
| 0.11.11 Monitoring virtual circuits .....                    | 296 |
| 0.11.12 Monitoring the effects of controls .....             | 299 |
| 0.11.13 Monitoring optimization reports .....                | 307 |
| 0.11.14 Monitoring Exinda Appliance system performance ..... | 312 |
| 0.11.15 Viewing monitoring statistics .....                  | 321 |

## 0.11.1 Dashboards

The Exinda Web UI provides dashboards you can use to monitor the operation of an Exinda Appliance. One dashboard displays system health and status information about the Exinda Appliance. The other dashboard provides statistical data to show the benefits and impact of the Exinda Appliance in your network.

### System dashboard

The System dashboard shows system information, the state of system alarms as well as a summary of other Exinda appliances and their respective reduction statistics. The dashboard answers questions, such as "Are there any issues with the NICs, or CPU utilization, or SMB signed connections, etc? What is this appliance licensed for? What is this appliance's host ID? What is the reduction ratio between this appliance and others that it's accelerating with?"

| Hostname: exinda-22061 |                                                    | Alarm                                                | Status        | Last Triggered | Count |
|------------------------|----------------------------------------------------|------------------------------------------------------|---------------|----------------|-------|
| Hardware Series:       | 2061                                               | <a href="#">CPU Utilization</a>                      | OK            |                |       |
| Licensed Model:        | <a href="#">Exinda 2861 (1.000/0.020Gbps - HP)</a> | <a href="#">System Disk Full</a>                     | OK            |                |       |
| SS Expiry Date:        | Aug 31, 2018                                       | <a href="#">Memory Paging</a>                        | OK            |                |       |
| Host ID:               | 00900b2e7a32                                       | <a href="#">Bridge Link</a>                          | OK            |                |       |
| Timezone:              | <a href="#">Etc/UTC</a>                            | <a href="#">Bridge Direction</a>                     | OK            |                |       |
| System Uptime:         | 6d 20h 55m 26.592s                                 | <a href="#">Link Negotiation</a>                     | OK            |                |       |
| Scheduled Jobs:        | No scheduled jobs.                                 | <a href="#">NIC Problems</a>                         | OK            |                |       |
| Memory Usage:          | <a href="#">86.29% of 3816MB</a>                   | <a href="#">NIC Collisions</a>                       | OK            |                |       |
| CPU Usage:             | <a href="#">7%</a>                                 | <a href="#">NIC Dropped Packets</a>                  | OK            |                |       |
| Database Status:       | <b>Running</b>                                     | <a href="#">SMB Signed Connections</a>               | OK            |                |       |
|                        |                                                    | <a href="#">Redundant Power</a>                      | Not Available |                |       |
|                        |                                                    | <a href="#">Redundant Storage</a>                    | Not Available |                |       |
|                        |                                                    | <a href="#">Max Accelerated Connections Exceeded</a> | DISABLED      |                |       |
|                        |                                                    | <a href="#">Asymmetric Route Detection</a>           | OK            |                |       |
|                        |                                                    | <a href="#">MAPI Encrypted Connections</a>           | OK            |                |       |

Screenshot 84: The system dashboard displays information about an Exinda Appliance and its peers.

The status of the appliance database is displayed as **Database Status**. The possible statuses include:

- » **Starting**—The database is initializing, and it is waiting for a response from the system on available storage.
- » **Running**—The database is operating.
- » **Upgrading**—The database has started, but is being upgraded.

- » **Downgrading**—The database has started, but is being upgraded.
- » **Stopped**—The database is stopped.
- » **Error**—The database cannot be accessed. This typically appears when there is a problem with the upgrade or downgrade of the database.
- » **Unknown**—The state of the database is unknown.

## Benefits Dashboard

The Benefits Dashboard exposes a set of widgets arranged on a dashboard that shows high level information about your network traffic. The dashboard answers questions, such as "What are the dominant application groups on my network? Are recreational apps using a large amount of my bandwidth? Is my link saturated? Am I getting the reduction that I'm expecting?" The dashboard may also show a recommendation. The Exinda Appliance analyzes network traffic and makes recommendations based on what it learns.

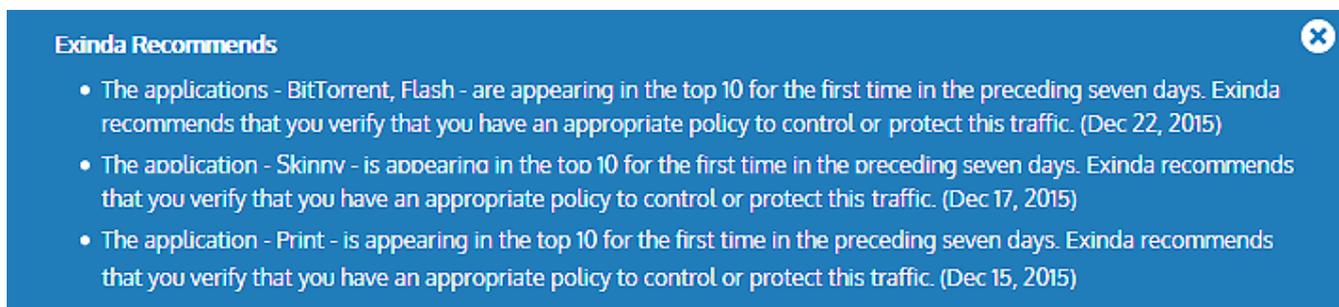
Widgets can be hidden to customize the dashboard to only include widget(s) relevant to you. To add a hidden widget, click the '**Add More**' link at the top right of the dashboard. If the '**Add More**' link is not visible, then all available widgets are displayed. Widget settings and layouts are retained between log-ins.

The dashboard can be captured and converted to PDF by clicking on the PDF icon at the top-right of the interface.

## Exinda recommends

Every night after midnight, your Exinda Appliance analyzes the traffic it saw during the previous day and, if there was something remarkable or unusual, it makes a recommendation, displays it on the dashboard and sends it to the email addresses configured in **Network Setup > Email**.

Each recommendation includes the date of the traffic data. Dismiss the recommendation by clicking the close button. To view the last three recommendations made, double-click the Exinda logo in the header bar on the dashboard.



**Exinda Recommends**

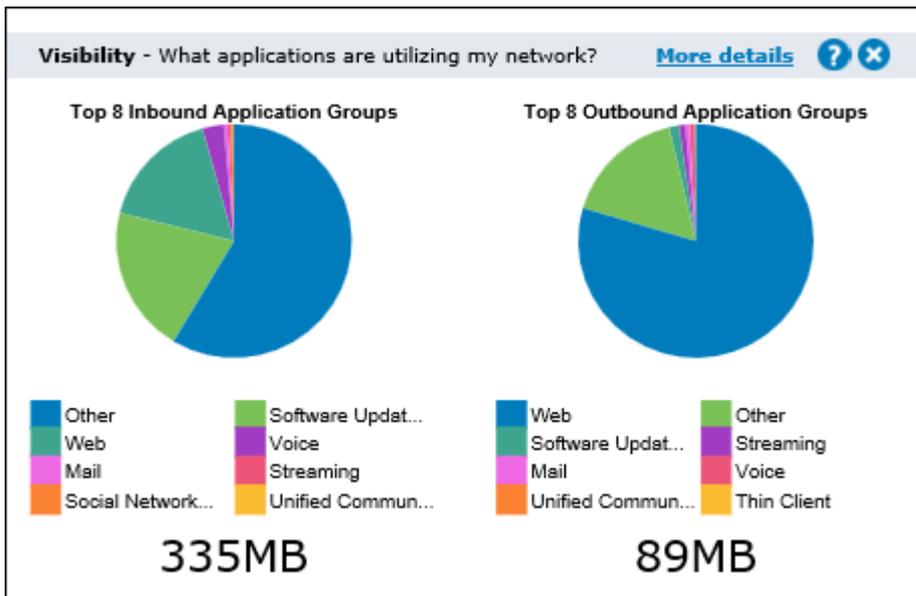
- The applications - BitTorrent, Flash - are appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you verify that you have an appropriate policy to control or protect this traffic. (Dec 22, 2015)
- The application - Skinnv - is appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you verify that you have an appropriate policy to control or protect this traffic. (Dec 17, 2015)
- The application - Print - is appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you verify that you have an appropriate policy to control or protect this traffic. (Dec 15, 2015)

Screenshot 85: Example of Exinda recommendation messages in the dashboard

## Visibility

Visibility gives you insight into the traffic on your network so you can effectively control or protect it. The visibility graphs show application groups utilizing the network. These graphs answer questions such as, "Are streaming applications for music and videos choking the network? Are data backups overrunning the network?"

Click the drill down link to see which apps are in an application group.



### Reduction

Reduction measures the amount of redundant data removed from the network, increasing free capacity. Data previously seen by the system is "remembered" and delivered from the local appliance rather than end-to-end from server to client resulting in a reduction in the amount of data sent across the network. The reduction ratio compares After Exinda to Before Exinda.

Reduction Ratio = (Data Transfer Size Before Exinda - Data Transfer Size After Exinda) / Data Transfer Size Before Exinda.

| Reduction - How much bandwidth have I saved? |             |             |                                                                                     | <a href="#">More details</a> ? x |
|----------------------------------------------|-------------|-------------|-------------------------------------------------------------------------------------|----------------------------------|
| Application                                  | LAN Data    | WAN Data    | Reduction Ratio (%)                                                                 |                                  |
|                                              | <b>16MB</b> | <b>12MB</b> |  | <b>21%</b>                       |
| <a href="#">HTTP</a>                         | 4MB         | 1MB         |  | 59.24%                           |
| <a href="#">CIFS</a>                         | 12MB        | 11MB        |  | 9.19%                            |
| <a href="#">Ports</a>                        | 5MB         | 2MB         |  | 36.6%                            |

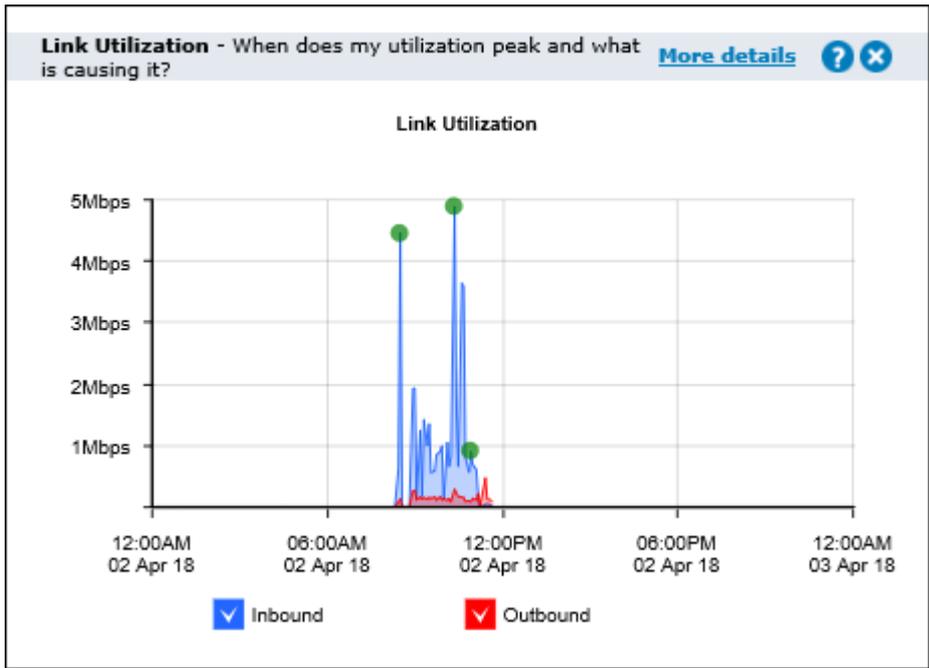
#### EXAMPLE

A ratio of 40% means a transfer that once put 100MB of load onto the WAN, now puts 60MB of load on the WAN. I.e. 40% less.

### Link Utilization

The link utilization chart shows the throughput through the Exinda Appliance over time for inbound traffic and outbound traffic. It also shows the top three conversations at each of the utilization peaks. It answers questions such as, "Are you receiving the bandwidth that you are paying for? When is my throughput peaking? What could be causing bottlenecks on my network?"

Hover over the green dots on the throughput peaks to see the top conversations which may be the cause of the network bottlenecks. This level of visibility allows IT professionals to address root causes instead of mistakenly treating symptoms, e.g. buying more bandwidth to cope with peak load.



### Recreational

Having visibility into key recreational applications is the first step to managing them. These applications are generally undesirable because they can impact the performance of key business applications, negatively impact customer experiences, reduce productivity, introduce viruses to the network and enable downloading of illegal or copyrighted material.

**Recreational** - How much recreational usage is there? [More details](#) ? X

| Application                       | Hosts    | Time          | Data       |
|-----------------------------------|----------|---------------|------------|
|                                   | <b>3</b> | <b>5m 30s</b> | <b>3MB</b> |
| <a href="#">Games</a>             | 1        | 20s           | 0MB        |
| <a href="#">Instant Messaging</a> | 0        | 0s            | 0MB        |
| <a href="#">P2P</a>               | 0        | 0s            | 0MB        |
| <a href="#">Social Networking</a> | 2        | 1m 50s        | 1MB        |
| <a href="#">Streaming</a>         | 2        | 3m 20s        | 2MB        |

### Prioritization

This ratio tells you how often critical applications were prioritized (also referred to re-ordering or re-queuing). A high percentage means that the system is prioritizing more often to ensure performance of your applications. A high percentage also means that by turning off optimization there is a higher probability that your critical applications will suffer.

$$\text{Prioritization Ratio} = \text{Number of Packets Re-ordered} / \text{Number of Total Packets}$$

**Prioritization** - How often were my critical applications prioritized? [More details](#) ? x

**Prioritization Rate: 6.6%**

**EXAMPLE**

A ratio of 40% means 40% of the packets on your network were re-ordered. That means that non critical data was queued so that business critical data could jump the queue and be delivered in the order that the business requires.

**Time Savings**

This table shows the improvement in transfer time due to WAN optimization. The Before time is the total amount of time an application would have taken to transfer data without WAN optimization. The After value is the actual amount of time taken with WAN optimization. The difference between the Before and After values is then shown as a time value and as a percentage.

$$\text{Before} = \text{LAN Data} / \text{WAN Throughput}$$

**Time Savings** - How much time is my optimization saving? [More details](#) ? x

| Application  | Before     | After      | Saving     | %          |
|--------------|------------|------------|------------|------------|
| Quicktime    | 9m         | 4m         | 5m         | 51%        |
| Vimeo        | 10m        | 5m         | 5m         | 50%        |
| Youtube      | 24m        | 18m        | 6m         | 27%        |
| <b>Total</b> | <b>43m</b> | <b>27m</b> | <b>16m</b> | <b>36%</b> |

**EXAMPLE**

A Savings value of 50% means that the time taken to transfer an application's data was reduced by half.

**Recommendations**

The Exinda appliance analyzes the traffic and makes recommendations based on what it learned about your traffic. These recommendations appear on the dashboard.

**Exinda Recommends** x

- The applications - BitTorrent, Flash - are appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you verify that you have an appropriate policy to control or protect this traffic. (Dec 22, 2015)
- The application - Skinnv - is appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you verify that you have an appropriate policy to control or protect this traffic. (Dec 17, 2015)
- The application - Print - is appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you verify that you have an appropriate policy to control or protect this traffic. (Dec 15, 2015)

The following are examples of the kinds of recommendations that the Exinda may make:

» "The applications "app-name-1" and "app-name-2" are appearing in the top 10 for the first time in the preceding seven days. Exinda recommends that you create a policy to control or protect this traffic."

Every night, the Exinda appliance looks at the top ten applications by data volume and determines if any of the applications are new to the top ten. That is, when looking at the data for the day, have any of the top ten applications not been in the daily top ten for the seven preceding days.

» "The circuit "circuit-name" has traffic that was not caught by a virtual circuit. Exinda recommends that you investigate by looking at the Virtual Circuits monitor or Real Time monitor, then redefine your virtual circuits to capture all of the circuit data."

Every night, the Exinda appliance looks at the traffic and determines whether traffic is showing up in any Auto Catch-all virtual circuits. It then reports traffic in the Auto-Catch-all virtual circuit if > 1% of the traffic of your entire appliance is caught in the single Auto Catch-all virtual circuit.

#### NOTE

When traffic is processed by the Optimizer policy tree, if traffic is not caught by a virtual circuit, it will be counted against an **Auto Catch-all** virtual circuit. There can be one Auto Catch-all virtual circuit per circuit. Also if traffic is not caught by a circuit, it will be counted against an **Auto Catch-all** virtual circuit within an **Auto Catch-all** circuit. To learn more about the policy tree and the Auto Catch-all circuits and virtual circuits, see [Policy Tree](#).

To diagnose the issue, you can see the Auto Catch-all circuit and virtual circuit on the **Virtual Circuit** monitor page. Auto Catch-all virtual circuit traffic also appears on the **Real Time Conversations** monitor page. Toggle on **Show Policies** , to show which virtual circuit and policy the traffic is being caught by.

### 0.11.2 Monitoring network traffic in real time

This section describes real-time reporting with the Exinda Web UI. The real-time monitors display information related to traffic that has passed through monitored links during the previous 10 seconds.

There are several views to help you understand real-time network traffic. These include traffic by applications, by hosts (and users), by conversations, and by reduction per application. Typically, conversations in real time are the most valuable view, as each conversation is shown separately rather than collapsing across an application or host. Also the conversation view allows you to filter the view by IP address or subnet.

When investigating a current issue, the real time monitors allow you to answer questions like:

- » My link is congested; which conversations, applications, or hosts may be contributing to the congestion?
- » I know I have an issue with a particular host or subnet; what traffic is that host handling?

#### Monitoring network applications in real time

The Applications in Real Time monitor shows the top applications by throughput observed during the last 10 seconds. This report answers questions such as:

- » My link is congested; which applications are on my network right now?
- » How much bandwidth is BitTorrent using right now?

The Applications in Real Time monitor shows inbound application traffic separately from outbound application traffic. Traffic is sorted by transfer rate. The packet rate and a number of flows for each application in that 10 second period are also shown. The Distribution percentage shows the proportion of bandwidth consumption of each application relative to all applications.

You can set the chart to refresh frequently or infrequently or not at all. Each refresh shows 10 seconds of data.

| Inbound Applications   |                      |                   |            |                  |
|------------------------|----------------------|-------------------|------------|------------------|
| Application Name       | Transfer Rate (Mbps) | Packet Rate (pps) | Flows      | Distribution (%) |
| <b>Total</b>           | <b>89.947</b>        | <b>9346</b>       | <b>268</b> |                  |
| HTTP                   | 45.729               | 4003              | 58         |                  |
| FTP                    | 39.354               | 3252              | 4          |                  |
| MySQL                  | 4.569                | 2044              | 1          |                  |
| Google Shared Services | 0.290                | 41                | 10         |                  |
| MSRPC                  | 0.002                | 3                 | 8          |                  |
| ICMPV6                 | 0.001                | 1                 | 1          |                  |
| NetBIOS                | 0.001                | 1                 | 4          |                  |
| HTTPS                  | 0.001                | 1                 | 8          |                  |
| CIFS                   | 0.000                | 0                 | 1          |                  |

Screenshot 86: The Inbound Applications monitor

| Outbound Applications |                      |                   |            |                  |
|-----------------------|----------------------|-------------------|------------|------------------|
| Application Name      | Transfer Rate (kbps) | Packet Rate (pps) | Flows      | Distribution (%) |
| <b>Total</b>          | <b>464.078</b>       | <b>378</b>        | <b>149</b> |                  |
| Unclassified          | 273.666              | 75                | 13         |                  |
| FTP                   | 95.968               | 171               | 4          |                  |
| HTTP                  | 52.698               | 94                | 5          |                  |
| MSRPC                 | 15.014               | 7                 | 12         |                  |
| DNS                   | 14.562               | 21                | 82         |                  |
| NetBIOS               | 7.850                | 8                 | 17         |                  |
| RIP                   | 2.160                | 1                 | 1          |                  |
| ICMP                  | 0.816                | 1                 | 4          |                  |
| DHCP                  | 0.554                | 0                 | 1          |                  |
| SNMP                  | 0.397                | 0                 | 4          |                  |
| NTP                   | 0.301                | 0                 | 1          |                  |
| CIFS                  | 0.094                | 0                 | 4          |                  |

Screenshot 87: The Outbound Applications monitor

To find this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Real Time > Applications**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

Related topic

[Control Anonymous Proxy Traffic](#)

## Monitoring hosts and users in real time

The Hosts/Users widgets in the Realtime Monitor shows the top internal hosts by bandwidth consumption observed during the last 10 seconds. The data displayed answers questions such as:

» My link is congested. Which hosts are on my network right now?

The Realtime Monitor separates inbound and outbound host/user traffic. The traffic is sorted by transfer rate. The packet rate and a number of flows in the preceding 10 second period are shown. The user name of the internal hosts will also be displayed if configured.

The Distribution percentage shows the proportion of bandwidth consumption of each host relative to all hosts for the period. You can set the chart to refresh frequently or infrequently or not at all. Each refresh shows 10 seconds of data.

| Outbound Hosts/Users                 |                      |                   |             |                  |
|--------------------------------------|----------------------|-------------------|-------------|------------------|
| IP Address                           | Transfer Rate (kbps) | Packet Rate (pps) | Flows       | Distribution (%) |
| <b>Total</b>                         | <b>1894.540</b>      | <b>2094</b>       | <b>1344</b> |                  |
| 192.168.17.99                        | 540.700              | 965               | 5           |                  |
| 192.168.10.27                        | 353.330              | 380               | 53          |                  |
| 192.168.10.2                         | 299.813              | 122               | 54          |                  |
| 192.168.10.19                        | 227.189              | 303               | 837         |                  |
| 192.168.10.90                        | 114.034              | 84                | 13          |                  |
| 192.168.10.9                         | 72.630               | 79                | 78          |                  |
| 192.168.10.3                         | 46.414               | 28                | 6           |                  |
| 2001:470:1f0b:992:217:a4ff:fe4a:81e7 | 29.577               | 5                 | 6           |                  |
| 2607:f298:1:106::746:b2c             | 23.178               | 3                 | 3           |                  |
| 192.168.10.27                        | 22.739               | 8                 | 4           |                  |
| 2620:101:8001:5::2:6                 | 19.135               | 3                 | 2           |                  |

Screenshot 88: Monitoring inbound hosts/users report

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Real Time > Hosts/Users**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To show the user associated with the internal hosts, check the **Show Users** checkbox.

### NOTE

Active Directory must be configured on the Exinda Appliances before usernames can be displayed in reports. See [Integrate with Active Directory](#) (page 553).

## Monitoring conversations in real time

The Realtime Conversations monitor shows the top conversations by throughput observed by the Exinda Appliance during the last 10 seconds. This report answers questions such as:

- » My link is congested; who's doing what on my network right now?
- » I think I have a problem with a particular host or subnet; what is that host or subnet doing right now?
- » Is network traffic being accelerated or processed by Edge Cache properly?
- » Is network traffic passing through my High Availability or Cluster correctly?"

Inbound and outbound conversation traffic is displayed separately. Conversations are represented by external IP address, internal IP address, and application. Some traffic types show extra information (like URL) in square brackets following the application.

Traffic is sorted by transfer rate. The packet rate and number of flows for each conversation in the preceding 10 second period is shown. You can set the chart to refresh frequently or infrequently or not at all. Each refresh shows 10 seconds of data.

The Realtime Conversations monitor helps you diagnose issues by:

- » Filtering the conversations by IP address or subnet
- » Showing the user name associated with the internal IP address
- » Showing which policy the conversation falls into
- » Allowing connections within a flow to either be shown individually or to be grouped together
- » Highlighting accelerated conversations in yellow and indicating the acceleration technique used
- » Highlighting conversations processed by Edge Cache (in blue)
- » Indicating how the conversations is flowing through the high availability cluster
- » Indicating asymmetric traffic

| Inbound Conversations                          |                 |                    |                      |                   |           |
|------------------------------------------------|-----------------|--------------------|----------------------|-------------------|-----------|
| External IP                                    | Internal IP     | Application        | Transfer Rate (Mbps) | Packet Rate (pps) | Flows     |
| <b>Total</b>                                   |                 |                    | <b>23.866</b>        | <b>4393</b>       | <b>32</b> |
| <b>Internet: Web - Guarantee High 10%-100%</b> |                 |                    |                      |                   |           |
| 192.168.13.1                                   | 192.168.13.99   | HTTP[192.168.13.1] | 8.936                | 762               | 1         |
| 204.200.202.152                                | 172.14.1.169    | HTTPS              | 0.002                | 1                 | 1         |
| 172.14.1.169                                   | 207.5.72.97     | HTTPS              | 0.001                | 2                 | 4         |
| 172.14.1.169                                   | 204.200.202.152 | HTTPS              | 0.001                | 1                 | 1         |
| <b>Internet: Files - Guarantee Med 8%-100%</b> |                 |                    |                      |                   |           |
| 192.168.13.1                                   | 192.168.13.99   | FTP                | 8.453                | 698               | 2         |
| 172.14.1.169                                   | 172.14.1.36     | CIFS               | 0.002                | 2                 | 1         |
| 172.14.1.169                                   | 172.14.1.59     | CIFS               | 0.001                | 1                 | 1         |
| 172.14.1.169                                   | 172.14.1.128    | CIFS               | 0.000                | 0                 | 1         |
| 172.14.1.169                                   | 172.14.1.190    | CIFS               | 0.000                | 0                 | 1         |
| 172.14.1.169                                   | 172.14.1.70     | CIFS               | 0.000                | 0                 | 1         |
| <b>Internet: ALL - Guarantee Low 5%-100%</b>   |                 |                    |                      |                   |           |
| 172.14.1.169                                   | 172.14.1.145    | MySQL              | 6.465                | 2918              | 1         |
| 172.14.1.169                                   | 172.14.1.128    | MSRPC              | 0.004                | 6                 | 14        |
| 172.14.1.169                                   | 172.14.1.255    | NetBIOS            | 0.001                | 1                 | 1         |
| 172.14.1.169                                   | 172.14.1.36     | NetBIOS            | 0.000                | 0                 | 1         |
| 172.14.1.169                                   | 172.14.1.36     | ICMP[Echo Reply]   | 0.000                | 0                 | 1         |

Screenshot 89: The Conversation monitor report displays information about network traffic.

Where do I find this report?

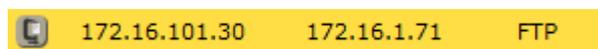
1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Real Time > Conversations**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

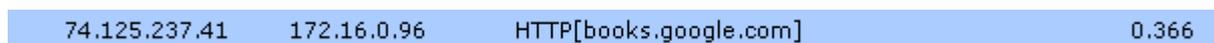
To understand the acceleration and high availability icons & coloring

Accelerated conversations are highlighted in yellow and the application acceleration technologies applied to that conversation are displayed on the left-hand side as a series of icons.

For example, the FTP connection below is accelerated and processed by WAN Memory, indicated by the icon.



Conversations processed by Edge Cache (see the *Edge Cache "How-to" Guide*), are highlighted in blue.



The following legend describes the meaning of each acceleration icon.

|  |                                                                           |
|--|---------------------------------------------------------------------------|
|  | WAN Memory: The connection is been processed by WAN Memory.               |
|  | CIFS Acceleration: The connection is been processed by CIFS Acceleration. |
|  | SSL Acceleration: The connection is been processed by SSL Acceleration.   |
|  | NCP Acceleration: The connection is been processed by NCP Acceleration.   |
|  | MAPI Acceleration: The connection is been processed by MAPI Acceleration. |

When an appliance is deployed in a High Availability (HA) or Clustering mode, the following icons may appear next to conversations.

|  |                                                                                                                                                               |
|--|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Asymmetric: The traffic is asymmetric, and is not being accelerated.                                                                                          |
|  | Locally Bridged: The connection is passing through this appliance in the cluster.<br>If the row is colored yellow, then the flow is remotely accelerated.     |
|  | Remotely Bridged: The connection is passing through another appliance in the cluster.<br>If the row is colored yellow, then the flow is remotely accelerated. |



Locally/Remotely Bridged: The connection is passing through both this and other appliances in the cluster. (Asymmetric traffic)  
If the row is colored yellow, then the flow is remotely accelerated.



Locally Bridged + Locally Accelerated: The connection is passing through this appliance in the cluster and is being accelerated on this appliance.



Remotely Bridged + Locally Accelerated: The connection is passing through another appliance in the cluster and is being accelerated on this appliance.



Locally/Remotely Bridged + Locally Accelerated: The connection is passing through both this and other appliances in the cluster and is being accelerated on this appliance. (Asymmetric traffic)

To filter by IP address or subnet

Type the IP address or subnet in the **IP/Subnet Filter** field at the top of the screen and click **Apply**. Conversations can be filtered by IPv4 or IPv6 addresses.

To show the user associated with the internal hosts

Check the **Show Users** checkbox.

#### NOTE

You must have Active Directory configured to display user names. See For more information, refer to [Integrate with Active Directory](#) (page 553)..

To group individual connections within a flow as a single line item or to show each connection as a separate line item

Toggle on or off the **Group** checkbox at the top of the screen.

To display the policy the conversation falls into

Select **Show Policies** checkbox at the top of the screen.

Troubleshooting virtual circuits and policies

If you are unsure whether traffic is being processed properly by the virtual circuit or the policies within your virtual circuit, use real time monitoring to determine if traffic is hitting the virtual circuit as you expect.

1. On the **Real Time Conversations** screen.
2. Select **Show Policies** checkbox at the top of the screen to group the conversations by virtual circuit and policy.
3. Look at the traffic falling under your virtual circuit to decide if the correct traffic is in there.
4. If traffic that you expected is not there, look in other virtual circuit groups to find the traffic.
5. If traffic is falling into the wrong virtual circuit, check the definition of the affected virtual circuits and ensure the most specific virtual circuit is higher in the policy tree.

### Monitoring edge cache traffic in real time

The Realtime Conversations monitor shows the top conversations by throughput observed by the Exinda Appliance during the last 10 seconds. This report answers questions such as:

- » Is traffic being processed by Edge Cache properly?

For more information, refer to [Monitoring conversations in real time](#) (page 246).

| Inbound Conversations                          |                 |                    |                      |                   |           |
|------------------------------------------------|-----------------|--------------------|----------------------|-------------------|-----------|
| External IP                                    | Internal IP     | Application        | Transfer Rate (Mbps) | Packet Rate (pps) | Flows     |
| <b>Total</b>                                   |                 |                    | <b>23.866</b>        | <b>4393</b>       | <b>32</b> |
| <b>Internet: Web - Guarantee High 10%-100%</b> |                 |                    |                      |                   |           |
| 192.168.13.1                                   | 192.168.13.99   | HTTP[192.168.13.1] | 8.936                | 762               | 1         |
| 204.200.202.152                                | 172.14.1.169    | HTTPS              | 0.002                | 1                 | 1         |
| 172.14.1.169                                   | 207.5.72.97     | HTTPS              | 0.001                | 2                 | 4         |
| 172.14.1.169                                   | 204.200.202.152 | HTTPS              | 0.001                | 1                 | 1         |
| <b>Internet: Files - Guarantee Med 8%-100%</b> |                 |                    |                      |                   |           |
| 192.168.13.1                                   | 192.168.13.99   | FTP                | 8.453                | 698               | 2         |
| 172.14.1.169                                   | 172.14.1.36     | CIFS               | 0.002                | 2                 | 1         |
| 172.14.1.169                                   | 172.14.1.59     | CIFS               | 0.001                | 1                 | 1         |
| 172.14.1.169                                   | 172.14.1.128    | CIFS               | 0.000                | 0                 | 1         |
| 172.14.1.169                                   | 172.14.1.190    | CIFS               | 0.000                | 0                 | 1         |
| 172.14.1.169                                   | 172.14.1.70     | CIFS               | 0.000                | 0                 | 1         |
| <b>Internet: ALL - Guarantee Low 5%-100%</b>   |                 |                    |                      |                   |           |
| 172.14.1.169                                   | 172.14.1.145    | MySQL              | 6.465                | 2918              | 1         |
| 172.14.1.169                                   | 172.14.1.128    | MSRPC              | 0.004                | 6                 | 14        |
| 172.14.1.169                                   | 172.14.1.255    | NetBIOS            | 0.001                | 1                 | 1         |
| 172.14.1.169                                   | 172.14.1.36     | NetBIOS            | 0.000                | 0                 | 1         |
| 172.14.1.169                                   | 172.14.1.36     | ICMP[Echo Reply]   | 0.000                | 0                 | 1         |

Screenshot 90: The Conversation monitor report displays information about network traffic.

When a conversation has been processed by Edge Cache (see the [Edge Cache "How to" guide](#)), it is highlighted in blue.

|               |             |                        |       |
|---------------|-------------|------------------------|-------|
| 74.125.237.41 | 172.16.0.96 | HTTP[books.google.com] | 0.366 |
|---------------|-------------|------------------------|-------|

#### NOTE

All conversations evaluated by Edge Cache will be highlighted in blue even if the object is set for exclusion from Edge Cache storage.

To learn more about this monitor, see [Conversations in Real Time](#).

### Monitoring reductions in real time

The Realtime Reductions monitor shows reduction by application during the last 5 minutes. This report answers questions such as:

- » Am I getting the reduction, freeing the network capacity I expect?

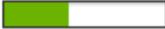
The monitor filter shows inbound application traffic, outbound application traffic, or bi-directional traffic.

The amount of data on the LAN-side and WAN-side of the appliance is shown along with the reduction ratio during the previous 5 minute sliding window.

You can set the chart to refresh frequently, infrequently, or not at all. Each refresh shows 5 minutes of data ending on the most recent 10 second boundary. That is, if it refreshes at 9:38:17, the chart will show data from 9:33:10 - 9:38:10.

## NOTE

This report is only shown when there are two or more Exinda Appliances.

| Reduction Statistics (Last 5 Minutes) |                 |               |               |                                                                                          |
|---------------------------------------|-----------------|---------------|---------------|------------------------------------------------------------------------------------------|
|                                       | Application     | LAN Data (MB) | WAN Data (MB) | Reduction Ratio (%)                                                                      |
| 1                                     | URL             | 0.0008        | 0.0007        |  12.50 |
| 2                                     | EMC Replication | 255.71        | 211.34        |  17.35 |
| 3                                     | HTTP            | 0.82          | 0.48          |  41.46 |
| 4                                     | Oracle          | 0.0010        | 0.0009        |  10.00 |
| 5                                     | LotusNotes      | 0.05          | 0.03          |  40.00 |

Screenshot 91: The reduction statistics report.

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Real Time > Reduction**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

The reduction ratio calculation:

$$\text{Reduction Ratio} = (\text{Data Transfer Size Before Exinda} - \text{Data Transfer Size After Exinda}) / \text{Data Transfer Size Before Exinda}$$
$$\text{Reduction Ratio} = \text{LAN Data} - \text{WAN Data} / \text{LAN Data}$$

## Monitoring application response in real time

The Realtime Application Response monitor shows the slowest applications by round-trip-time observed by the Exinda Appliance during the last 10 seconds.

This report can answer questions such as:

- » Which applications may be having problems?
- » What are my poorest performing applications?
- » Why is the application performing poorly; could it be due to network delay or server delay?

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Real Time > Application Response**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

The monitor shows application response metrics like round-trip time (RTT), normalized network delay, normalized server delay, normalized total delay, network delay, server delay, transaction delay, transaction count, and flow count by application. Traffic is sorted by round-trip-time.

You can set the chart to refresh frequently or infrequently or not at all. Each refresh shows 10 seconds of data.

| Application Response |          |                            |                           |                                |              |             |                        |                   |       |
|----------------------|----------|----------------------------|---------------------------|--------------------------------|--------------|-------------|------------------------|-------------------|-------|
| Application Name     | RTT (ms) | Normalized Network (ms/kb) | Normalized Server (ms/kb) | Normalized Delay Total (ms/kb) | Network (ms) | Server (ms) | Transaction Delay (ms) | Transaction Count | Flows |
| FTP                  | 1.94     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 1     |
| HTTP                 | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 2     |
| MSRPC                | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 12    |
| NTP                  | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 1     |
| DHCP                 | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 1     |
| DNS                  | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 56    |
| SNMP                 | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 4     |
| RIP                  | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 1     |
| ICMP                 | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 4     |
| NetBIOS              | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 16    |

Screenshot 92: The Application Response monitor displays response by RTT.

#### NOTE

These statistics are only available if the Performance Metrics ASAM Module is enabled on the [System > Setup > Monitoring](#) page.

Related topic

[Monitor the real time application response](#)

[Configuring a policy to accelerate traffic](#)

[Configuring adaptive response limits](#)

[Monitoring real time application response](#)

The APM values are available as a real time display. The real time display shows the APM values by application for the selected time period. As well as the APM values, the number of flows and the number of transactions are shown.

Display the report in the Exinda Web UI

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Monitor > Real Time** and switch to the **Application Response** tab.

The following report opens:

| Application Response |          |                            |                           |                                |              |             |                        |                   |       |
|----------------------|----------|----------------------------|---------------------------|--------------------------------|--------------|-------------|------------------------|-------------------|-------|
| Application Name     | RTT (ms) | Normalized Network (ms/kb) | Normalized Server (ms/kb) | Normalized Delay Total (ms/kb) | Network (ms) | Server (ms) | Transaction Delay (ms) | Transaction Count | Flows |
| HTTP                 | 3074.50  | 1.94                       | 0.98                      | 2.92                           | 73.19        | 2.16        | 75.36                  | 38                | 79    |
| FTP                  | 9.81     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 6     |
| mDNS                 | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 4     |
| ICMPV6               | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 1     |
| HTTPS                | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 4     |
| DNS                  | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 202   |
| SMTP                 | 0.00     | 0.00                       | 0.00                      | 0.00                           | 0.00         | 0.00        | 0.00                   | 0                 | 2     |

6. To change how often the table is refreshed, select an **Auto-Refresh Rate** from the list.

Display the report in the Exinda CLI

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Configuration > System > Tools > Console**.
5. Type the appliance username and password at the prompts. Do one of the following:
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.
7. To display real time APM data from the CLI, use the following command:  
`(config) # show realtime apm applications`

The following results are displayed:

```
ex-240 (config) # show realtime apm applications
```

| Application  | RTT (ms) | Network (ms) | Server (ms) | Transaction (ms) | Transactions | Flows |
|--------------|----------|--------------|-------------|------------------|--------------|-------|
| ExindaWM     | 956.04   | 77706.24     | 206863.37   | 226125.26        | 48           | 4     |
| Unclassified | 459.74   | 35040.99     | 15000.30    | 37512.24         | 8            | 44    |
| Replify      | 292.75   | 2660.00      | 0.00        | 2655.70          | 4            | 1     |
| HTTP         | 256.16   | 202.86       | 147.08      | 338.41           | 10           | 9     |
| HTTPS        | 217.45   | 97.34        | 26.83       | 124.18           | 10           | 6     |
| CIFS         | 108.53   | 186.69       | 89.73       | 231.30           | 2            | 2     |
| SSH          | 71.48    | 386.28       | 0.00        | 336.24           | 2            | 1     |
| ExindaCom    | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 16    |
| mDNS         | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 3     |
| ICMP         | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 7     |
| ssdp         | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 1     |
| IGMP         | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 15    |
| NTP          | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 2     |
| slm          | 0.00     | 0.00         | 0.00        | 0.00             | 0            | 1     |

```
ex-240 (config) # █
```

## Monitoring host health in real time

The Realtime Host Health monitor shows unhealthy hosts as measured by the number of retransmitted bytes during the last 10 seconds. This report answers questions such as:

- » Which internal hosts are having the most difficulty with successfully transmitting traffic?

The monitor separates internal and external hosts and displays metrics like number of retransmitted bytes, number of aborted connections, number of refused connections, number of ignored connections, and the flow count for each internal and external host monitored during the last 10 seconds.

## NOTE

A new internal mechanism was implemented in the ExOS 7.4.2 firmware that prevents Exinda Appliances from being affected by DDoS attacks. As a consequence, the "Ignored Connections" historical report is no longer available, but the Realtime Health report includes the number of ignored connections.

Traffic is sorted by the amount of retransmitted bytes. You can set the chart to refresh frequently or infrequently or not at all. Each refresh shows 10 seconds of data.

| Health         |                       |         |         |         |       |  |
|----------------|-----------------------|---------|---------|---------|-------|--|
| Internal IP    | Retransmitted (bytes) | Aborted | Refused | Ignored | Flows |  |
| 192.168.10.1   | 275132                | 0       | 0       | 0       | 35    |  |
| 192.168.10.28  | 75817                 | 0       | 2       | 0       | 13    |  |
| 192.168.10.6   | 5690                  | 0       | 0       | 0       | 14    |  |
| 192.168.10.5   | 5606                  | 0       | 0       | 0       | 78    |  |
| 192.168.10.62  | 0                     | 0       | 0       | 0       | 1     |  |
| 192.168.16.249 | 0                     | 0       | 0       | 0       | 1     |  |
| 192.168.10.24  | 0                     | 0       | 0       | 0       | 1     |  |

Screenshot 93: The Realtime Host Health report displays the number of retransmitted bytes.

## NOTE

These statistics are only available if the Performance Metrics ASAM Module is enabled on the [System > Setup > Monitoring](#) page.

Display the report in the Exinda Web UI

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Monitor > Real Time>Host Health**.

The reports contains the following status:

| Connection Status          | Description                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Aborted Connections</b> | Connections that were unexpectedly aborted by either the client or server sending a TCP reset.              |
| <b>Refused Connections</b> | Connections that were refused by the server (TCP SYN sent, received ICMP refused or TCP reset in response). |
| <b>Ignored Connections</b> | Connections that were ignored by the server (TCP SYN sent, received nothing in response).                   |

5. To change how often the table is refreshed, select an **Auto-Refresh Rate** from the list.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

Display the report in the Exinda CLI

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).

2. Key-in the **User** and **Password**.

3. Click **Login**.

4. Click **Configuration > System > Tools > Console**.

5. Type the appliance username and password at the prompts. Do one of the following:

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`

The `hostname #` prompt appears.

- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`

The `hostname (config) #` prompt appears.

6. To display realtime TCP health from the CLI, use the following command:

```
(config) # show realtime apm hosts
```

The following results are displayed:

```
ex-240 (config) # show realtime apm hosts
```

| Internal Host  | Retransmissions | Aborted | Refused | Ignored | Flows |
|----------------|-----------------|---------|---------|---------|-------|
| 172.16.1.240   | 0               | 0       | 0       | 0       | 13    |
| 192.168.0.176  | 0               | 0       | 0       | 0       | 1     |
| 172.16.0.213   | 0               | 0       | 0       | 0       | 1     |
| 192.168.50.147 | 0               | 0       | 0       | 0       | 1     |
| 192.168.0.179  | 0               | 0       | 0       | 0       | 1     |
| 172.16.0.63    | 0               | 0       | 2       | 0       | 3     |
| 172.16.1.242   | 0               | 0       | 0       | 0       | 1     |
| 192.168.40.96  | 0               | 0       | 0       | 0       | 1     |
| 192.168.0.178  | 0               | 0       | 0       | 0       | 6     |
| 0.0.0.0        | 0               | 0       | 0       | 0       | 1     |
| 192.168.0.209  | 0               | 0       | 0       | 0       | 1     |
| 192.168.50.143 | 0               | 0       | 0       | 0       | 1     |
| 172.16.0.252   | 0               | 0       | 0       | 0       | 1     |
| 172.16.0.108   | 0               | 0       | 0       | 0       | 3     |
| 172.16.1.149   | 0               | 0       | 0       | 0       | 3     |
| 172.16.0.67    | 0               | 0       | 0       | 0       | 5     |
| 172.16.0.190   | 0               | 1       | 0       | 0       | 4     |
| 192.168.0.118  | 0               | 0       | 0       | 0       | 1     |
| 192.168.0.145  | 0               | 0       | 0       | 0       | 1     |
| 192.168.0.207  | 0               | 0       | 0       | 0       | 1     |

Screenshot 94: Realtime TCP health from the CLI

[Related topic](#)

[Monitor the real time application response](#)

### 0.11.3 Monitoring network interfaces

Interface reports allow you to view the volume of traffic flowing in and out of your network. The Throughput report displays interface and bridge throughput. The Packets Per Second report displays the outbound packet rate from your network. These reports provide answers to important questions about your network traffic.

#### Monitoring interface throughput

The Interfaces Throughput report shows throughput for a defined period of time for a particular interface or for all interfaces aggregated. This report answers questions such as:

- » Is my link congested?
- » Which bridge is congested?
- » Was there a brief maximum burst or is the traffic consistently over average?"

The throughput for inbound traffic is shown separately from outbound traffic. The charts show maximum throughput and average throughput.

Maximum throughput is determined by looking at 10 second samples of data. Average throughput is calculated by averaging the amount of data transferred over the time period specified by the granularity.

For instance, when looking at a day of traffic, each data point represents 5 minutes of data. The average shows the average data transferred over each 5 minute period. You can better understand the throughput distribution by adding a percentile marker line. The line shows that the  $n^{\text{th}}$  percentile of (maximum) throughput observations exceeds the specified throughput rate. For instance, if the 95<sup>th</sup> percentile is at 55 Mbps, then 5 out of 100 maximum throughput observations appeared above 55 Mbps.

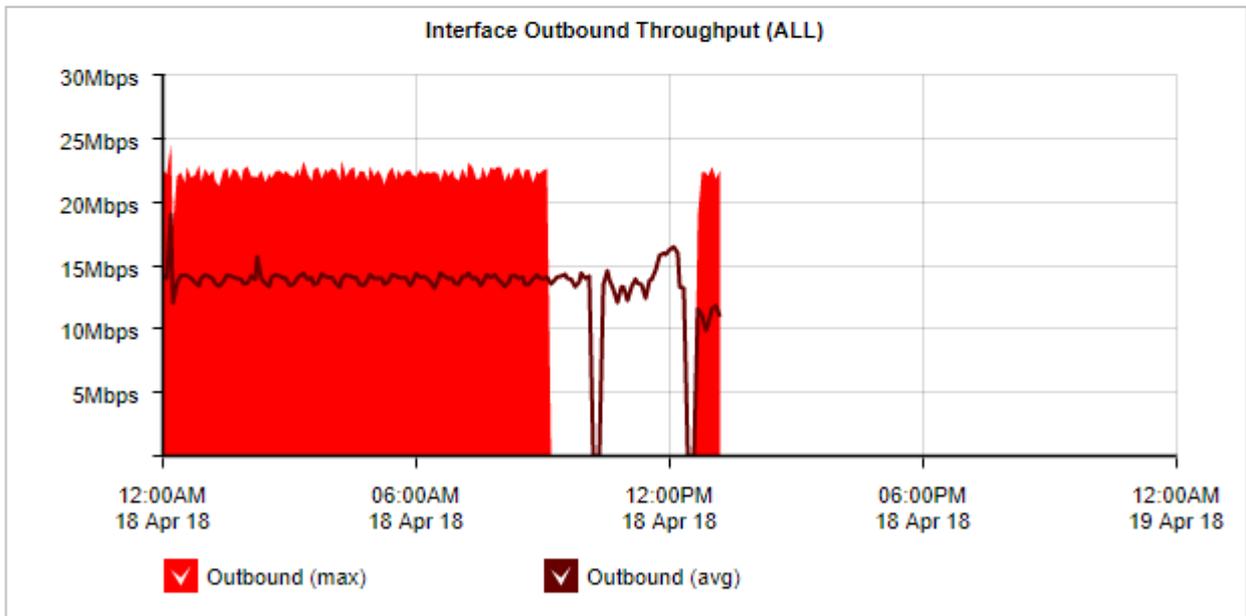
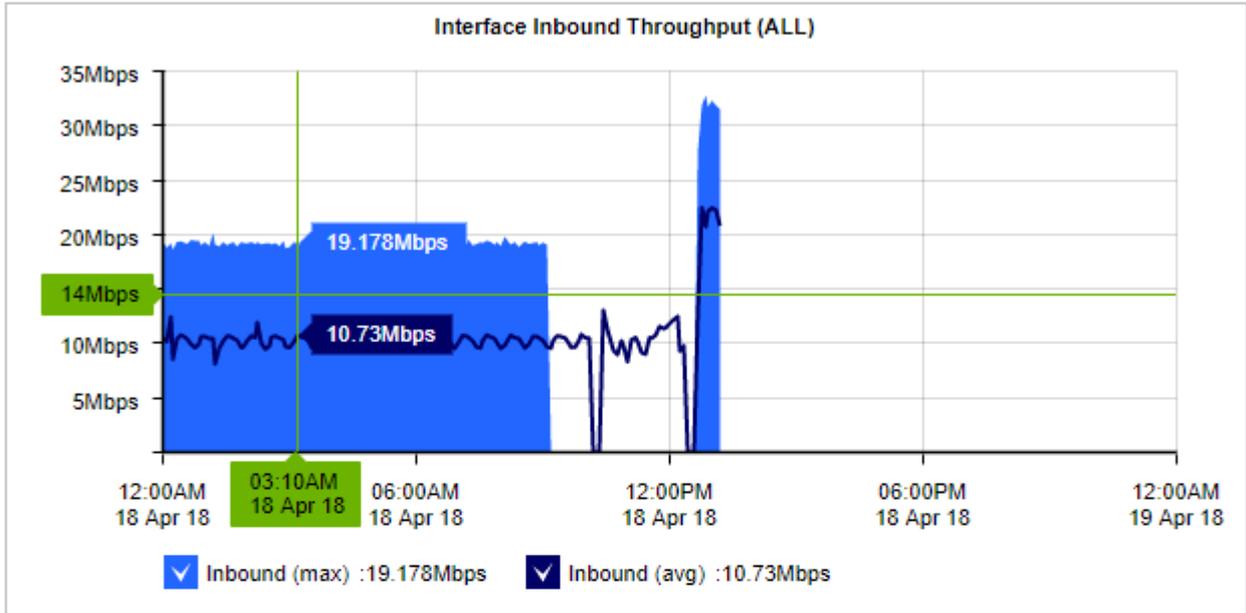
The table at the bottom of the report shows the total amount of data transferred into and out of the WAN-side interface (s), and the maximum and average throughput rates for the selected time period.

#### NOTE

Because this report shows all data on the wire, it may include traffic not seen on the WAN, such as local LAN broadcasts, etc.

WAN/Out-of-path Interface Selection: ALL

Select Percentile Marker to Display: None



| WAN Interface Throughput Summary (ALL) |                 |                       |                       |
|----------------------------------------|-----------------|-----------------------|-----------------------|
| Data Direction                         | Total Data (MB) | Throughput Avg (Mbps) | Throughput Max (Mbps) |
| Inbound                                | 58812.22        | 10.75                 | 32.76                 |
| Outbound                               | 75744.97        | 13.84                 | 24.59                 |

Screenshot 95: The Interface Throughput report displays inbound and outbound network traffic.

Where do I find this report?

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Interfaces > Throughput**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

### To filter the data in the chart by bridge or out-of-path interface

Select the desired bridge (WAN interface), policy-based routing interfaces, or WCCP interface from the **WAN/Out-of-path-Interface Selection** selector. Selecting **ALL** includes all interfaces available.

### To identify how much traffic falls above a specific percentile

Select the desired percentile level from the **Select Percentile Marker to Display** selector.

## Monitoring interface packets per second

The Packets Per Second monitor shows the packet rate over time for a particular interface or for all interfaces aggregated. This report answers questions such as:

- » Is an unacceptably high packet rate possibly indicating a problem such as a denial of service attack?

Only the packet rate for outbound traffic is shown. The charts show maximum and average packet rates.

The maximum packet rate is determined by looking at 10 second samples of data. The average packet rate is calculated by averaging the packet rate over the time period specified by the granularity.

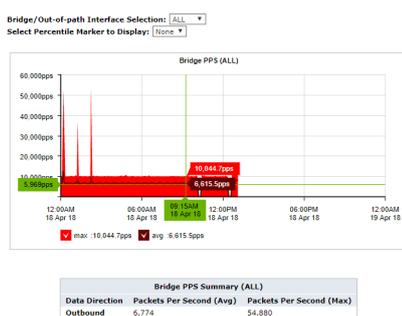
For instance, when looking at a day of traffic, each data point represents 5 minutes of data. The average shows the average packet rate over each 5 minute period. You can better understand the packet rate distribution by adding a percentile marker line.

The line shows that the nth percentile of (maximum) packet rate observations exceeds the specified packets per second rate. For instance, if the 95th percentile is at 40,000 pps, then 5 out of 100 packet rate observations appeared above 40,000 pps.

The table at the bottom of the report shows the average packets per second and maximum packets per second transferred out of the WAN-side interface(s) for the selected time period. The values in the table are automatically updated when the interactive graphs are manipulated.

### NOTE

Because this report shows all data on the wire, it include traffic not seen on the WAN such as local LAN broadcasts, etc.



Screenshot 96: The Interfaces Packet Per Second report displays packet rates per second.

## Where do I find this report?

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Interfaces > Packets Per Second**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

## To filter the data in the chart by bridge or out-of-path interface

Select the desired bridge (WAN interface), policy-based routing interfaces, or WCCP interface from the **WAN/Out-of-path-Interface Selection** selector. Selecting **ALL** includes all interfaces available.

## To identify which packet rate falls above a specific percentile

Select the desired percentile level from the **Select Percentile Marker to Display** selector.

### 0.11.4 Monitoring network throughput

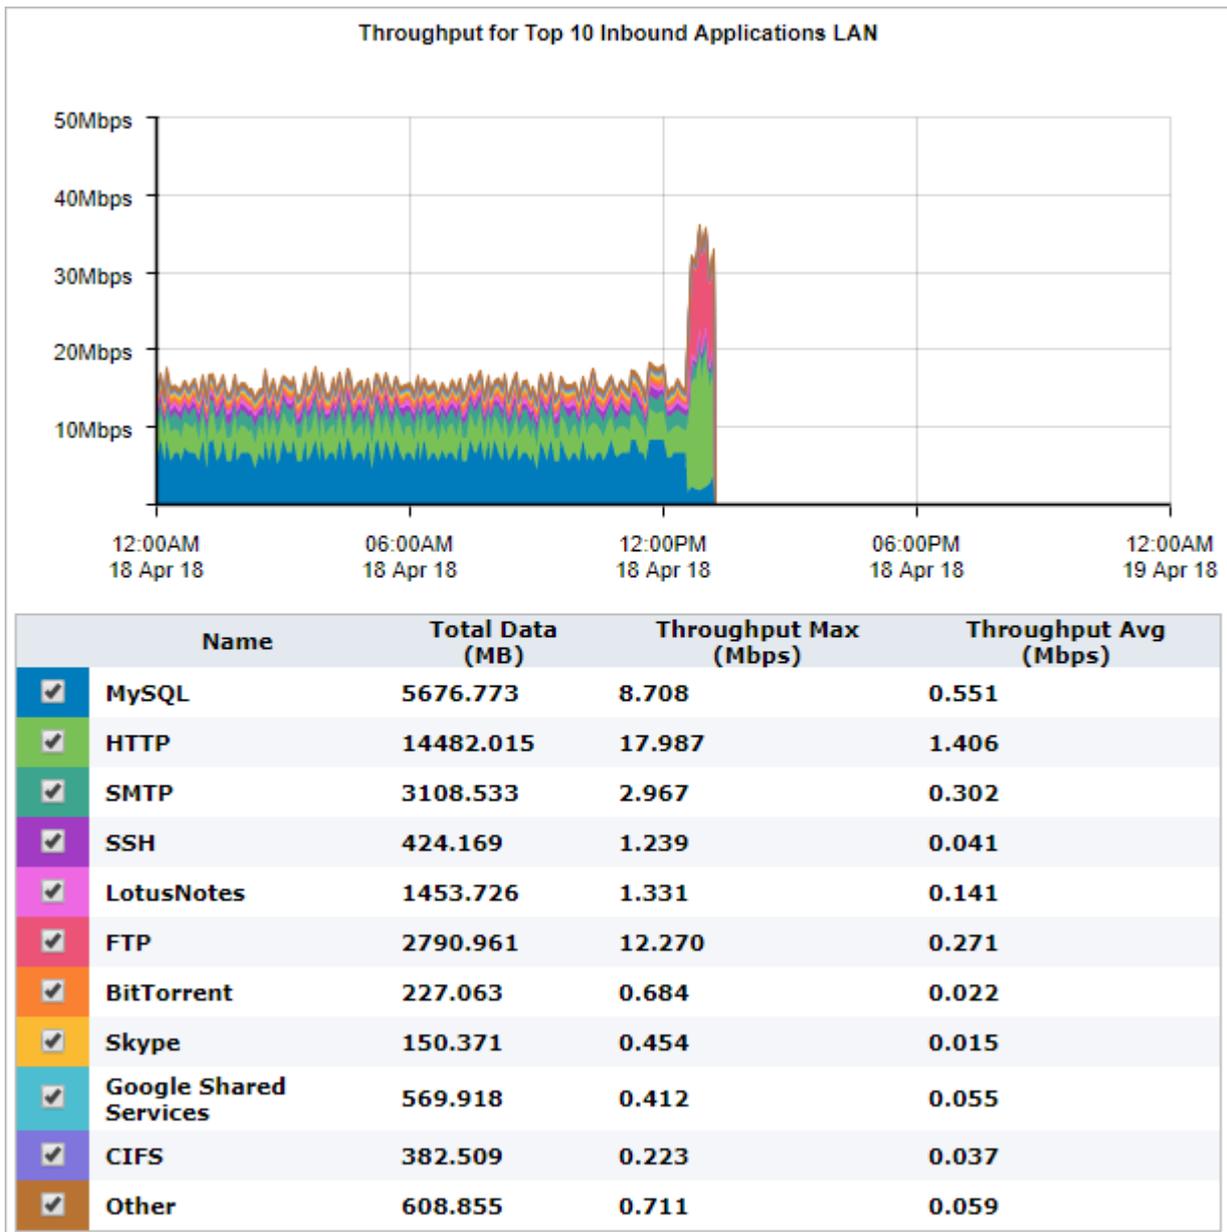
The Network Summary report shows traffic throughput over time by application, application groups, internal or external hosts, internal or external users, conversations, or URLs. You can remove items from the chart to isolate traffic patterns and sources.

This report answers questions such as:

- » What is the pattern of throughput for particular apps, app groups, users, hosts, etc.?
- » Are there any spikes and what type of traffic may be causing the spikes?
- » What would happen to the throughput if I created a policy to block a particular app, application group, user, or host?

The charts help you diagnose issues and perform what-if scenarios to determine the right size of your network.

The report shows LAN-side and WAN-side charts for both inbound traffic and outbound traffic. The total data volume, maximum throughput, and average throughput is also shown in tables below each chart. The charts aggregate data outside the top 10 in a category named "Other".



Screenshot 97: The Network Summary report displays LAN traffic volume for the top 10 inbound applications.

Where do I find this report?

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Network**.

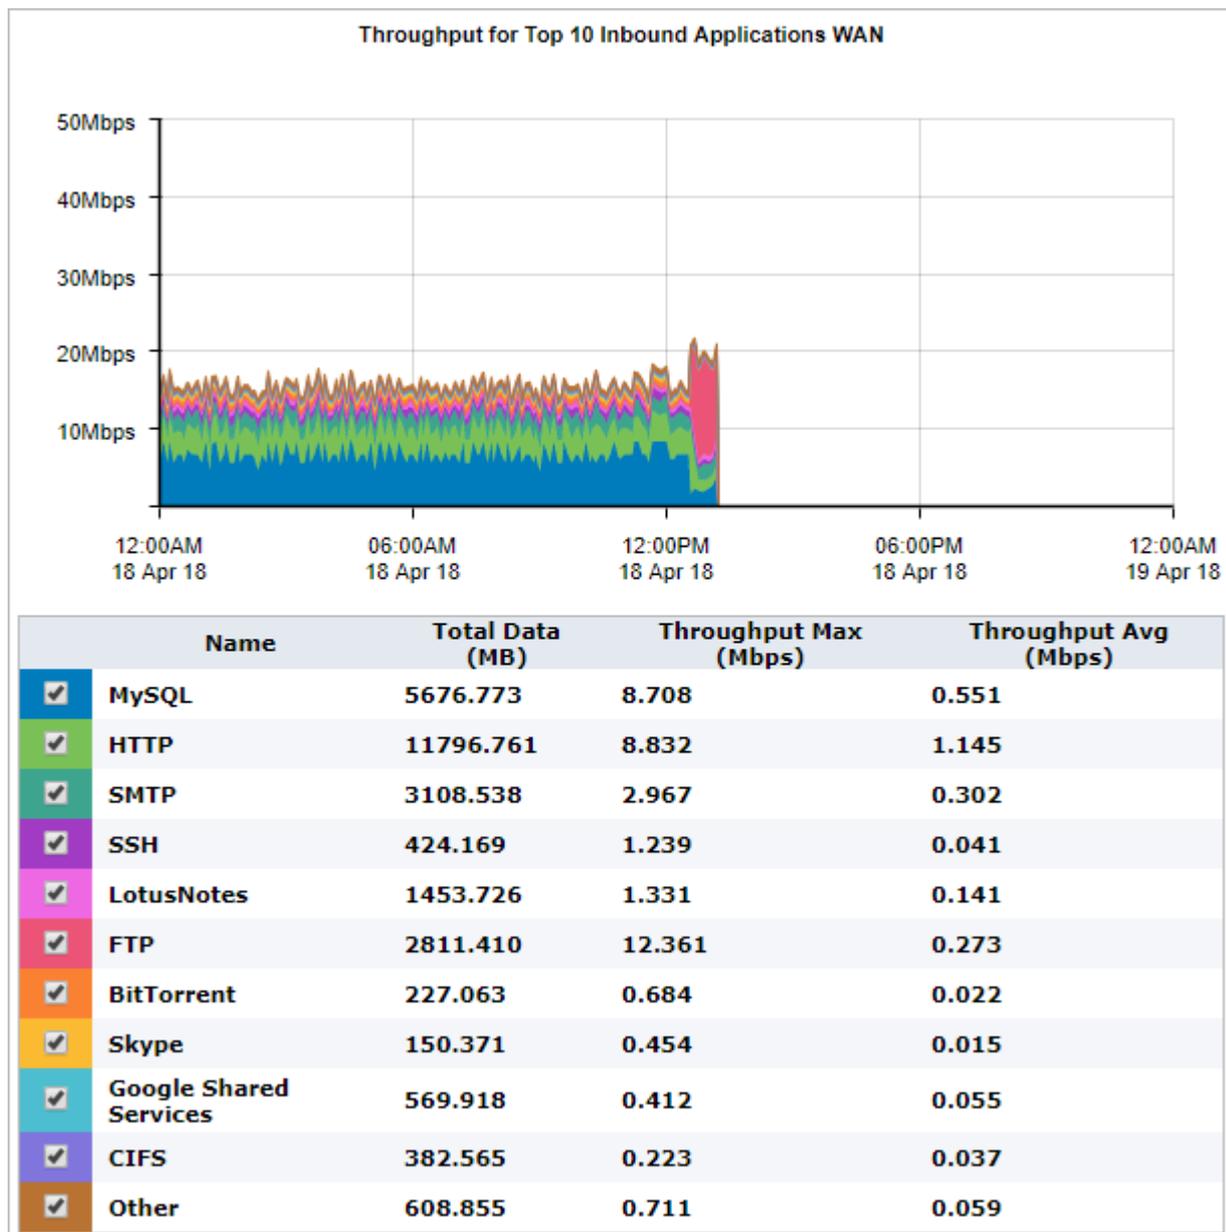
Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To determine the right size of your network (i.e. remove items from the chart)

Remove specific types of traffic from the graph by deselecting their checkbox in the legend below the graph. The remaining traffic models what your network traffic would look like if you blocked that type of traffic. You can then determine an appropriate amount of bandwidth required.

To identify which throughput falls above a specific percentile

Select the desired percentile level from the **Select Percentile Marker to Display** selector.



Screenshot 98: The Network Summary report displays WAN traffic volume for the top 10 inbound applications.

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### 0.11.5 Monitoring service levels

Learn how to view application performance reports, the availability of your ISP, and the health and efficiency of TCP traffic.

## Monitoring application performance scores

The Application Performance Score (APS) report shows scores that assess network performance and user experience when using business-critical applications.

These charts can answer questions such as:

- » Are my important applications performing well from a network perspective for my network users?
- » Has this been a persistent problem or is it getting worse?
- » If an application is not performing well, what might be causing the problem?

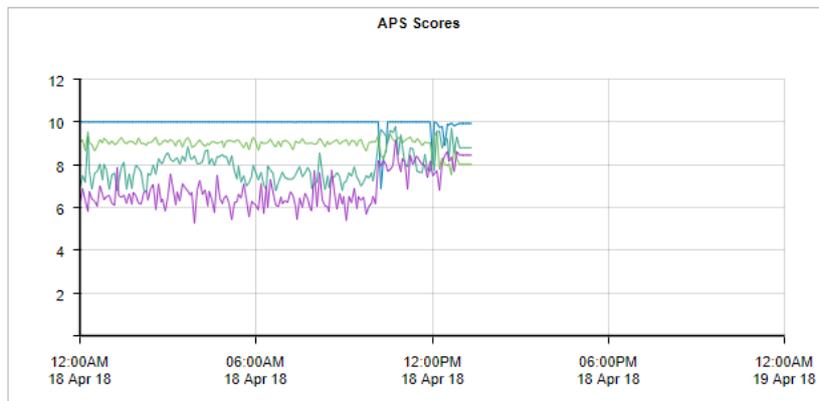
To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Service Levels > Application Performance Score (APS)**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

A score between 0 and 10, where 0 is poor and 10 is excellent, indicates whether the app is performing well or poorly.

Scores are graphed over time to show how the scores changes and trends. The underlying metrics and measures used to calculate scores are shown in the table below the graph. You can drill into the details of the APS by clicking an application name.



| Name                                         | Score | Normalized Delays (ms/kb)                       |        | Transaction Delays (ms) |        | Jitter (ms) | Loss (%) |          | RTT (ms) |
|----------------------------------------------|-------|-------------------------------------------------|--------|-------------------------|--------|-------------|----------|----------|----------|
|                                              |       | Network                                         | Server | Network                 | Server |             | Inbound  | Outbound |          |
|                                              |       | <input checked="" type="checkbox"/> Lotus Notes | 9.93   | 742.72                  | 1.92   |             | 362.37   | 2.55     |          |
| <input checked="" type="checkbox"/> HTTP     | 8.97  | 189.22                                          | 29.14  | 753.09                  | 85.40  | 2928.31     | 8.10     | 15.40    | 1351.23  |
| <input checked="" type="checkbox"/> CIFS     | 7.90  | 3718.46                                         | 42.58  | 1750.19                 | 13.49  | 3538.62     | 23.40    | 33.60    | 2060.20  |
| <input checked="" type="checkbox"/> SMTP     | 6.77  | 92136.15                                        | 7.18   | 20724.26                | 41.22  | 31884.15    | 14.60    | 5.10     | 184.56   |
| <input checked="" type="checkbox"/> Database | -     | -                                               | -      | -                       | -      | -           | -        | -        | -        |

Screenshot 99: The Application Performance Score displays scores from 0 - 10 over time.

A score includes input from one or more of the following metrics:

- » Network delay – the time taken for data to traverse the network (on the wire)
- » Server delay – the time taken for a server to respond to the request

- » Normalized network delay – the time taken for data to traverse the network, where the delay is measured independent of the transaction size by assuming a normalized packet size of 1024 bytes
- » Normalized server delay – the time taken for a server to respond to the request, where the delay is measured independently of the transaction size, by assuming a normalized packet size of 1024 bytes
- » Round-trip time – the time taken for
- » Jitter – the measure of variability of network delay, defined as one standard deviation of network delay
- » Inbound loss – the percentage of packet loss on inbound traffic
- » Outbound loss – the percentage of packet loss on outbound traffic

Each metric that contributes to the score has a threshold value set. The threshold may have been [set manually](#) or may have been determined automatically by the Exinda Appliance observing the traffic for the period of time to determine a baseline threshold values. The table below the chart indicates the current observed values for these metrics and whether that value is considered good or not.

- » If the observed traffic is within the threshold, it is considered good and is colored green in the APS Scores table.
- » If the observed traffic is above the threshold but not above 4 x the threshold, it is considered tolerable and is colored yellow.
- » If the observed traffic is above 4 x the threshold, it is considered poor and is colored red.
- » If there is no color in the table for a particular metric, it indicates that the metric is not contributing to the calculation of the APS score.

Use this information to determine which metrics contribute to an application's performance score.

### Generate a PDF report of APS results

Create a report that contains the APS, TCP health, and TCP efficiency for a specified period of time.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Report** and switch to the **PDF Reports** tab.
6. Click **Add New PDF Report**.
7. In the Report Selection area select **APS, TCP Health**, and **TCP Efficiency**.
8. In the Report Details area, type a name for the report.
9. Specify how often the report will be generated.
10. Click **Add New Report**.
11. To generate the report, locate the report in the list and click **PDF**.

### What to expect

#### If an APS report is not showing data

Either the APS object does not have thresholds set and therefore the score cannot be calculated or there is no traffic for the specified application on the network for the period that is shown on the screen.

#### If the thresholds were set by using the baselining feature

You should get an application performance score of 9.0 if the same traffic were to be observed.

## Evaluating the APS

If the thresholds were set by using the baselining feature, a score of 8.5 or higher is considered a good score. The thresholds are automatically set to be slightly higher than the average observed measures so that good would be at the top end of the range of scores.

If you set the thresholds manually such that your thresholds lie at the boundary of what you consider good or not good, then you should consider APS scores greater than 5.0 good, as statistically half the time you would expect the observed values to be slightly above your threshold and half the time the observed values would be slightly below the threshold.

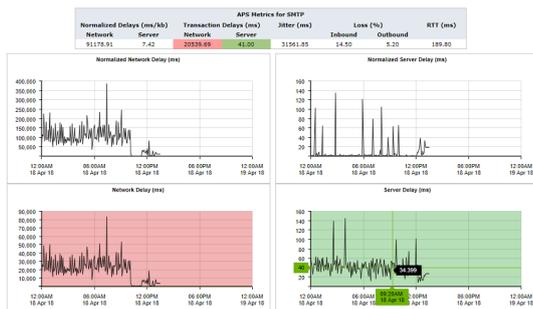
## Looking at the results

### Determining what might be causing a low APS score

Click the APS name in the table. A new screen that shows charts for each underlying metric appears. The background of each chart is colored to represent value within threshold, within 4 x threshold, and above.

If the background of the chart is not colored, that metric does not contribute to the calculation of the APS score. You can zoom into these charts by clicking and dragging within a chart to zoom into the x-axis range. All other charts will synchronize their zoom ranges to the specified zoom range.

The metrics showing poor values could indicate a problem to investigate. For example, if the network delay is good, but the server delay is poor, you know that the network is not to blame and that the server administrator should take a look at the application server.



### Determining if a problem has been persistent

Look at the APS score time line. If the score has been low for an extended period or if it looks like the score is dropping, you'll know that this is a persistent problem that needs addressing.

### Determining if you should pay attention to the normalized delays or the transaction delays

Generally, you should use the transaction delays unless the protocol that is being monitored has large or variable sized packets. The normalized delay measure normalizes the score to reflect a 1024 packet size allowing you to more easily compare delays when the packets are variable in size.

### Configuring the system to notify you if the APS score drops too low

You can configure the system to send an email if the APS score drops below an APS value you specify and remains below that value for a specified duration. For example, you can set it to notify you if it drops below 7.0 and stays below 7.0 for 5 minutes.

### Making the APS chart easier to read by removing score lines

» You can temporarily remove lines from the APS Scores chart by clearing the checkboxes next to the APS name in the table.

» You can zoom into an area of interest by clicking and dragging in the chart to select a smaller time range. This often has the effect of flattening the lines so that it appears less cluttered.

### Related Topics

- » [Configure Application Performance Score Objects](#)
- » [How an Application Performance Score is calculated](#)

### Calculating an application performance score

The Application Performance Score object defines the application traffic that will be monitored and which application performance metrics to evaluate. It also provides application performance thresholds to be used in the evaluation.

For each metric, the observed traffic is compared against the threshold and is classified into one of three categories:

- » Good — The baseline for the application is good, which indicates that the application is performing within the expected levels (below the threshold). Users should be happy with application performance.
- » Tolerated — The performance of the application is less than expected, but still performing within a range that users should be able to tolerate (between the threshold and four times the threshold).
- » Frustrated — The application is performing poorly (more than four times the threshold). Users will be frustrated.

The number of good observations for all metrics with a threshold are totaled and given a full weighting; the number of tolerated observations for all metrics with a threshold are totaled and given a half weighting; and all frustrated observations are given a zero weighting. These weighted totals are summed and divided by the total observations.

$$\text{aps} = 10 * ((1 * \text{number of satisfied samples}) + (0.5 * \text{number of tolerated samples}) + (0 * \text{number of frustrated samples})) / \text{total samples}$$

#### EXAMPLE

For HTTP, a threshold is configured for Network Delay as  $T = 300 \text{ msec}$  and a threshold is configured for round-trip time (RTT) as  $T = 40 \text{ msec}$ .

In one 10s period, 11 flows are sampled for HTTP with the following results:

- » 2 flow samples have a network delay of  $> 1200 \text{ ms}$  (frustrated samples)
- » 3 flow samples have a network delay of  $> 300 \text{ ms}$  but  $< 1200 \text{ ms}$  (tolerated samples)
- » 6 flow samples have a network delay of  $< 300 \text{ ms}$  (satisfied samples)
- » 1 flow sample has a RTT of  $> 40 \text{ ms}$  but  $< 160 \text{ ms}$  (tolerated samples)
- » 10 flow samples have a RTT of  $< 40 \text{ ms}$  (satisfied samples)

The APS score is calculated as follows:

$$\text{aps} = 10 * (1 * (6 + 10) + 0.5 * (3 + 1) + 0 * 2) / 22 = 8.1$$

### Setting thresholds

The appropriate thresholds for an application is unique for each network environment. Thresholds can be set manually when configuring an APS object or the Exinda appliance can analyze the traffic for an application for a baseline period and create a recommended set of thresholds.

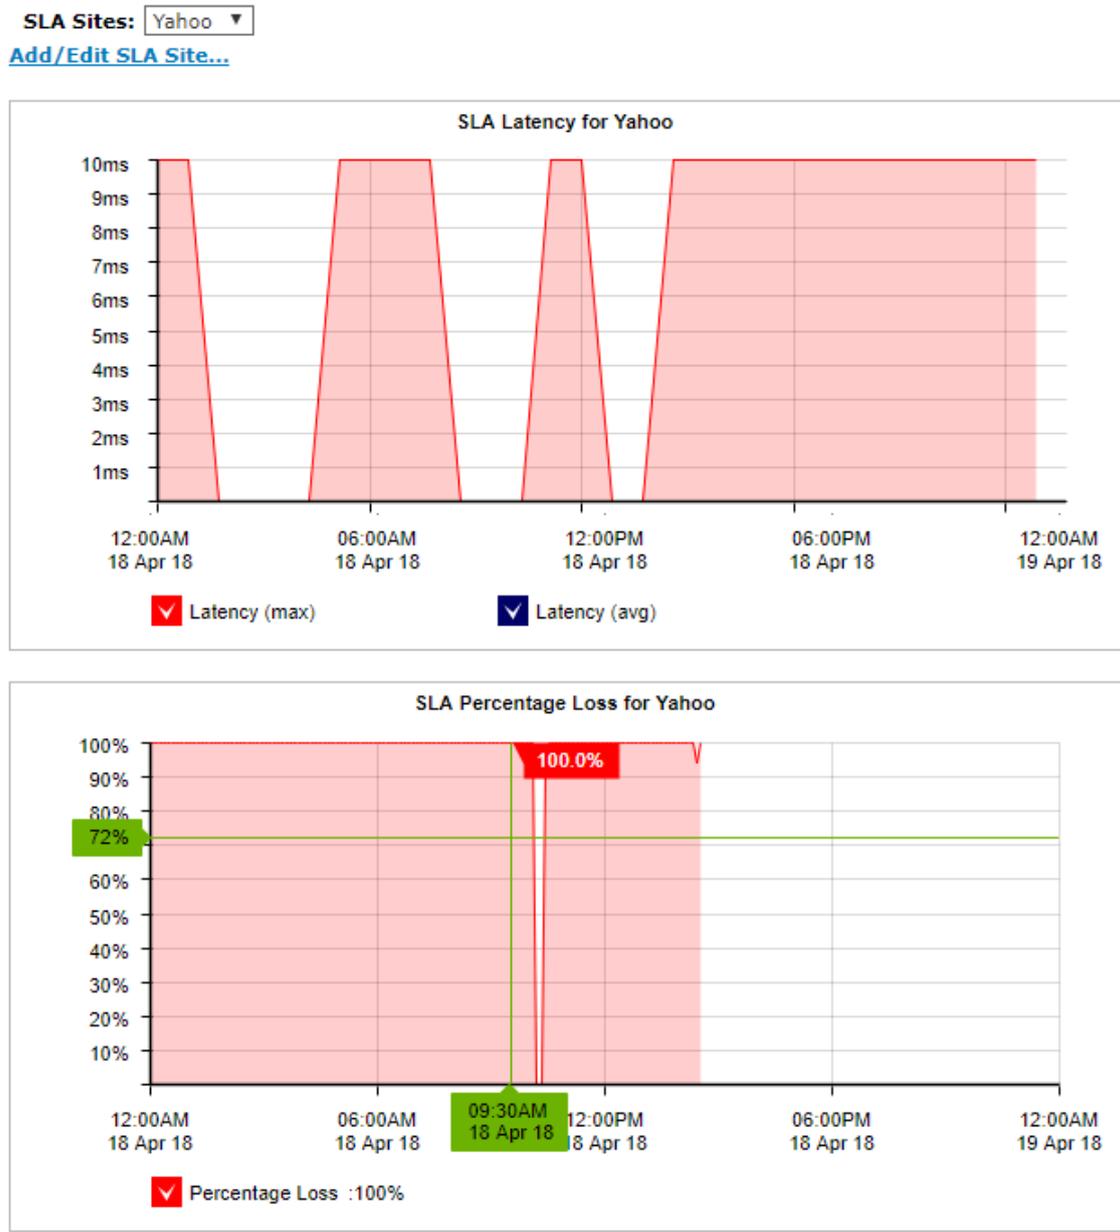
For more information, refer to [Configuring application performance score objects](#) (page 223).

### Monitoring network response SLA

The SLA monitor reports the performance of your ISP against a set of predefined criteria. The SLA monitor sends 1 64-bit long ICMP ping every 10 seconds to the remote site. It reports the maximum and average latency and the percentage

loss of the pings over time. This report answers questions such as:

- » Is my ISP always available?
- » What is the latency of my ISP?



| SLA Statistics for Yahoo |            |             |              |                  |                  |                  |
|--------------------------|------------|-------------|--------------|------------------|------------------|------------------|
| Site Name                | IP Address | Domain name | Availability | Min Latency (ms) | Avg Latency (ms) | Max Latency (ms) |
| Yahoo                    | 72.30.2.43 | none        | 1.00 %       | 0.00             | 0.00             | 0.00             |

Screenshot 100: The SLA monitor tracks latency and percentage loss over time.

For each SLA object, the Exinda tracks the IP address, percentage of availability, minimum and maximum and average latency in the table below the charts.

- » Availability is the percentage of time a resource is reachable by the Exinda appliance.
- » Latency is the delay in getting an ICMP echo reply for an ICMP echo request generated from the Exinda appliance. It represents both the delay from the local Exinda appliance to a remote host and back again.

| SLA Statistics for Yahoo |            |             |              |                  |                  |                  |
|--------------------------|------------|-------------|--------------|------------------|------------------|------------------|
| Site Name                | IP Address | Domain name | Availability | Min Latency (ms) | Avg Latency (ms) | Max Latency (ms) |
| Yahoo                    | 72.30.2.43 | none        | 1.00 %       | 0.00             | 0.00             | 0.00             |

Where do I find this report?

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Service Levels > Network Response (SLA)**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To add an SLA Site

Click the **Add/Edit SLA Site...** link. See TO-DO for details of configuring an SLA object.

To view the chart for a different SLA Site

Select the desired site from the **SLA Sites** selector.

How do I interact with the interactive flash time graphs?

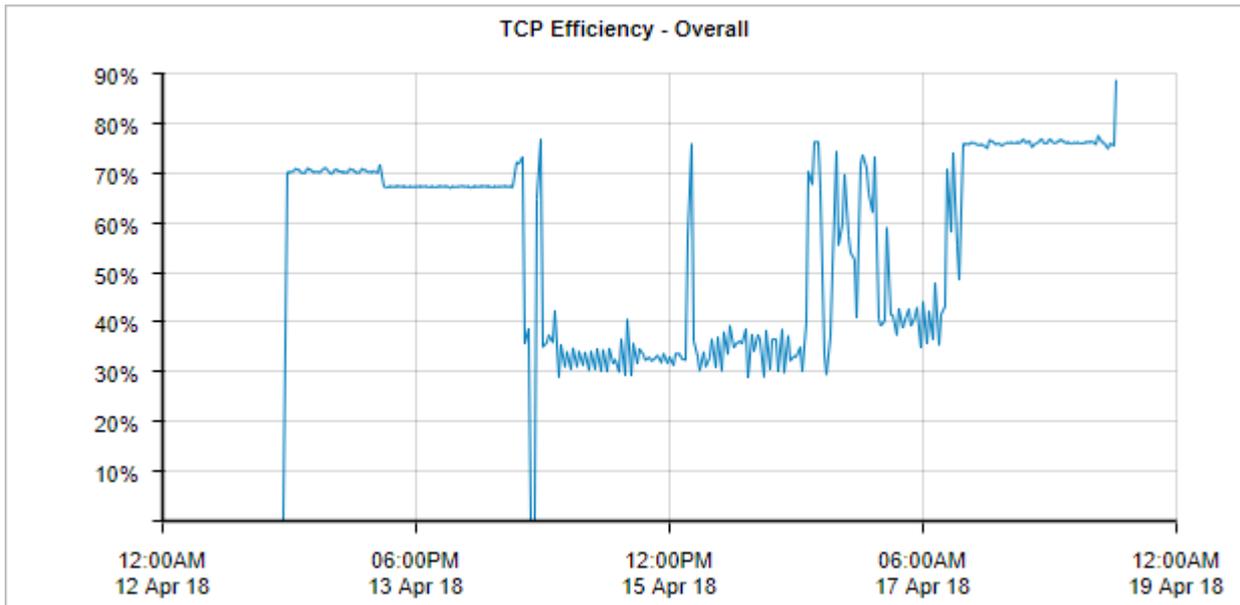
- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## Monitoring TCP efficiency

The TCP Efficiency report shows the total efficiency of all TCP connections over time. The report data can be categorized by applications, internal hosts or external hosts. You can drill into particular apps or hosts to view the efficiency for particular traffic.

This report answers questions such as:

- » Are we experiencing network delays due to TCP inefficiencies?
- » Is a particular application or host having troubles due to retransmissions?



Screenshot 101: The TCP Efficiency report displays TCP connection efficiency over time.

TCP Efficiency is calculated using the formula below:

$$\text{TCP Efficiency} = (\text{Total Bytes} - \text{Bytes Retransmitted}) / \text{Total Bytes}$$

The table below shows both retransmitted bytes and efficiency per Application or Host. Each item in the table below can be drilled down to view TCP Efficiency details and a graph for that item.

| Top 50 Least Efficient Applications    |                    |           |                     |          |                |
|----------------------------------------|--------------------|-----------|---------------------|----------|----------------|
|                                        | Bytes Inbound (MB) |           | Bytes Outbound (MB) |          | Efficiency (%) |
|                                        | Retransmitted      | Total     | Retransmitted       | Total    |                |
| <a href="#">Skype</a>                  | 170.713            | 171.466   | 0.674               | 5.866    | 3.35           |
| <a href="#">CIFS</a>                   | 382.379            | 429.043   | 36.456              | 167.053  | 29.74          |
| <a href="#">HTTPS</a>                  | 96.687             | 308.138   | 465.100             | 624.949  | 39.79          |
| <a href="#">Flash</a>                  | 54.247             | 88.122    | 1.411               | 5.380    | 40.47          |
| <a href="#">BitTorrent</a>             | 94.545             | 255.452   | 244.879             | 413.697  | 49.28          |
| <a href="#">Google Shared Services</a> | 153.453            | 632.337   | 21.121              | 131.789  | 77.15          |
| <a href="#">HTTP</a>                   | 3220.013           | 14015.730 | 352.107             | 1854.006 | 77.49          |
| <a href="#">msmq</a>                   | 0.034              | 0.400     | 0.000               | 0.037    | 92.29          |
| <a href="#">LDAP</a>                   | 1.574              | 21.612    | 1.041               | 16.318   | 93.10          |
| <a href="#">SMTP</a>                   | 99.226             | 3431.120  | 33.630              | 1962.290 | 97.54          |

Screenshot 102: The TCP Efficiency report displays the 50 least efficient applications.

Where do I find this report?

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Service Levels > TCP Efficiency**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## Monitoring TCP health

The TCP Health report shows the number of aborted, refused, and ignored connections over time. The report can be categorized by applications, internal hosts, or external hosts. You can drill into particular apps or hosts to view the health for particular traffic.

This report can answer questions such as:

- » Why are there so many retransmissions for a particular application or host?"

The definitions of aborted, refused and ignored connections used by the TCP Health report:

- » **Aborted**— Connections were established, but were closed by a RST (reset) issued by either the client or server rather than a clean close. High numbers of aborted connections can point to network or server problems.
- » **Refused**— A SYN packet was observed and a RST or ICMP "connection refused" message was received in response. This usually means the server is up, but the application is unavailable or not working correctly. It can also indicate a TCP port scan is occurring.
- » **Ignored**— A SYN packet was observed, but no SYN-ACK response was received. This usually means the server is not responding, does not exist, is not accessible, or is ignoring the connection request. It can also indicate a TCP port scan is occurring.

### VERSION INFO

A new internal mechanism was implemented in the ExOS 7.4.2 firmware that prevents Exinda Appliances from being affected by DDoS attacks. As a consequence, the "Ignored Connections" historical report is no longer available, but the Real Time TCP Health report still includes the number of ignored connections.



Screenshot 103: The TCP Health report displays data about connections over time.

The most unhealthy applications or hosts are shown in the table below the charts. The table shows the number of connections, number of aborted, ignored, and refused connections. You can click the name of the application or host to view the TCP Health details and a graph for that item.

| Top 50 Applications                    |             |         |         |         |
|----------------------------------------|-------------|---------|---------|---------|
|                                        | Connections | Aborted | Ignored | Refused |
| <a href="#">ExindaWM</a>               | 5051599     | 0       | 24184   | 5020297 |
| <a href="#">HTTP</a>                   | 8613838     | 2719847 | 2832    | 1451120 |
| <a href="#">BitTorrent</a>             | 1632737     | 631253  | 2735    | 63598   |
| <a href="#">Google Shared Services</a> | 399416      | 127797  | 86      | 45309   |
| <a href="#">HTTPS</a>                  | 997130      | 100701  | 371     | 49815   |
| <a href="#">Telnet</a>                 | 4100589     | 78683   | 45      | 58080   |
| <a href="#">CIFS</a>                   | 323069      | 36379   | 70      | 26454   |
| <a href="#">FTP</a>                    | 43439       | 16450   | 6       | 3561    |
| <a href="#">MSRPC</a>                  | 279522      | 10935   | 220     | 8678    |
| <a href="#">Wikipedia</a>              | 32573       | 11720   | 0       | 2312    |
| <a href="#">LotusNotes</a>             | 269426      | 3574    | 1421    | 1361    |
| <a href="#">Yahoo</a>                  | 15582       | 3834    | 44      | 2050    |
| <a href="#">Flash</a>                  | 16130       | 4742    | 0       | 66      |
| <a href="#">Bebo</a>                   | 13258       | 2944    | 0       | 1728    |
| <a href="#">SunRPC</a>                 | 3623        | 0       | 0       | 3599    |
| <a href="#">LDAP</a>                   | 52083       | 2331    | 78      | 9       |
| <a href="#">Windows Updates</a>        | 2902        | 1988    | 0       | 4       |
| <a href="#">PPTP</a>                   | 4035        | 450     | 0       | 1502    |
| <a href="#">SMTP</a>                   | 46698       | 46      | 417     | 1472    |
| <a href="#">Amazon Cloud</a>           | 5907        | 743     | 137     | 579     |
| <a href="#">SOAP</a>                   | 1651        | 890     | 0       | 0       |
| <a href="#">xbox-live</a>              | 343         | 270     | 0       | 255     |
| <a href="#">Kerberos</a>               | 724         | 362     | 0       | 0       |
| <a href="#">SSL</a>                    | 1364        | 191     | 5       | 153     |
| <a href="#">Facebook</a>               | 33          | 0       | 6       | 0       |

Screenshot 104: The TCP Health report displays the applications with the most connections.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Service Levels > TCP Health**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## 0.11.6 Monitoring applications

This section provides information about various reports that deal with the performance of your application groups, individual applications, and URLs.

### Monitoring application performance on the network

Analyzing the performance of networked applications is a common task for network administrators because every organization relies on these applications to conduct business operations. Too often, the root causes of poor application performance is misunderstood. And when the root cause is unknown or misdiagnosed, solutions typically involve expensive upgrades to increase and enhance network capacity.

The Exinda Appliance is designed to detect network problems, show them to you and help you uncover root causes, so you can take full advantage of the network hardware and capacity you have and only invest in more when it's truly required.

Exinda Appliances monitors and collects several properties of TCP flows of an application converts them to metrics. These metrics are compared to an established threshold and given a score between one and ten, known as the Application Performance Score (APS). The appliance also monitors single metric values within TCP flows for a specified application, known as Application Performance Metrics (APM).

This allows IT departments to use the Application Performance Score (APS) to determine what is performing well, and what is not. APS and APM have thresholds that identify acceptable performance levels for the applications. When the metric values cross the configured threshold, notifications are sent alerting the necessary users so they can review the issue and make the necessary modifications to allow the applications to perform within the threshold level.

Application Performance Score reports can be easily communicated to senior management and to users to help explain how the applications are performing. The reports can also be used to diagnose and determine where issues are in the network. For each APS score, the results for the metrics can identify the specific area within the network that is affecting the performance of the application, for instance server delays, network delays, or jitter. This makes it easier to fix any network issues and get the application back to optimum performance levels.

For more information refer to:

- » [Configure Application Performance Metric Objects](#)
- » [Application Performance Score Reports](#)
- » [How an Application Performance Score is calculated](#)

---

## Related Tasks

Configuring an application performance metric object.....236

### Monitoring application groups traffic

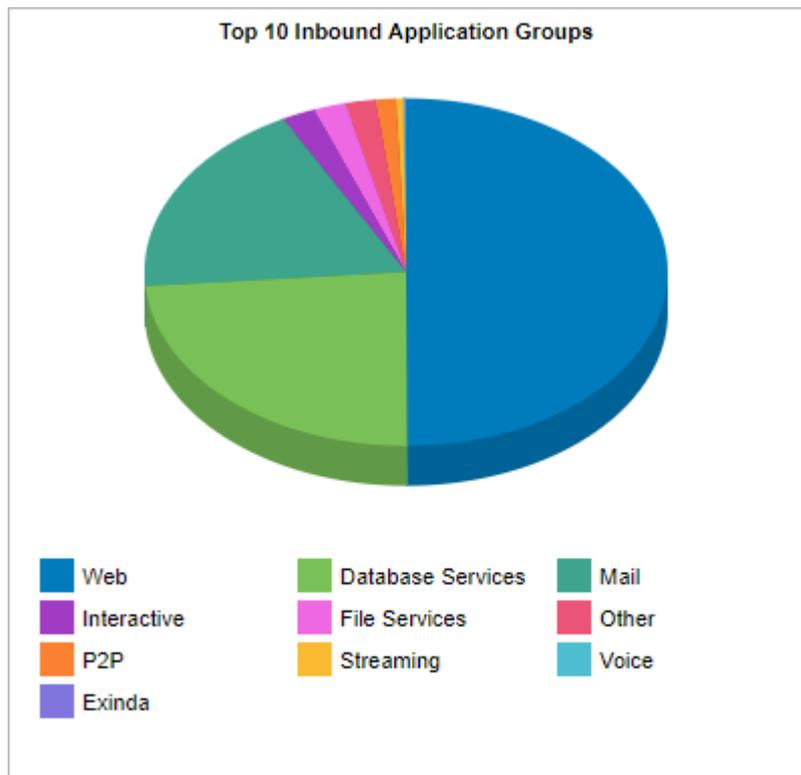
The Traffic Analysis Applications Groups report shows the top application groups by data volume for a selected time period. Inbound and outbound traffic are shown separately.

This report answers questions such as:

- » Which application groups may be overrunning my network?
- » Is the proportion of traffic for a particular application group what I expect?

Use this information to determine if you need to create policies to control or protect high data volume application groups.

You can drill into the application group by clicking on the application group name in the tables below the charts. This shows the [Hosts Report](#) which lists hosts in the selected application group. You can then drill into a particular application to see the hosts using that application.



Screenshot 105: The Applications Group report displays the top 10 inbound application groups.

The tables at the bottom of the report show for each of the top application groups, the total amount of data, and the maximum and average throughput rates, the number of packets, and the number of flows for the selected time period. More network metrics, such as, round-trip time (RTT), network and server delays, and TCP efficiency can be shown by clicking on the **Show Details** link in the tables.

| Top 50 Inbound Application Groups |          |           |                   |         |       |
|-----------------------------------|----------|-----------|-------------------|---------|-------|
| Name                              | Packets  | Data (MB) | Throughput (kbps) |         | Flows |
|                                   |          |           | Average           | Max     |       |
| <a href="#">Web</a>               | 13940088 | 11700.145 | 54.20             | 2037.70 | 3444  |
| <a href="#">Database Services</a> | 23016431 | 6204.919  | 4287.53           | 7552.39 | 13    |
| <a href="#">Mail</a>              | 9025249  | 4715.167  | 60.30             | 2969.56 | 1119  |
| <a href="#">Interactive</a>       | 12130955 | 884.652   | 28.59             | 3709.27 | 542   |
| <a href="#">Other</a>             | 3449552  | 831.518   | 2.01              | 375.37  | 10058 |
| <a href="#">File Services</a>     | 1027800  | 334.036   | 19.35             | 2000.69 | 235   |
| <a href="#">Streaming</a>         | 175030   | 185.987   | 39.31             | 271.30  | 86    |
| <a href="#">P2P</a>               | 279777   | 45.698    | 0.36              | 8.27    | 4299  |
| <a href="#">Exinda</a>            | 756548   | 42.328    | 24.44             | 127.82  | 24    |
| <a href="#">Social Networking</a> | 53292    | 38.503    | 11.12             | 86.47   | 72    |
| <a href="#">Voice</a>             | 60164    | 6.074     | 9.05              | 22.37   | 17    |
| <a href="#">Software Updates</a>  | 5333     | 3.391     | 5.12              | 12.78   | 12    |
| <a href="#">Thin Client</a>       | 3732     | 0.339     | 0.58              | 1.37    | 12    |

Screenshot 106: The Application Groups report displays traffic volume from the top application groups.

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Applications > Application Groups**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

You can customize the applications objects included in an application group. For more information, refer to [Adding and updating application group objects](#) (page 200).

To interact with the pie-based reports, you can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie. Note that the pie is showing only the top items, so the proportion is relative to the top items - not relative to all the traffic through the appliance. That is, if one wedge showed 50% of the traffic, that means it is 50% of the top items, not 50% through the appliance.

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

#### Related topic

[View a network summary of application groups](#)

#### Viewing a network summary of application groups

Each table shows the top application groups together with the number of packets, number of flows data transferred and throughput statistics.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Monitor > Application Groups**.
6. To expose Round trip time, Normalized Delays, Transaction Delays, and Efficiency statistics for each Application Group, click **Show Details**.

| Top 50 Inbound Application Groups |          |           |                   |          |       |
|-----------------------------------|----------|-----------|-------------------|----------|-------|
| Name                              | Packets  | Data (MB) | Throughput (kbps) |          | Flows |
|                                   |          |           | Average           | Max      |       |
| <a href="#">[+] Show Details</a>  |          |           |                   |          |       |
| <a href="#">Web</a>               | 16428058 | 14509.372 | 61.90             | 35854.61 | 4042  |
| <a href="#">Database Services</a> | 24611856 | 6626.318  | 4201.48           | 10742.92 | 16    |
| <a href="#">Mail</a>              | 9729809  | 5122.511  | 60.46             | 5186.96  | 1305  |
| <a href="#">File Services</a>     | 2165485  | 1889.490  | 100.80            | 36089.38 | 282   |
| <a href="#">Interactive</a>       | 12651096 | 924.558   | 27.49             | 3709.27  | 638   |
| <a href="#">Other</a>             | 3681629  | 890.287   | 1.98              | 510.46   | 12778 |
| <a href="#">Streaming</a>         | 198242   | 212.644   | 40.56             | 465.38   | 105   |
| <a href="#">P2P</a>               | 339351   | 57.992    | 0.42              | 42.56    | 5033  |
| <a href="#">Exinda</a>            | 813537   | 45.508    | 24.09             | 127.82   | 28    |
| <a href="#">Social Networking</a> | 57116    | 41.221    | 10.86             | 86.47    | 92    |
| <a href="#">Voice</a>             | 68218    | 10.905    | 14.75             | 453.74   | 23    |
| <a href="#">Software Updates</a>  | 5615     | 3.572     | 5.04              | 12.78    | 14    |
| <a href="#">Thin Client</a>       | 5280     | 0.633     | 0.93              | 8.28     | 16    |
| <a href="#">VPN</a>               | 192      | 0.053     | 2.35              | 5.89     | 5     |
| <a href="#">Games</a>             | 26       | 0.005     | 1.43              | 1.94     | 1     |

| Top 50 Outbound Application Groups |          |           |                   |          |       |
|------------------------------------|----------|-----------|-------------------|----------|-------|
| Name                               | Packets  | Data (MB) | Throughput (kbps) |          | Flows |
|                                    |          |           | Average           | Max      |       |
| <a href="#">[+] Show Details</a>   |          |           |                   |          |       |
| <a href="#">Interactive</a>        | 14405208 | 18129.099 | 82.13             | 50807.61 | 3369  |
| <a href="#">Web</a>                | 13127597 | 2408.708  | 13.40             | 1221.63  | 5332  |
| <a href="#">Mail</a>               | 2837334  | 2015.518  | 46.42             | 1585.91  | 1310  |
| <a href="#">Database Services</a>  | 25238049 | 1922.948  | 290.54            | 2082.32  | 126   |
| <a href="#">Other</a>              | 6142829  | 1854.472  | 2.87              | 913.88   | 27337 |
| <a href="#">P2P</a>                | 2049480  | 400.489   | 2.66              | 15.52    | 5934  |
| <a href="#">File Services</a>      | 1123370  | 163.823   | 9.91              | 155.03   | 469   |
| <a href="#">Exinda</a>             | 536095   | 39.993    | 30.14             | 57.53    | 28    |
| <a href="#">Voice</a>              | 128276   | 24.122    | 14.05             | 32.09    | 78    |
| <a href="#">Streaming</a>          | 60829    | 7.571     | 3.15              | 45.74    | 105   |
| <a href="#">Social Networking</a>  | 38351    | 6.943     | 2.38              | 7.34     | 120   |
| <a href="#">Thin Client</a>        | 27843    | 1.971     | 0.36              | 6.66     | 146   |
| <a href="#">Software Updates</a>   | 5102     | 0.977     | 1.44              | 1.90     | 28    |
| <a href="#">VPN</a>                | 234      | 0.044     | 2.65              | 6.01     | 5     |
| <a href="#">Games</a>              | 38       | 0.008     | 2.34              | 3.46     | 1     |

7. To view the data for individual applications within a group, click the application group name.

### Viewing application traffic volume

The Applications report shows the top applications by volume and their average throughput. Volume and throughput data for single applications can be graphed by clicking the filter icon for the desired application in the data table below the graphs. Inbound and outbound LAN application traffic is reported separately.

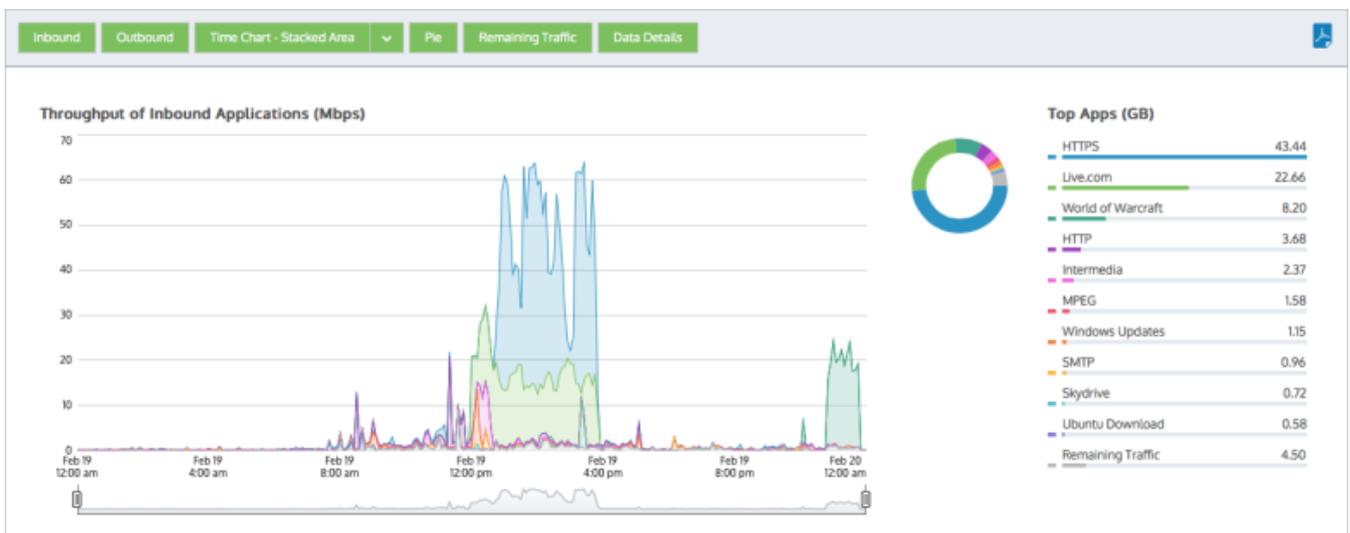
To show all application traffic, [add a category](#) to represent the remaining application traffic on your network. Doing this allows the cumulative stack on the throughput chart to represent all the application traffic moving through the appliance.

This will help you understand the significance of the top applications relative to the whole. In addition to showing a stacked cumulative view, you can display the throughput as a line chart with a common zero baseline. You can also display the application volumes as a pie chart.

These charts can answer questions such as:

- » What are the top applications on my network?
- » Are those top applications significant relative to the entire traffic?
- » How much bandwidth does my FTP application typically take?
- » Could one application be choking out the other application traffic?
- » Do any of my top applications appear to be limited?

Using this information, you can determine if you need to create policies for high data volume applications and applications that tend to have large data volume spikes. You may want to create protection policies for your business critical apps and limiting policies for high volume non-business critical applications like recreational applications. For more information, refer to [Adding a policy to the policy tree](#) (page 372).



Screenshot 107: The Applications report shows traffic volume graphed over time.

#### NOTE

Average bandwidth is calculated as the total bits observed in the charting interval divided by the number of seconds in that interval. E.g. For a chart with an hour of data, the intervals are five minutes.

If you drilled into the applications chart from any of the virtual circuit, subnets, or hosts charts, then the relevant virtual circuit, subnet, or host will be shown on the filter bar below the button bar. To turn off the filtering, click on the close 'x' in the filter tag.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Applications > Applications**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To filter the report data

Various components on the screen can be toggled on and off by clicking buttons above the charts. Note that when generating a pdf report of this screen, the toggle states are taken into account. That is, if you had toggled off the outbound charts, they will not be present in the pdf.

- » **Host Type:** When you first load the Hosts report, Internal hosts are graphed by default. Click the Internal hosts button and then select External hosts to change the type. Note that you cannot graph both internal and external hosts at the same time.
- » **Traffic Type (Inbound/Outbound):** By default, both Inbound and Outbound traffic is graphed. Click either the Inbound or Outbound option to hide the data, including all the carts and the data tables below the charts.
- » **Chart Type:** Toggles on or off the time series charts and allows selecting stacked area display versus a line chart display.
- » **Pie:** Toggles on or off a colour-coded Pie chart to the left of the Top Listeners and Top Talkers list.
- » **Remaining Traffic:** Append or hide the Remaining Traffic category below the Top Listeners and Top Talkers lists. This option toggles on or off the presence of a category for all the remaining applications summed together.
- » **Data Details:** Toggles on or off the data tables below the time series charts.
- » **Mouse Hover:** Hover the mouse pointer over the graph to view data throughput at a given date and time.

Turning on or off the remaining traffic category

Toggle the Remaining Traffic button on the button bar. When toggled on, a grey chart series will appear in all charts (throughput, pie, top applications) representing all the applications on your network that are not explicitly represented in the top applications. If the remaining applications show vastly more data volume than the top applications, then the top applications may look insignificant relative to the total, and so you may need to toggle off the remaining traffic category to see the relative differences and usage patterns of the top applications.

Changing the throughput chart to stacked area charts or line charts

Press the down arrow next to the dropdown list at the top of the page and choose which chart type to show. The line chart shows the applications against the common zero baseline so they can be compared to each other and the pattern of a specific application is clearer. You can look for particular patterns such as spikes or flat tops.

Determining if one or more applications may be choking out the other application traffic

View the throughput charts with the remaining traffic category toggled on. For any periods where the cumulative throughput is especially high (compared to the pipe size that this appliance is managing) is there an application or two that is consuming a significant portion of the bandwidth. If so, that application may be choking out other application and would be a candidate for control. Please note that you may need to go to the virtual circuit chart and filter the applications by the individual virtual circuits to understand whether an application is choking out others since the virtual circuits share bandwidth and an application may overrun one virtual circuit but not others.

Determining if any of the top applications appear to be limited

View the throughput charts as a line chart with the remaining traffic category toggled off. If any of the lines representing the applications have raised flat tops, this may represent that the application is being limited by a policy or the capacity of your pipe.

Charting any single application

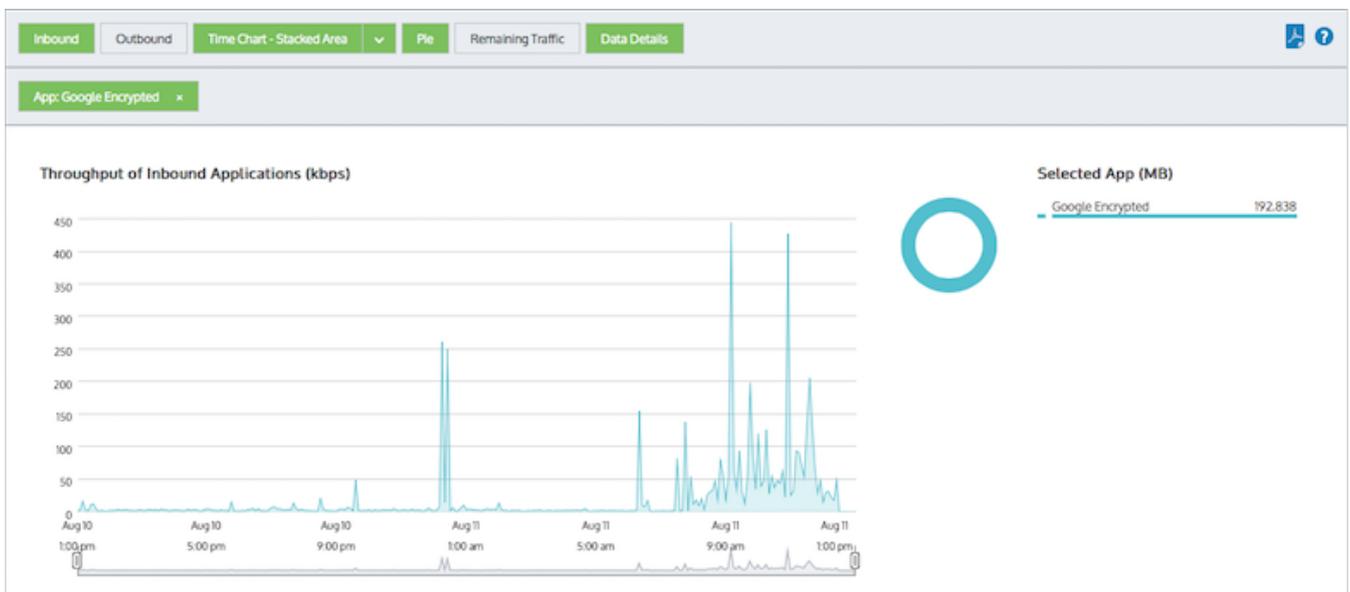
## VERSION INFO

Filtering to show any single application is available in 7.0.1 and above.

In the data table, each application has a filter icon on the right-hand-side of the row. By clicking on the filter icon, as shown below for Google Encrypted, only the selected application will be charted.

|                  |          |            |   |
|------------------|----------|------------|---|
| Grooveshark      | 0.379 GB | 0.036 Mbps | 🔼 |
| Google Encrypted | 0.194 GB | 0.018 Mbps | 🔽 |
| DropBox          | 0.142 GB | 0.013 Mbps | 🔼 |
| MPEG             | 0.098 GB | 0.009 Mbps | 🔼 |

Screenshot 108: Click filter icon beside desired application to chart that application.



Screenshot 109: Applications are filtered to only show 'Google Encrypted'

While in filter mode, any other application filter icon can be clicked to change which application is charted.

To remove this filter and return to the top application set, click the 'x' on the green 'App: Google Encrypted' filter tag.

To show more applications or fewer applications in the top applications chart and the throughput chart

The number of applications shown are configurable by using the **Chart Items** setting on the **Configuration > System > Setup > Monitoring** page. Please note that this configuration applies to all charts on the appliance. See [Monitoring Configuration](#).

How do I interact with the new time-series & bar chart reports?

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how the charts interact and what the toggle buttons do, see [Understanding How Charts Relate](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand the difference between inbound and outbound traffic, see [Understanding Traffic Direction](#).

- » To understand how many data points are shown for each time period, see [Understanding Traffic Granularity](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

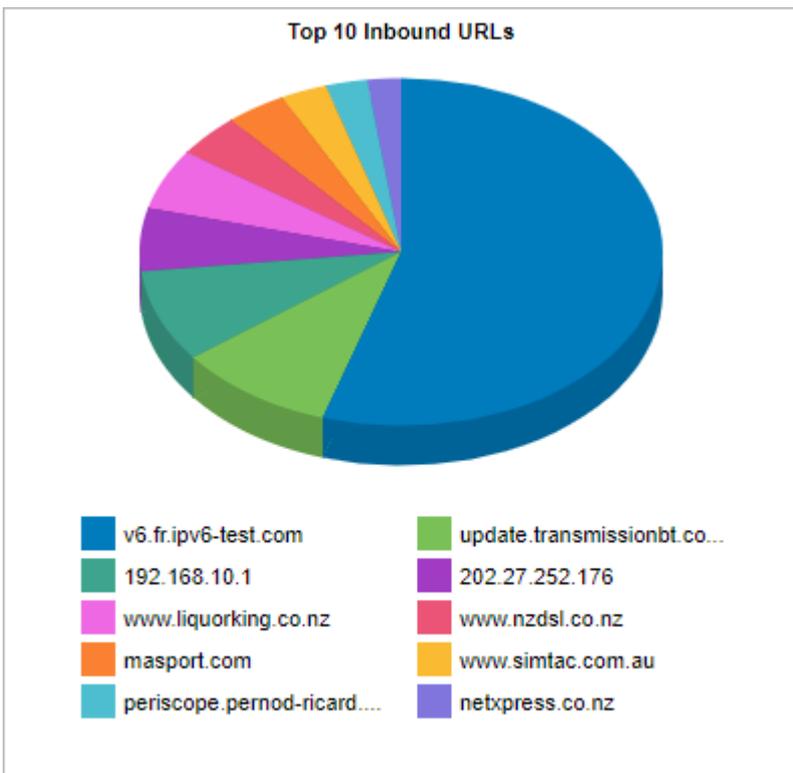
### Monitoring URLs visited

The URLs report shows the top URLs visited by data volume for the selected time period. The URLs report shows inbound traffic separately from outbound traffic. This report answers questions such as:

- » Which websites are generating the most traffic?

Using this information you can determine if you need to create applications based on URLs and create policies to control or protect high data volume URLs.

The URL names are represented as a domain/host name. Drill into the URLs by clicking on the URL name in the tables below the charts. This will show the [Hosts Report](#) which lists hosts that visited the URL.



The tables at the bottom of the report displays the total amount of data and the maximum and average throughput rates, the number of packets, and the number of flows for the selected time period for the top URLs. More network metrics like round-trip time (RTT), network and server delays, and TCP efficiency can be shown by clicking on the **Show Details** link in the tables.

| Top 50 Inbound URLs                         |          |           |                   |          |       |
|---------------------------------------------|----------|-----------|-------------------|----------|-------|
| Name                                        | Packets  | Data (MB) | Throughput (kbps) |          | Flows |
|                                             |          |           | Average           | Max      |       |
| <a href="#">v6.fr.ipv6-test.com</a>         | 34291129 | 42353.671 | 4360.44           | 20961.79 | 6     |
| <a href="#">update.transmissionbt.com</a>   | 5243367  | 7538.356  | 3560.60           | 9427.21  | 6     |
| <a href="#">192.168.10.1</a>                | 4796627  | 6636.374  | 35458.56          | 17163.58 | 1     |
| <a href="#">202.27.252.176</a>              | 3559228  | 4535.962  | 1160.78           | 3286.27  | 12    |
| <a href="#">www.liquorking.co.nz</a>        | 3655802  | 4348.822  | 895.23            | 2441.56  | 12    |
| <a href="#">www.nzdsl.co.nz</a>             | 2879458  | 3091.721  | 1873.93           | 3856.20  | 6     |
| <a href="#">masport.com</a>                 | 2644319  | 2890.873  | 590.03            | 1624.00  | 6     |
| <a href="#">www.simtac.com.au</a>           | 2985241  | 2221.415  | 226.18            | 608.41   | 12    |
| <a href="#">periscope.pernod-ricard.com</a> | 2356400  | 1993.273  | 92.50             | 441.25   | 76    |
| <a href="#">netxpress.co.nz</a>             | 1922936  | 1619.214  | 268.28            | 556.58   | 6     |
| <a href="#">219.88.73.253</a>               | 1680382  | 1434.416  | 657.53            | 1692.41  | 6     |

Screenshot 110: The URLs report displays traffic volume by inbound URL.

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.

Go to **Monitor > Applications > URLs**.

To interact with the pie-based reports, you can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie. Note that the pie is showing only the top items, so the proportion is relative to the top items - not relative to all the traffic through the appliance. That is, if one wedge showed 50% of the traffic, that means it is 50% of the top items, not 50% through the appliance.

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

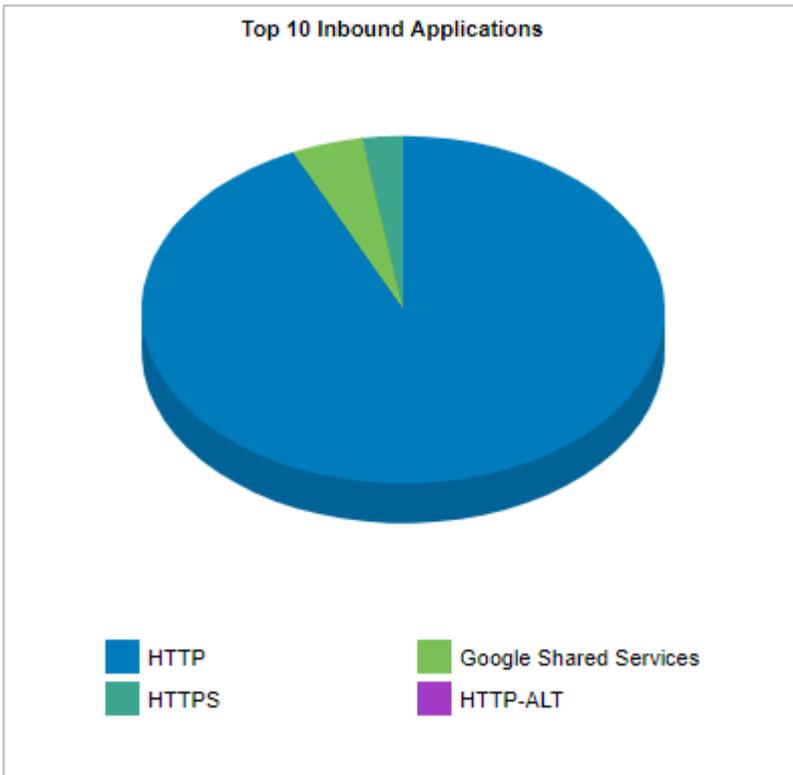
Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

## Drilling into application data

The Application Drill-in report shows application specific traffic data volume for a selected time period. Inbound and outbound traffic are shown separately. This report answers questions such as:

- » Which applications are part of the application group that I clicked on?
- » Which applications did a particular user or host use?"

You can drill into the application by clicking on the application name in the tables below the charts. This will show the [Hosts Report](#) which lists hosts that used the application.



Screenshot 111: The Applications report displays a graph of traffic volume by application.

The tables at the bottom of the report show the total amount of data, the maximum and average throughput rates, the number of packets, and the number of flows by application for the selected time period. Click on the **Show Details** link in the Name column to see more metrics like round-trip time (RTT), network and server delays, and TCP efficiency.

| Top 50 Inbound Applications              |           |            |                   |          |       |
|------------------------------------------|-----------|------------|-------------------|----------|-------|
| Name<br><a href="#">[+] Show Details</a> | Packets   | Data (MB)  | Throughput (kbps) |          | Flows |
|                                          |           |            | Average           | Max      |       |
| <a href="#">HTTP</a>                     | 174038232 | 180238.195 | 112.50            | 25583.98 | 1832  |
| <a href="#">Google Shared Services</a>   | 10378299  | 8590.616   | 97.73             | 1699.17  | 106   |
| <a href="#">HTTPS</a>                    | 9544782   | 4873.738   | 49.26             | 699.53   | 100   |
| <a href="#">HTTP-ALT</a>                 | 57569     | 5.146      | 0.98              | 2.47     | 7     |

Access this report by drilling in from other reports, such as application group, hosts, users, conversations, subnets.

To interact with the pie-based reports, you can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie. Note that the pie is showing only the top items, so the proportion is relative to the top items - not relative to all the traffic through the appliance. That is, if one wedge showed 50% of the traffic, that means it is 50% of the top items, not 50% through the appliance.

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

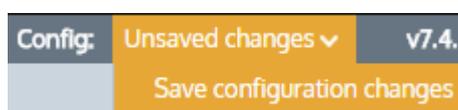
### Disabling calculations of application performance metrics

Stop the Exinda Appliance from calculating Round Trip Time (RTT), Network and Server Delay, Loss and Efficiency, and TCP Health.

## IMPORTANT

Application performance metrics must be enabled to calculate Application Performance Scores.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Setup** and switch to the **Monitoring** tab.
6. In the ASAM section, uncheck the **Performance Metrics** checkbox.
7. Click **Apply Changes**.
8. To save the changes to the configuration file, in the status bar click the **Unsaved changes** menu and select **Save configuration changes**.



### 0.11.7 Monitoring network users

The Users report shows the top users by data volume for a selected time period. Inbound and outbound traffic are reported separately. You can view internal and external users in the report, answering questions such as:

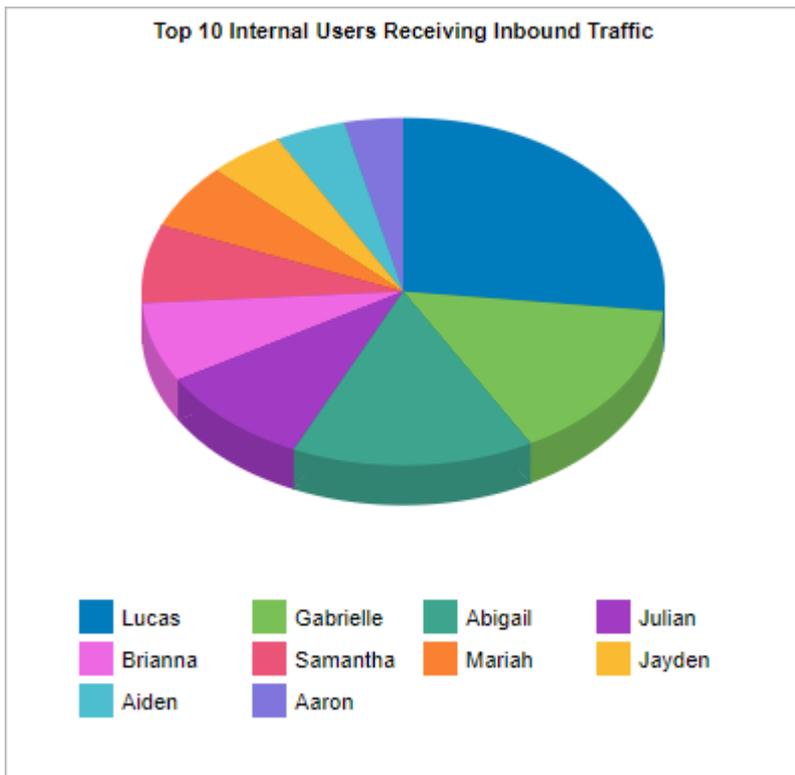
- » What internal users are the top talkers and top listeners?
- » Which external users are top talkers?
- » Which external users are top listeners?
- » Is one user choking the network?

Using this information, you can determine if you need to create policies for these high data volume users. You may want to create protection policies for your important users, like your CEO or finance department, or create control policies to limit users who are abusing the network.

In this report, users are associated with IP addresses. Network traffic flows from one host to another and typically, one host is considered internal to your network while the other is considered external.

Hosts that fall into a [network object](#) defined as internal are considered internal to your network. Hosts that fall into a network object defined as external are considered external to your network. Keep in mind that the traffic is inbound and outbound relative to your LAN – not relative to the host or user. Inbound traffic for an external user means a user was sending data into your network.

You can drill into the user by clicking on the user name in the tables below the charts. This will show the [Applications Report](#) for the user that you drilled into. You can then use the selector on the Applications report page to show URLs or conversations or hosts that involved the user.



Screenshot 112: The Users report displays traffic volume by user.

The tables at the bottom of the report shows for each of the top users, the total amount of data, and the maximum and average throughput rates, the number of packets, and the number of flows for the selected time period. More network metrics, such as, round-trip time (RTT), network and server delays, and TCP efficiency can be shown by clicking on the **Show Details** link in the tables.

| Top 50 Internal Users Receiving Inbound Traffic |          |           |                   |         |       |
|-------------------------------------------------|----------|-----------|-------------------|---------|-------|
| Name                                            | Packets  | Data (MB) | Throughput (kbps) |         | Flows |
|                                                 |          |           | Average           | Max     |       |
| <a href="#">Lucas</a>                           | 49129154 | 28960.976 | 6.32              | 9427.21 | 7800  |
| <a href="#">Gabrielle</a>                       | 14896978 | 16293.027 | 864.82            | 4323.26 | 14    |
| <a href="#">Abigail</a>                         | 13230315 | 16148.534 | 42.78             | 496.81  | 425   |
| <a href="#">Julian</a>                          | 8925277  | 10418.744 | 385.88            | 3286.27 | 33    |
| <a href="#">Brianna</a>                         | 9254896  | 8004.930  | 56.33             | 3856.20 | 195   |
| <a href="#">Samantha</a>                        | 11080874 | 7911.251  | 163.18            | 2434.68 | 76    |
| <a href="#">Mariah</a>                          | 5174481  | 6551.210  | 203.14            | 2441.56 | 48    |

Screenshot 113: The table on the Users report shows traffic volume metrics broken down by user.

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Users**.

To show only internal users or external users, use the **Select Users to View** selector at the top of the page.

To interact with the pie-based reports, you can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie. Note that the pie is showing only the top items, so the proportion is relative to the top

items - not relative to all the traffic through the appliance. That is, if one wedge showed 50% of the traffic, that means it is 50% of the top items, not 50% through the appliance.

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

## Related topic

[Set the Time Period Reflected in the Report](#)

## Setting the time period for a report

To limit report data to specific periods of time, set the date range. Viewing reports by date range is available on all reports except Realtime reports.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Select a report from the Monitor list.
6. Beside the title of the report, select the desired date range from the drop down list.

**Range:**  12:00AM 16/Nov/2009 - 12:00AM 17/Nov/2009

7. To specify a custom date range, in the drop down list select **Custom**. Select the start and end date and time to include in the report. After the date range is selected, the graphs and charts are immediately updated.

**Range:**  12:00AM 25/Oct/2010 - 12:00AM 26/Oct/2010

## Temporal granularity of stored data

The Exinda Appliance stores data for the following time intervals:

- » 2 years of data - this year, previous year & last 12 months
- » 2 months of data - this month, previous month & last 30 days
- » 2 weeks of data - this week, previous week & last 7 days
- » 2 days of data - today, yesterday & last 24 hours
- » 1 day of data - this hour, last hour & last 60 minutes, last 5 minutes

For the Applications, URLs, Users, Hosts, Conversations and Subnets Reports, the data is stored at:

- » Hourly granularity for up to 2 days (today, yesterday, this hour, previous hour)
- » Daily granularity for up to 2 months (this week, last week, this month and last month)
- » Monthly granularity for up to 2 years (this year, last year)

For the Interface, Network, Reduction, Optimizer, Service Levels, System the data is stored at:

- » 10 second granularity for 1 day (except Network)
- » 5 minute granularity for 2 weeks
- » 30 minute granularity for 2 months
- » 60 minute granularity for 6 months
- » 24 hour granularity for 2 years

### 0.11.8 Monitoring hosts traffic volume

The Hosts report shows the top hosts by data volume for the selected time period. For more information, refer to [Setting the time period for a report](#) (page 284).

Traffic inbound into your LAN is reported separately from the outbound traffic. You can view internal and external hosts and data is graphed separately for Top Listeners and Top Talkers. This allows multi-site enterprises to monitor corporate systems while excluding Internet servers.

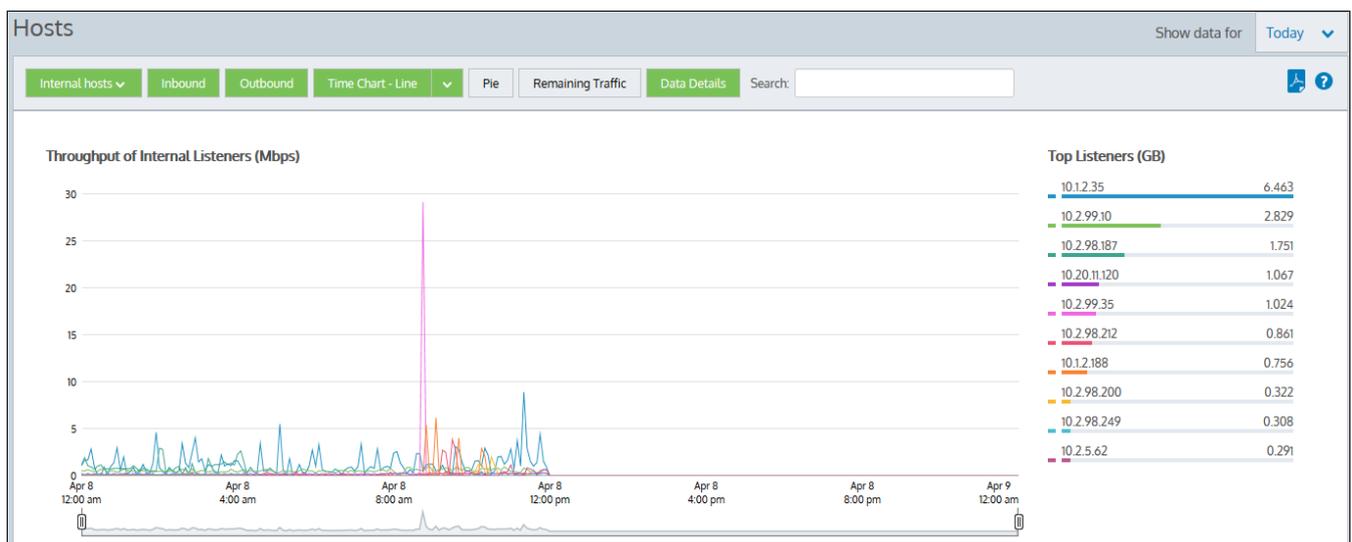
This report answers questions such as:

- » What internal hosts are the top talkers and top listeners?
- » Which external hosts are top talkers from which internal hosts are retrieving information?
- » Which external hosts are top listeners from which internal hosts are sending information to?
- » Could one host be choking out my network?

Use this information to determine if you need to create policies for these high data volume hosts. You may want to create protection policies for your business critical server machines or create control policies to limit hosts that are abusing the network.

#### VERSION INFO

The hosts report as a time-series is available in version 7.0.3 and higher.



Screenshot 114: The Hosts report displays traffic volume over time and top listeners.

#### AVERAGE BANDWIDTH

Average bandwidth is calculated as the total bits observed in the charting interval divided by the number of seconds in that interval. E.g. For a chart with an hour of data, the intervals are five minutes.

## What are hosts?

Hosts are IP Address endpoint's in IP transactions and are usually client PCs or servers. During a flow, traffic flows from one host to another. Typically, one host is considered internal to your network; the other is external:

- » Hosts that fall into a network object that was defined as **internal** are considered internal to your network
- » Hosts that fall into a network object that was defined as **external** are considered external to your network

Traffic is inbound and outbound relative to your LAN – not relative to the host. Therefore, inbound traffic for an external host means that host was sending data inbound into your network.

## Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Hosts**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To filter the report data

Toggle various chart elements on and off by clicking the buttons above the charts. Note that when generating a PDF report of this screen, the toggle states are taken into account. That is, if you had toggled off the outbound charts, they will not be present in the PDF.

- » **Host Type:** When you first load the Hosts report, Internal hosts are graphed by default. Click the Internal hosts button and then select External hosts to change the type. Note that you cannot graph both internal and external hosts at the same time.
- » **Traffic Type (Inbound/Outbound):** By default, both Inbound and Outbound traffic is graphed. Click either the Inbound or Outbound option to hide the data. When viewing Internal hosts, hiding the Inbound data toggles off the Top Listeners data from the graphs, whereas hiding the Outbound data toggles off the Top Talkers data. When viewing External hosts, the opposite is true.
- » **Chart Type:** The chart is initially mapped as a Stacked Area , but you can change the format to Line Chart if necessary.
- » **Pie:** Toggles on or off a colour-coded Pie chart to the left of the Top Listeners and Top Talkers list.
- » **Remaining Traffic:** Append or hide the Remaining Traffic data below the Top Listeners and Top Talkers lists. Remaining traffic represents the remaining application traffic on your network and the cumulative stack on the throughput chart represents all the hosts communicating through the appliance. If the remaining traffic show vastly more data volume than the top hosts, then the top hosts may look insignificant relative to the total, and so you may need to toggle off the remaining traffic category to see the relative differences and usage patterns of the top hosts.

### NOTE

If there are more than 100,000 hosts to display, it may take several minutes to render the screen when Remaining Traffic is enabled.

- » **Data Details:** Toggles on or off the data tables below the time series charts.
- » **Mouse Hover:** Hover the mouse pointer over the graph to view data throughput at a given date and time. Refer to Chart Interactions - Drill in & Data brush in [WUI Guided Tour](#) for details.



Screenshot 115: The Hosts report displays throughput by internal listeners over time broken down by top listeners and talkers.

### Drilling down into report data

Drill into the host data by clicking on a host in the Top Listeners or Top Talkers list (located to the right of the graphs). Click a particular host to view the [Applications Report](#) for the host that you selected. You can then use the selector on the Applications Report page to show URLs or conversations that involved the host.

The tables at the bottom of the Hosts report information for the top listeners and talkers and include the IP Address, the Total Volume of data, and the Average Throughput rates. Click on any entry in the table to open the Applications Report for that specific host.

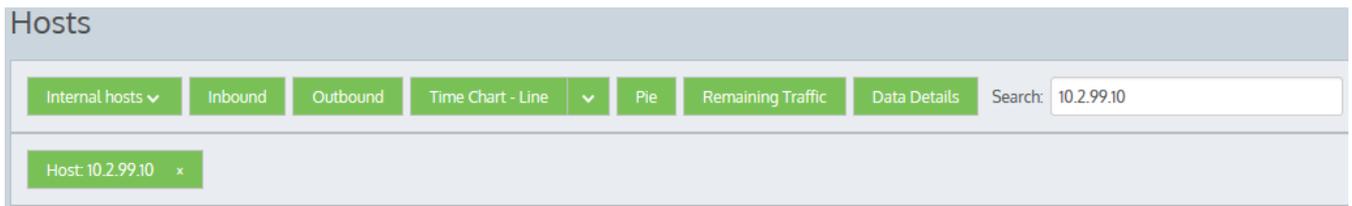
| Internal Listeners |              |                | Internal Talkers |              |                |
|--------------------|--------------|----------------|------------------|--------------|----------------|
| Name               | Total Volume | Avg Throughput | Name             | Total Volume | Avg Throughput |
| 10.1.2.35          | 9.266 GB     | 0.921 Mbps     | 10.2.99.10       | 73.596 GB    | 7.317 Mbps     |
| 10.3.141.155       | 4.715 GB     | 0.469 Mbps     | 10.3.139.172     | 7.328 GB     | 0.729 Mbps     |
| 10.2.98.200        | 2.318 GB     | 0.230 Mbps     | 10.1.2.35        | 1.918 GB     | 0.191 Mbps     |
| 10.2.98.212        | 2.228 GB     | 0.222 Mbps     | 10.10.10.10      | 1.423 GB     | 0.141 Mbps     |
| 10.2.99.10         | 1.863 GB     | 0.185 Mbps     | 10.11.23         | 0.438 GB     | 0.044 Mbps     |
| 10.2.135.93        | 1.451 GB     | 0.144 Mbps     | 10.10.7.77       | 0.321 GB     | 0.032 Mbps     |
| 10.20.11.120       | 1.379 GB     | 0.137 Mbps     | 10.2.98.173      | 0.285 GB     | 0.028 Mbps     |
| 10.2.98.187        | 1.174 GB     | 0.117 Mbps     | 10.2.98.200      | 0.241 GB     | 0.024 Mbps     |

Screenshot 116: Drilling down into hosts data.

### Searching for a specific host

If the host you are looking for is not listed in the Top hosts, you can use the search function to locate data for a single host only. Type a single IP Address in the Search field to locate data for a particular host. If entering an IPv6 host, use the full IPv6 address only. When the data is retrieved, the individual host is shown on the filter bar below the button bar. To

turn off the filtering, click on the close 'x' in the filter tag.



How do I interact with the new time-series & bar chart reports?

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how the charts interact and what the toggle buttons do, see [Understanding How Charts Relate](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand the difference between inbound and outbound traffic, see [Understanding Traffic Direction](#).
- » To understand how many data points are shown for each time period, see [Understanding Traffic Granularity](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### 0.11.9 Monitoring network conversations

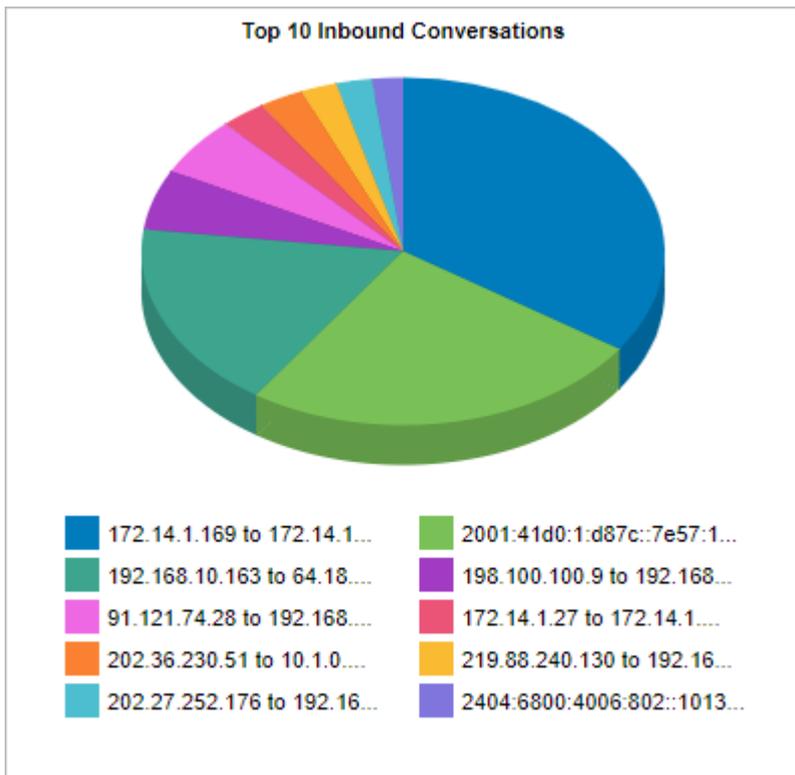
The Conversations report shows top conversations by data volume for a selected time period. Traffic inbound to your LAN is reported separately from the outbound traffic.

This report answer questions such as:

- » What are the top conversations on my network?
- » Could one conversation be choking out the other application traffic?

Use this information to determine if you need to create policies for high data volume conversations. You may want to create [limiting policies](#) for particular hosts or users accessing particular applications.

A conversation is defined as data transacted between two host machines using the same application within a specified time period. Conversations may also be referred to as sessions.



Screenshot 117: The Conversations report displays traffic volume data by conversation.

The tables at the bottom of the report shows for each of the top conversations, the total amount of data, and the maximum and average throughput rates, the number of packets, and the number of flows for the selected time period. More network metrics, such as, round-trip time (RTT), network and server delays, and TCP efficiency can be shown by clicking on the **Show Details** link in the tables.

| Top 50 Inbound Conversations                      |                                                |                                        |           |                   |          |       |
|---------------------------------------------------|------------------------------------------------|----------------------------------------|-----------|-------------------|----------|-------|
| External Host<br><a href="#">[+] Show Details</a> | Internal Host                                  | Application                            | Data (MB) | Throughput (kbps) |          | Flows |
|                                                   |                                                |                                        |           | Average           | Max      |       |
| <a href="#">172.14.1.169</a>                      | <a href="#">172.14.1.145</a>                   | <a href="#">MySQL</a>                  | 81113.665 | 7445.35           | 25103.24 | 6     |
| <a href="#">2001:41d0:1:d87c::7e57:1</a>          | <a href="#">dead:beef::7a2b:cbff:fe91:daab</a> | <a href="#">HTTP</a>                   | 58645.827 | 5196.54           | 23668.62 | 6     |
| <a href="#">192.168.10.163</a>                    | <a href="#">64.18.6.12</a>                     | <a href="#">SMTP</a>                   | 41148.742 | 3266.59           | 10095.75 | 6     |
| <a href="#">198.100.100.9</a>                     | <a href="#">192.168.10.71</a>                  | <a href="#">LotusNotes</a>             | 13228.160 | 1344.55           | 4323.26  | 6     |
| <a href="#">91.121.74.28</a>                      | <a href="#">192.168.10.168</a>                 | <a href="#">HTTP</a>                   | 12774.326 | 2605.37           | 9427.21  | 6     |
| <a href="#">172.14.1.27</a>                       | <a href="#">172.14.1.145</a>                   | <a href="#">SSH</a>                    | 6521.001  | 881.87            | 3716.86  | 6     |
| <a href="#">202.36.230.51</a>                     | <a href="#">10.1.0.224</a>                     | <a href="#">HTTP</a>                   | 6451.551  | 725.85            | 2286.82  | 6     |
| <a href="#">219.88.240.130</a>                    | <a href="#">192.168.10.2</a>                   | <a href="#">HTTP</a>                   | 5277.952  | 1096.99           | 3856.20  | 6     |
| <a href="#">202.27.252.176</a>                    | <a href="#">192.168.10.174</a>                 | <a href="#">HTTP</a>                   | 5204.047  | 1030.81           | 3286.27  | 6     |
| <a href="#">2404:6800:4006:802::1013</a>          | <a href="#">dead:beef::7a2b:cbff:fe91:daab</a> | <a href="#">Google Shared Services</a> | 4552.360  | 427.97            | 1699.17  | 6     |

To access this report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Conversations**.

To interact with the pie-based reports, you can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie. Note that the pie is showing only the top items, so the proportion is relative to the top items - not relative to all the traffic through the appliance. That is, if one wedge showed 50% of the traffic, that means it is 50% of the top items, not 50% through the appliance.

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

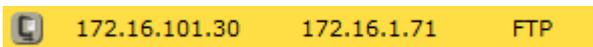
## Related topic

### Understanding the Conversation Report

#### Understanding the conversation report

Accelerated conversations are highlighted in yellow and the application acceleration technologies applied to that conversation are displayed on the left-hand side as a series of icons.

For example, the FTP connection below is accelerated and processed by WAN Memory, indicated by the icon.



Conversations processed by Edge Cache (see the *Edge Cache "How-to" Guide*), are highlighted in blue.



The following legend describes the meaning of each acceleration icon.

|  |                                                                           |
|--|---------------------------------------------------------------------------|
|  | WAN Memory: The connection is been processed by WAN Memory.               |
|  | CIFS Acceleration: The connection is been processed by CIFS Acceleration. |
|  | SSL Acceleration: The connection is been processed by SSL Acceleration.   |
|  | NCP Acceleration: The connection is been processed by NCP Acceleration.   |
|  | MAPI Acceleration: The connection is been processed by MAPI Acceleration. |

When an appliance is deployed in a High Availability (HA) or Clustering mode, the following icons may appear next to conversations.

|  |                                                                                                                                                        |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Asymmetric: The traffic is asymmetric, and is not being accelerated.                                                                                   |
|  | Locally Bridged: The connection is passing through this appliance in the cluster. If the row is colored yellow, then the flow is remotely accelerated. |

|                                                                                   |                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Remotely Bridged: The connection is passing through another appliance in the cluster.<br>If the row is colored yellow, then the flow is remotely accelerated.                                           |
|  | Locally/Remotely Bridged: The connection is passing through both this and other appliances in the cluster. (Asymmetric traffic)<br>If the row is colored yellow, then the flow is remotely accelerated. |
|  | Locally Bridged + Locally Accelerated: The connection is passing through this appliance in the cluster and is being accelerated on this appliance.                                                      |
|  | Remotely Bridged + Locally Accelerated: The connection is passing through another appliance in the cluster and is being accelerated on this appliance.                                                  |
|  | Locally/Remotely Bridged + Locally Accelerated: The connection is passing through both this and other appliances in the cluster and is being accelerated on this appliance. (Asymmetric traffic)        |

### 0.11.10 Monitoring subnets

A subnet, a type of network object, can include multiple network subnets and/or multiple IP addresses. The Subnets report shows the top subnets by volume and their average throughput for the selected time period.

When subnets are defined, they can be specified as internal or external to your network. Inbound and outbound traffic for these subnets are reported separately. Inbound and outbound traffic is relative to the subnet, not relative to the Exinda Appliance.

Subnets are not required to be mutually exclusive. Traffic may be reported in more than one subnet. You can optionally show the top three applications for each of the top subnets.

These charts can answer questions such as:

- » What are the top subnets in my network?
- » How much bandwidth does my subnet for the New York branch or for my finance department or for my PBX phones typically consume?
- » Do each of my branches or departments (partitioned by subnet) have the same top applications?

Toggle chart components on and off by clicking the buttons at the top of the report. Note when generating a PDF report of this screen, toggle states are taken into account.



#### NOTE

Average bandwidth is calculated as the total bits observed in the charting interval divided by the number of seconds in that interval. E.g. For a chart with an hour of data, the intervals are five minutes.

Once Network Objects get modified, it can lead to some discrepancies on how Exinda Appliance shows data in the **Subnet** report. Some sections of the reports may become out of sync when drilling down into conversations, URLs, and hosts. It may happen because the report shows data for all the subnets, including the ones defined before and after the change. Which subnets are included in the reports depends on the covered time period and when exactly the change happened. On the other hand, the drilled-down sections always show data based on the newer subnet definitions. Thus, it will include only data collected after the modifications.

The **Subnet** report has a granularity of five minutes, but for some technical limitations, the drilled-down records have a lower granularity of one hour. It leads the drilled-down records to show different data for the selected period.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Subnets**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To configure a subnet for monitoring

Create a network object. See For more information, refer to [Adding network objects](#) (page 178)..

I do not see my subnet data

If Network Object/Subnet statistics collection is disabled, the Subnets report will not include application data for the time period the collection was disabled. For more information, refer to [Adding network objects](#) (page 178).

If the Subnet Report checkbox is not enabled on the definition of the subnet, then the data will not be included in the report. If the data was collected, then enabling the Subnet Report will immediately show the data in the chart.

To change the throughput chart to stacked area charts or line charts

Press the down arrow next to the Stacked Chart button to select Line Chart to switch to the line chart. Conversely, press the down arrow next to the Line Chart button to select Stacked Chart and switch to the stacked area chart. The line chart shows the subnets against the common zero baseline so that the throughput of the subnets can be compared with one another and the pattern of a specific subnet is clearer. You can look for particular patterns such as spikes or flat tops. If your subnets are not defined to be mutually exclusive, displaying the throughput in a line chart with a common zero baseline may make the most sense, as the cumulative values chart will double count some data and may not be meaningful. However, if you have defined your subnets to be mutually exclusive then stacked area charts is an option.

To show the data volume of the subnets as a pie chart

Toggle on the pie chart by clicking the Pie button. Note that if your subnets are not defined to be mutually exclusive, that is, data is captured in more than one subnet, then the pie chart does not hold much meaning.

To show more or fewer subnets in the top subnets chart and the throughput chart

The number of subnets shown are configurable. Note that this configuration applies to all charts on the appliance. For more information, refer to [Monitoring Configuration](#) (page 621).

Should subnet totals match virtual circuit totals when the virtual circuit and subnet are based on the same network object?

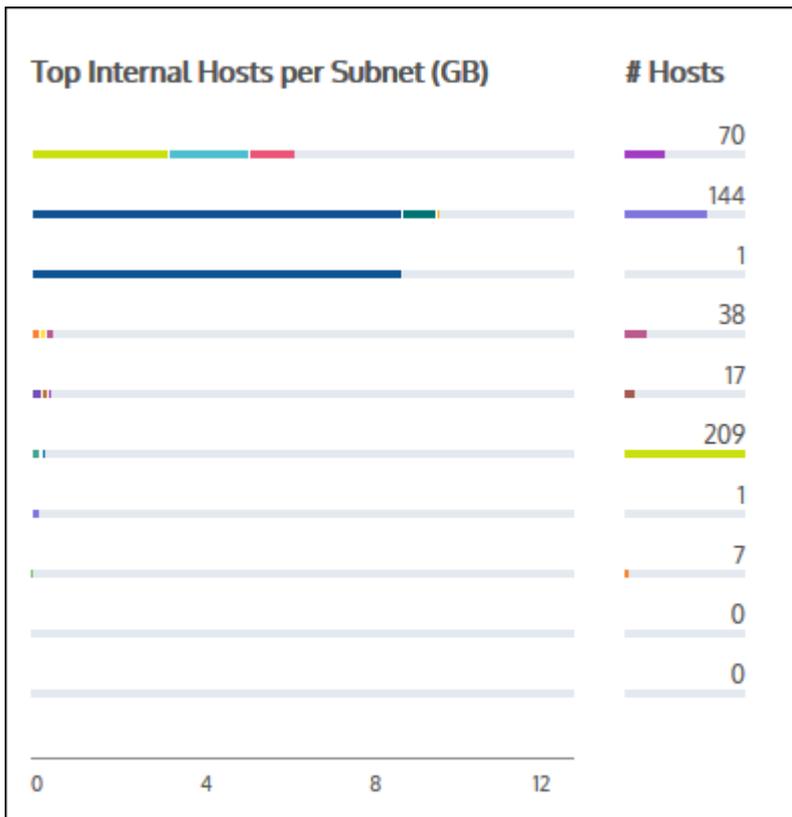
In general, yes. However, there are some cases where the traffic direction is different for subnets versus virtual circuits and so the totals will not match. For more information, refer to [Determining traffic direction and the implications of directional flow on reports](#) (page 295).

Can I view the Top Internal or External Hosts per Subnet in this report?

By default, this report displays the Top Apps per subnet, but you can change the view to Top Internal or Top External Hosts per Subnet. Click the drop-down arrow beside the Top Apps per Subnet button to view these other options. When the display updates, the Top hosts data is mapped to a bar graph. You can brush over any host to view its IP Address and throughput data.

**NOTE**

Toggling the Top Internal Hosts per Subnet and Top External Hosts Per Subnet within this report is available in version 7.0.3 and higher.



How can I drill down in this report?

You can drill into the applications for a specific subnet by clicking on the subnet name in the Top Subnets chart or by clicking on the subnet name in the table below the charts. You can also drill into the hosts, or users, or conversations for a particular subnet by clicking on the **View Users**, **View Conversations**, **View URLs** links in the table. The applications, hosts, users, conversations, or URLs graph will be shown filtered for the specified subnet.

How do I interact with the new time-series & bar chart reports?

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how the charts interact and what the toggle buttons do, see [Understanding How Charts Relate](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand the difference between inbound and outbound traffic, see [Understanding Traffic Direction](#).
- » To understand how many data points are shown for each time period, see [Understanding Traffic Granularity](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

#### Related Topics

- » [Create a Detailed Subnet Activity Report](#)
- » [How Traffic Direction is Determined and its Implications](#)

#### Creating a detailed subnet report

The Detailed Subnet report lists all network activity for Applications, Conversations, Hosts, URLs, and Users on the selected subnets.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Report > PDF Reports**.
5. Click **Add New PDF Report**.
6. In the Reports Selection area, select **Detailed Subnet Reports**.
7. From the subnet list, select a subnet to add to the report and click **Add subnet to report**. Repeat this for each subnet to include in the report.
8. In the Subnets Selected area, select the network traffic to include in the report.
9. In the Report Details area, specify the name of the report, what time period the report should reflect, and an email address where the report can be sent.

#### NOTE

Reports can be sent to multiple recipients by separating email addresses with a comma or semi-colon.

10. Click **Add New Report**.

#### NOTE

The total traffic may not be the same as the subnets summary total as provided in this report. In traffic flow, if multiple subnets are involved, the traffic is counted multiple times in the subnet summary. However, the same is counted only once in the total as provided in this report.

## Determining traffic direction and the implications of directional flow on reports

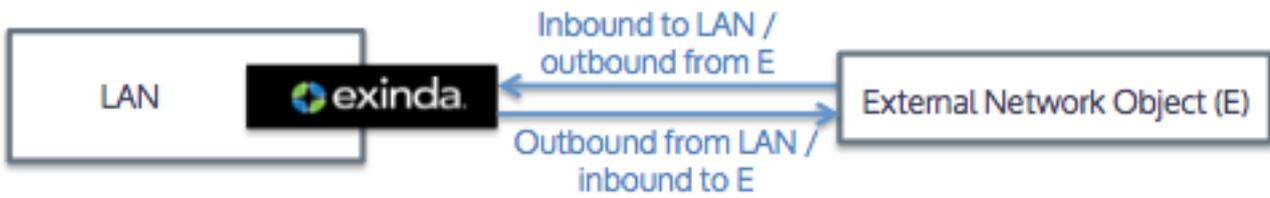
On all reports except the subnet report, traffic direction is determined by the direction of the traffic through the LAN port and WAN port on the appliance. If traffic flows from LAN-side to WAN-side, then the traffic is outbound. If the traffic flows from WAN-side to LAN-side, then the traffic is inbound.

On the subnet report, traffic direction is determined relative to the subnet network object. Traffic originating from the network object is outbound. Traffic destined for the network object is inbound.

Because of these differences, when virtual circuits are based exclusively on a network object, you should generally expect the totals for that network object on the subnet report and the virtual circuit to match. However, there are a few cases where the totals will not match.

When the network object is marked as external, the inbound and outbound traffic are flipped, that is, the inbound virtual circuit traffic will match the outbound subnet traffic. This is because traffic direction for virtual circuits is reported relative to the internal network as determined by the WAN and LAN ports of the appliance, whereas traffic direction for the subnet is reported relative to the location of the network object as determined by the location setting on the network object.

Consider the scenario where the external network object defines a virtual circuit as shown in the figure below. Traffic direction from the LAN to the external network object will be reported as inbound on the Subnet report and outbound on the Virtual Circuit report.



Screenshot 118: Traffic inbound to the External Network Object is outbound from the internal LAN

When a network object is defined on both the LAN and WAN side of the appliance, the Subnet report will double count the traffic, but the virtual circuit report will not .

Consider the scenario where 3 MB of traffic flows from host A to B and both hosts are defined within an internal network object, yet reside on either side of the appliance.

Traffic from host A on the LAN-side to host B on the WAN-side counts on the Subnet report as both 3 MB outbound from the network object as it leaves host A and 3 MB inbound to the network object as it arrives to host B. Traffic from A to B on the virtual circuit report will be counted only as 3MB outbound traffic since the traffic flowed from the LAN-side of the appliance to the WAN-side.



Screenshot 119: Traffic from a network object to itself will be counted as both inbound and outbound traffic on the subnet, but only one direction on the virtual circuit

### 0.11.11 Monitoring virtual circuits

The Virtual Circuits report shows the top virtual circuits by volume and their average throughput. Virtual circuit traffic inbound into your LAN is reported separately from outbound traffic.

You can choose show the entire circuit traffic by adding in a category to represent the remaining virtual circuit traffic on your network; the cumulative stack on the throughput chart represents all the virtual circuit traffic through the appliance. This will help you understand the significance of the top virtual circuits relative to the whole.

In addition to showing the data as a stacked cumulative display, you can choose to display the data as a line chart with a common zero baseline. You can optionally show the top three applications for each of the top virtual circuits.

This report answers questions such as:

- » What are the top virtual circuits in my network?
- » How much bandwidth does my virtual circuit for the New York branch typically consume?
- » I reserved 200 Kbps for this virtual circuit; is that amount of bandwidth really needed or is it too little?
- » Do each of my branches (partitioned by virtual circuit) have the same top applications?
- » What are the top apps per virtual circuit?
- » What are the top internal and external hosts per virtual client?

Drill into the applications for a specific virtual circuit by clicking on the virtual circuit name in the charts or the tables below the charts. The [applications graph](#) will be shown filtered for the specified virtual circuit.



Screenshot 120: Virtual circuit details

**NOTE**

Average bandwidth is calculated as the total bits observed in the charting interval divided by the number of seconds in that interval. E.g. For a chart with an hour of data, the intervals are five minutes.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Virtual Circuits**.

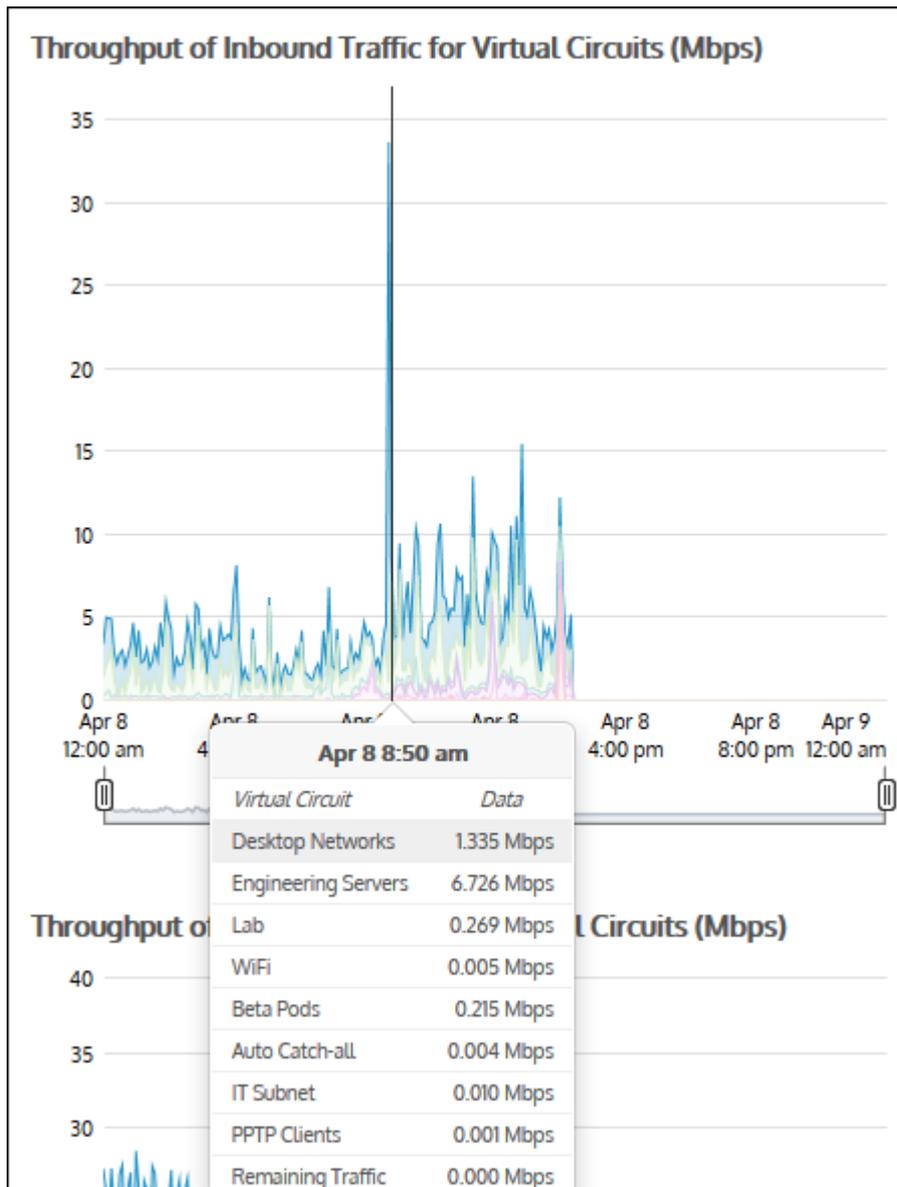
Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To filter the report data

Various components on the screen can be toggled on and off by clicking buttons above the charts. Note that when generating a pdf report of this screen, the toggle states are taken into account. That is, if you had toggled off the outbound charts, they will not be present in the pdf.

- » **Circuit Type:** When you first load the report, only Internet traffic (HTTP) is graphed by default. Click the Internet button to select Auto Catch-all to view all other traffic types graphed instead.
- » **Traffic Type (Inbound/Outbound):** By default, both Inbound and Outbound traffic graphs are displayed. Click either the Inbound or Outbound option to hide the specific graph.
- » **Chart Type:** The chart is initially mapped as a Stacked Area chart, but you can change the format to a Line chart if necessary.
- » **Pie:** Toggles on or off a colour-coded Pie chart to the left of the Top Virtual Circuit chart.

- » **Remaining Traffic:** Append or hide the Remaining Traffic data below the Top Virtual Circuit chart to represent the remaining traffic on your network and so the cumulative stack on the throughput chart represents all the virtual circuits communicating through the appliance.
- » **Data Details:** Toggles on or off the data tables below the time series charts.
- » **Top Apps/Internal Hosts/External Hosts:** Choose one of three display options to view the data per Virtual Circuit: Top Apps per Virtual Circuit (shown by default), Top Internal Hosts per Virtual Circuit, or Top External Hosts per Virtual Circuit.
- » **Mouse Hover:** Hover the mouse pointer over the graph to view data throughput at a given date and time. Refer to Chart Interactions - Drill in & Data brush in [WUI Guided Tour](#) for details.



To determine if a particular virtual circuit is consuming an appropriate amount of bandwidth

View the throughput charts as a line chart with the remaining traffic category toggled off so that each virtual circuit has a common zero baseline. Compare the line for the virtual circuit relative to the bandwidth that was assigned to that virtual circuit.

To configure a virtual circuit

Create a virtual circuit. See [Virtual Circuits](#).

To show more or fewer virtual circuits in the top virtual circuits chart and the throughput chart

The number of virtual circuits shown are configurable. Note that this configuration applies to all charts on the appliance. For more information, refer to [Monitoring Configuration](#) (page 621)..

Should subnet totals match virtual circuit totals when the virtual circuit and subnet are based on the same network object?

In general, yes. However, there are some cases where the traffic direction is different for subnets versus virtual circuits and so the totals will not match. For more information, refer to [Determining traffic direction and the implications of directional flow on reports](#) (page 295).

How do I interact with the new time-series & bar chart reports?

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how the charts interact and what the toggle buttons do, see [Understanding How Charts Relate](#).
- » To understand how to drill into the data to find particular filtered data, see [Drilling into the Data](#).
- » To understand the difference between inbound and outbound traffic, see [Understanding Traffic Direction](#).
- » To understand how many data points are shown for each time period, see [Understanding Traffic Granularity](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### 0.11.12 Monitoring the effects of controls

This section of the chapter deals with the reports available in the **Control** interface. These reports provide information related the effect of applied policies, the discarded packets, and how prioritization has worked to assist your preferred applications.

#### Monitoring the effects of policy controls

The Policies report shows the average throughput and the peak throughput for a circuit, a virtual circuit, or a policy over time. However, only the traffic that fell into a bandwidth shaping (control) policy is reported. The traffic inbound to your LAN is shown separately from the outbound traffic.

If your virtual circuit was defined to provide fair sharing among hosts, the Dynamic Virtual Circuit cart will show the number of active hosts and the number of hosts that exceeded the limit and therefore were not processed by the virtual circuit.

This report answers questions such as:

- » Are my control policies having the effect I need them to?
- » How much traffic is being controlled?
- » Are my policies actively limiting the traffic?
- » How much traffic is hitting my 'Choke P2P' policy?
- » Are hosts being excluded from the dynamic virtual circuit?

#### NOTE

The Policies report only shows traffic for control and bandwidth shaping policies. Accelerated traffic is not included in this report.

## VERSION INFO

Showing the peak throughput is available in version 7.0.1 Update 1 and later.

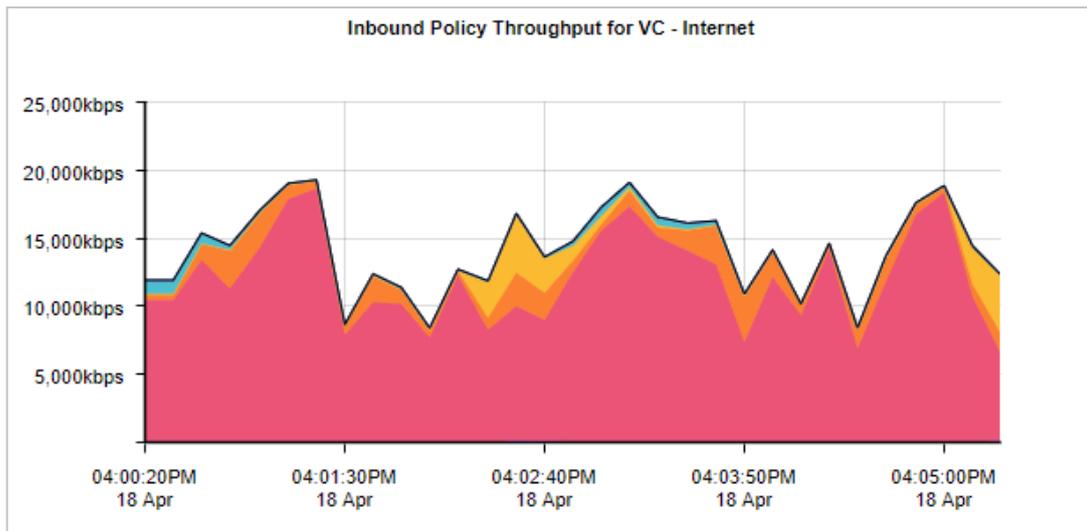
When showing the circuit, the average throughput of the underlying virtual circuits are shown in a stacked graph. The peak throughput for the circuit is shown as a line. Any of the virtual circuits or the peak throughput for the circuit can be removed from the chart by clearing the checkbox in the table below the graph.

The table beneath the graph shows some additional information for the selected time period.

- » Maximum Bandwidth is the amount of bandwidth specified in the virtual circuit. The maximum, average, and current rates are measured as per-second throughput.
- » Average Rate is the average throughput for the time specified in the time range, that is the total volume of data divided by the number of seconds in the specified time period.
- » Max Rate is the maximum throughput or peak bandwidth.

## NOTE

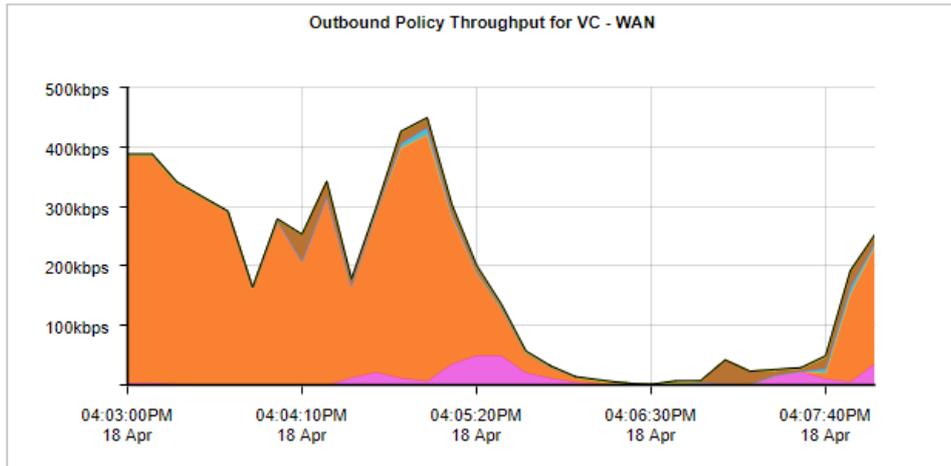
Each time period is broken into 10 second intervals. The throughput for each 10 second sample is the total volume of data seen in that sample divided by 10 seconds to yield a per-second throughput rate. Max Rate is the maximum 10-second throughput. Current Rate is the throughput averaged over the last 10 seconds. Utilization is the percentage when you consider the current rate as compared to the maximum bandwidth.



| Inbound Policies Summary (Internet)                                            |                            |                                                   |  |
|--------------------------------------------------------------------------------|----------------------------|---------------------------------------------------|--|
| Policy Name<br><a href="#">[+] Show Details</a>                                | Avg Rate / Max Rate (kbps) | Current Rate (kbps) / Utilization (%)             |  |
| <input checked="" type="checkbox"/> Peak for Internet                          | 14,277.00 / 19,288.00      | <div style="width: 74%;"></div> 12323.000 / 60.17 |  |
| <input checked="" type="checkbox"/> 10 - P2P - Choke 1%-3%                     | 5.00 / 43.00               | <div style="width: 0%;"></div> 0.000 / 0.00       |  |
| <input checked="" type="checkbox"/> 20 - Recreational - Limit Low 2%-10%       | 27.00 / 180.00             | <div style="width: 15%;"></div> 180.000 / 8.79    |  |
| <input checked="" type="checkbox"/> 30 - Software Updates - Limit Med 3%-50%   | 0.00 / 1.00                | <div style="width: 0%;"></div> 0.000 / 0.00       |  |
| <input checked="" type="checkbox"/> 40 - Voice - Guarantee Critical 15%-100%   | 0.00 / 2.00                | <div style="width: 0%;"></div> 0.000 / 0.00       |  |
| <input checked="" type="checkbox"/> 50 - Thin Client - Guarantee High 10%-100% | 0.00 / 4.00                | <div style="width: 0%;"></div> 0.000 / 0.00       |  |
| <input checked="" type="checkbox"/> 60 - Files - Guarantee Med 8%-100%         | 12,080.00 / 18,635.00      | <div style="width: 65%;"></div> 6349.000 / 31.00  |  |
| <input checked="" type="checkbox"/> 70 - Web - Guarantee High 10%-100%         | 1,297.00 / 3,392.00        | <div style="width: 38%;"></div> 1462.000 / 7.14   |  |
| <input checked="" type="checkbox"/> 80 - Mail - Guarantee Med 8%-100%          | 639.00 / 4,259.00          | <div style="width: 15%;"></div> 4242.000 / 20.71  |  |
| <input checked="" type="checkbox"/> 200 - ALL - Guarantee Low 5%-100%          | 223.00 / 982.00            | <div style="width: 23%;"></div> 87.000 / 0.42     |  |
| <input checked="" type="checkbox"/> Auto Catch-all (N/A)                       | 0.00 / 0.00                | <div style="width: 0%;"></div> 0.000 / 0.00       |  |

Screenshot 121: Control graph for a specified circuit

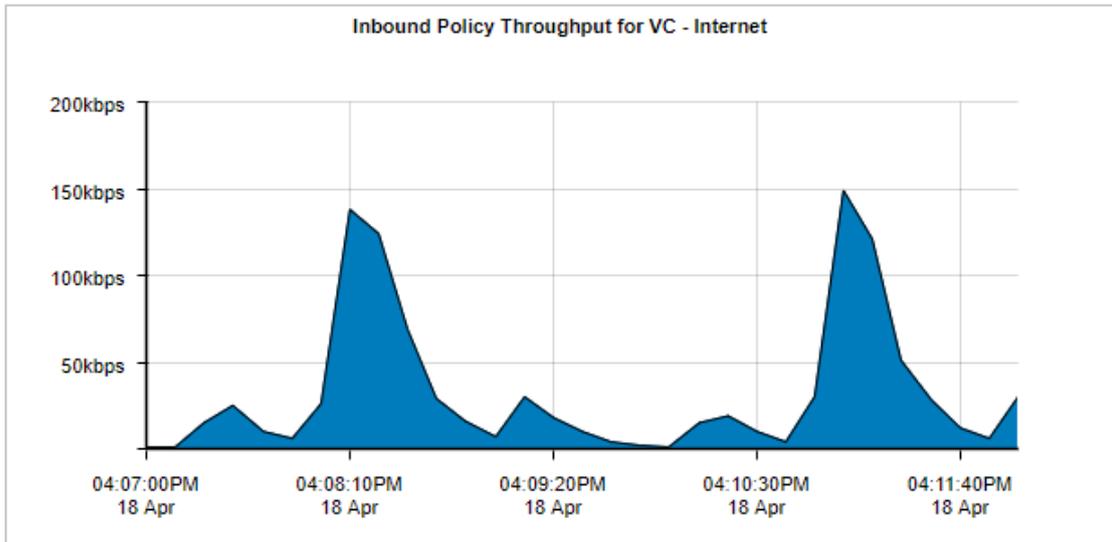
When showing a particular virtual circuit, the average throughput of the underlying top policies are shown in a stacked graph. The peak throughput for the virtual circuit is shown as a line.



| Outbound Policies Summary (WAN)                                                                |                            |                                       |  |                |
|------------------------------------------------------------------------------------------------|----------------------------|---------------------------------------|--|----------------|
| Policy Name<br><a href="#">[+] Show Details</a>                                                | Avg Rate / Max Rate (kbps) | Current Rate (kbps) / Utilization (%) |  |                |
| <input checked="" type="checkbox"/> Peak for WAN                                               | 170.00 / 449.00            | <input type="text"/>                  |  | 253.000 / 2.47 |
| <input checked="" type="checkbox"/> 10 - P2P - Choke 1%-3%                                     | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 20 - Recreational - Limit Low 2%-10%                       | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 30 - Software Updates - Guarantee Low 5%-100% - Accelerate | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 40 - Voice - Guarantee Critical 15%-100%                   | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 50 - Interactive and Secure - Guarantee High 10%-100%      | 11.00 / 50.00              | <input type="text"/>                  |  | 37.000 / 0.36  |
| <input checked="" type="checkbox"/> 60 - Thin Client - Guarantee High 10%-100%                 | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 70 - Files - Guarantee Med 8%-100% - Accelerate            | 146.00 / 414.00            | <input type="text"/>                  |  | 197.000 / 1.92 |
| <input checked="" type="checkbox"/> 80 - Web - Guarantee Med 8%-100% - Accelerate              | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 90 - Mail - Guarantee Low 5%-100% - Accelerate             | 1.00 / 12.00               | <input type="text"/>                  |  | 5.000 / 0.05   |
| <input checked="" type="checkbox"/> 100 - Database - Guarantee Med 8%-100% - Accelerate        | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |
| <input checked="" type="checkbox"/> 200 - ALL - Guarantee Low 5%-100%                          | 9.00 / 44.00               | <input type="text"/>                  |  | 13.000 / 0.13  |
| <input checked="" type="checkbox"/> Auto Catch-all (N/A)                                       | 0.00 / 0.00                | <input type="text"/>                  |  | 0.000 / 0.00   |

Screenshot 122: Control graph for a specified virtual circuit

When showing a particular policy, the average throughput for the selected policy is shown in the throughput chart. The peak throughput for the virtual circuit is shown as a line.

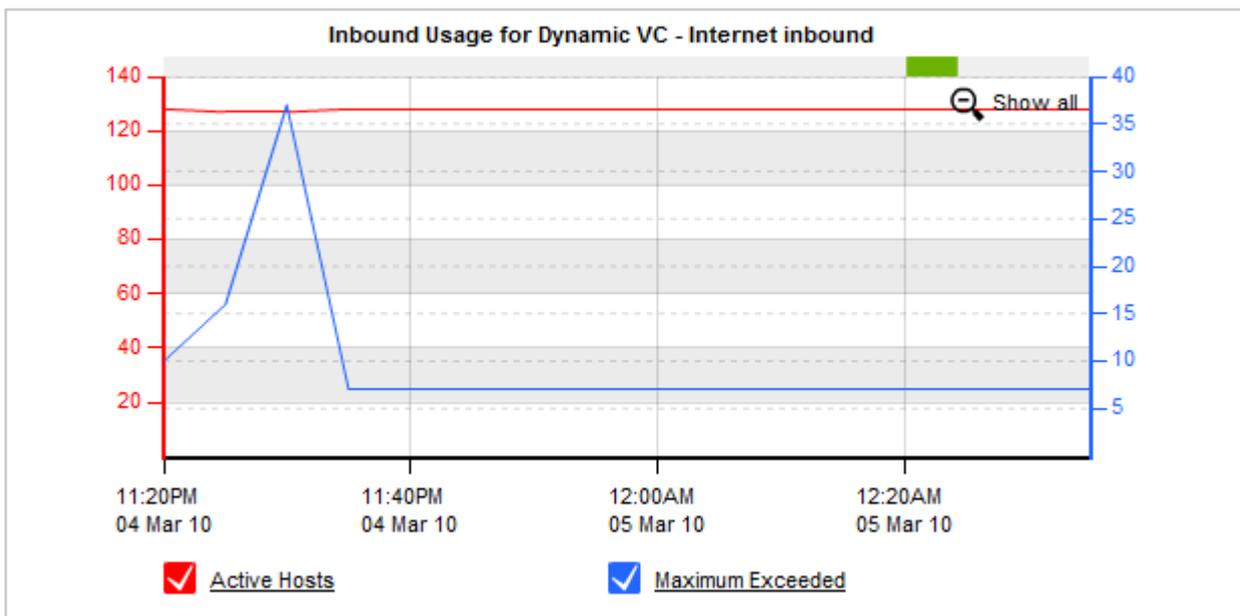


| Inbound Policies Summary (Internet)                                          |                            |                                       |  |
|------------------------------------------------------------------------------|----------------------------|---------------------------------------|--|
| Policy Name<br><a href="#">[+] Show Details</a>                              | Avg Rate / Max Rate (kbps) | Current Rate (kbps) / Utilization (%) |  |
| <input checked="" type="checkbox"/> Peak for Recreational - Limit Low 2%-10% | 33.00 / 149.00             | <input type="text"/> 30.000 / 1.46    |  |
| <input checked="" type="checkbox"/> 20 - Recreational - Limit Low 2%-10%     | 33.00 / 149.00             | <input type="text"/> 30.000 / 1.46    |  |

Screenshot 123: Control graph for a specified policy

Virtual circuits can be defined to provide fair share per host shaping for each host that falls into the virtual circuit. This dynamic virtual circuit dynamically divides the amount of bandwidth available for the number of active hosts. The policies page will show the number of active hosts and the number of hosts that have exceeded the hosts limit for a selected dynamic virtual circuit. The hosts that exceeded the hosts limit were not handled by this virtual circuit and would have been captured in another virtual circuit.

If the virtual circuit selected is a dynamic virtual circuit, then the following graph will appear above the throughput graph.



Screenshot 124: Control graph for a dynamic virtual circuit showing the number of hosts included (Active Hosts) and excluded (Maximum Exceeded) from the virtual circuit

The 'Active Hosts' line shows the number of hosts that were processed by this dynamic virtual circuit over time. The 'Maximum Exceeded' line shows the number of hosts that have exceeded the hosts limit for this dynamic virtual circuit.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Control > Policies**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

To show the control throughput for a circuit

Select the desired circuit. Set the desired virtual circuit to ALL. The average throughput for the circuit is shown in the throughput chart.

To show the control throughput for a particular virtual circuit

Select the desired circuit. Set the desired virtual circuit. The average throughput for the top policies is shown stacked in the throughput chart.

To show the control throughput for a particular policy

Select the desired circuit. Select the desired virtual circuit. Select the desired policy. The average throughput for the selected policy is shown.

How is the average bandwidth calculated?

The average bandwidth is calculated as the total bits observed in the charting interval and dividing by the number of seconds in that interval. E.g. For a chart with an hour of data, the intervals are five minutes. Thus, for each five-minute interval, the number of observed bits is divided by 300 seconds.

How is the peak bandwidth calculated?

The peak bandwidth is calculated as the maximum throughput observed in the charting interval, where the considered throughput is the average over a 10 second interval. E.g. For a chart with an hour of data, the intervals are five minutes. For a five minute interval, there are 30 possible 10-second intervals considered. The average bits per second are calculated for each 10-second interval. The maximum of these 30 possible calculations is considered the peak throughput for that five-minute interval.

To determine if any of the policies appear to be limited

Select the desired circuit, the desired virtual circuit, and the desired policy. The average throughput for the selected policy is shown in the throughput chart. If the throughput chart has any raised flat tops, this may represent that the policy is actively limiting the traffic.

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## Monitoring discarded network packets

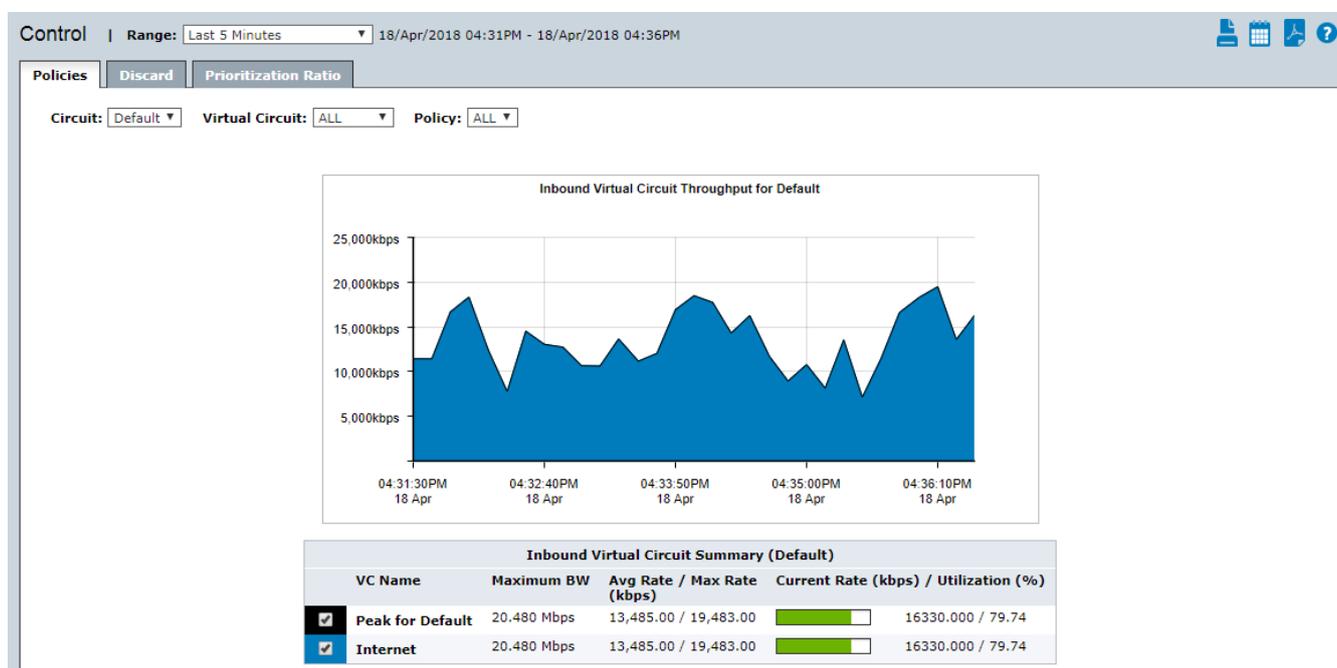
The Discard report shows the number of discarded packets over time for the specified span of time as a result of Optimizer policies configured to block traffic.

This report answers questions such as:

- » How many attempts to access blocked apps occurred?
- » When are people trying to access the blocked apps?

The graph shows the number of packets discarded over time. The table below the graph shows the total number of discarded packets over the selected time period.

Select a virtual circuit from the **Virtual Circuit** selector at the top of the graph to view the packet discards for it. Filter by policy by selecting a policy in the **Policy** selector. Only virtual circuits and policies that have reported packet discards will appear in the selectors at the point in time where the packets were discarded, which gives an indication where more traffic would have appeared on the network if that traffic had not been blocked.



Screenshot 125: The Discard tab on the Control report discard packets over time.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Control > Discard**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).

- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

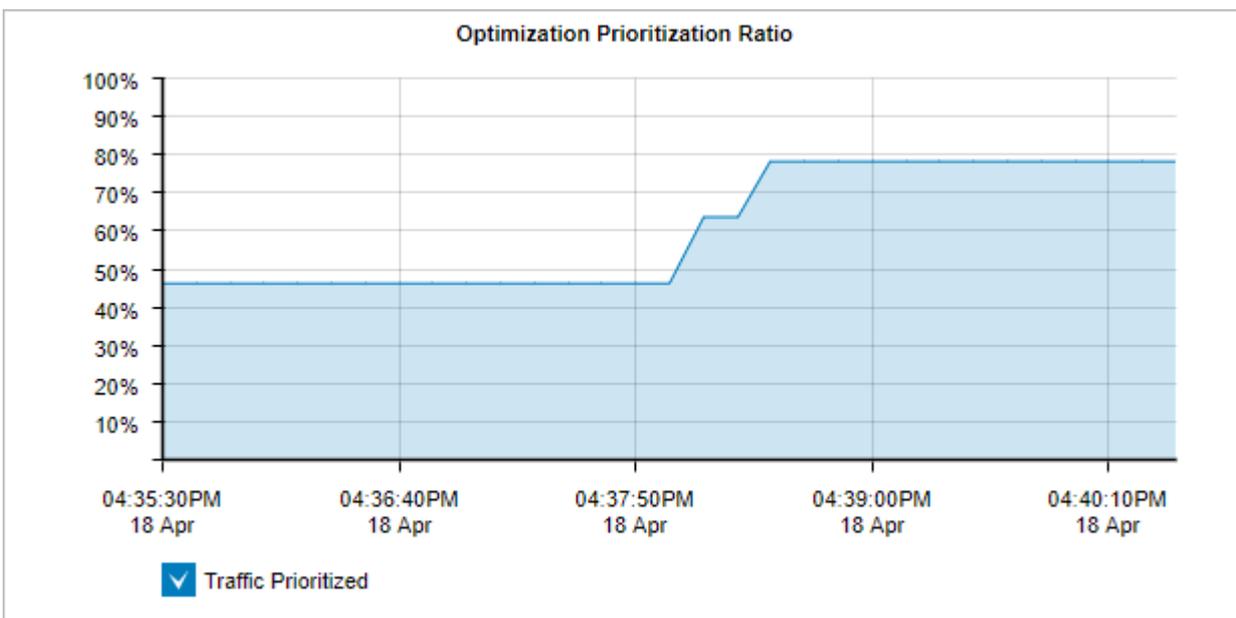
## Monitoring prioritization of applications

The Prioritization report shows how often applications were prioritized (also referred to as re-ordering or re-queuing).

This report answers questions such as:

- » I've created a number of policies with different priorities; how much of my traffic is re-ordered to ensure that my important traffic has priority on my network?

A high percentage means that the system is prioritizing more often to ensure performance of your applications. A high percentage also means that by turning off the optimizer there is a higher probability critical applications will suffer.



### EXAMPLE

A ratio of 40% means 40% of the packets on your network were re-ordered. That means that non critical data was queued so business critical data could jump the queue and be delivered according to business requirements.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Control > Prioritization Ratio**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How is the prioritization ratio calculated?

Prioritization Ratio =  $100 \times \text{Number of Packets Re-ordered} / \text{Number of Total Packets}$

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### 0.11.13 Monitoring optimization reports

Learn about the available optimization reports. These reports let you see the bandwidth savings achieved through using the WAN memory and Edge Cache configurations on your Exinda Appliances.

#### Monitoring traffic reduction

The Optimization Reduction report shows the amount of traffic reduction achieved due to WAN memory techniques. It also shows the amount of reduction per Exinda Appliance peer and per application.

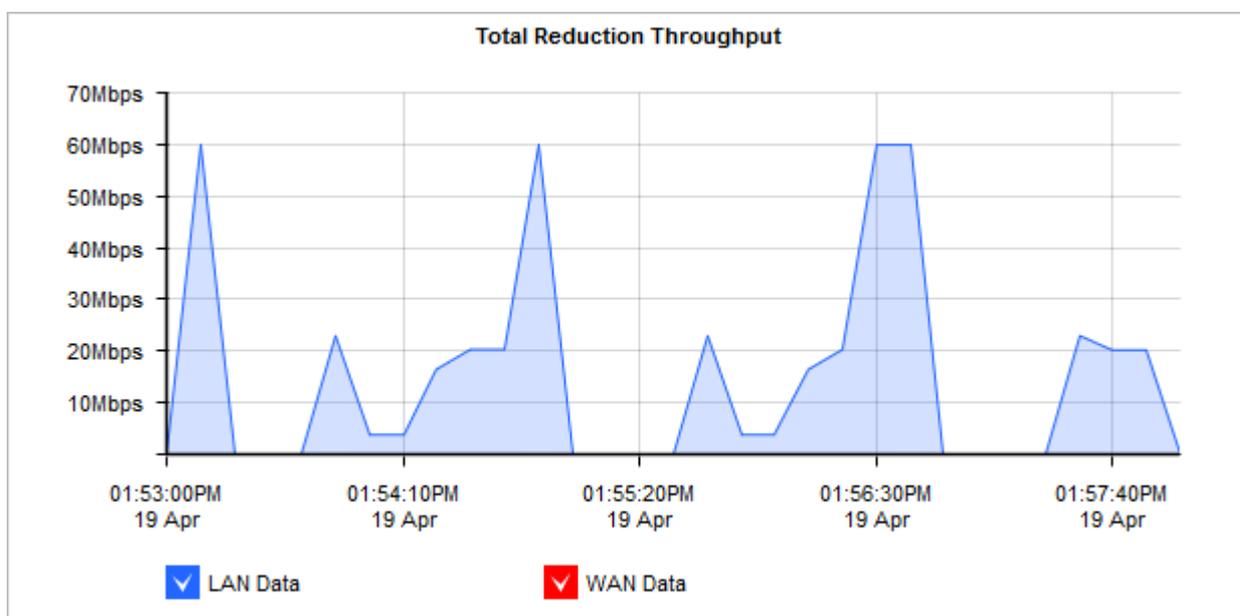
This report answers questions such as:

- » Am I achieving the amount of traffic reduction overall that I am expecting?
- » Is a specific application achieving the amount of traffic reduction that was expected?

The chart shows traffic reduction over time as either throughput or percentage reduction. You can also choose the traffic direction to report on - inbound, outbound, or bi-directional.

#### NOTE

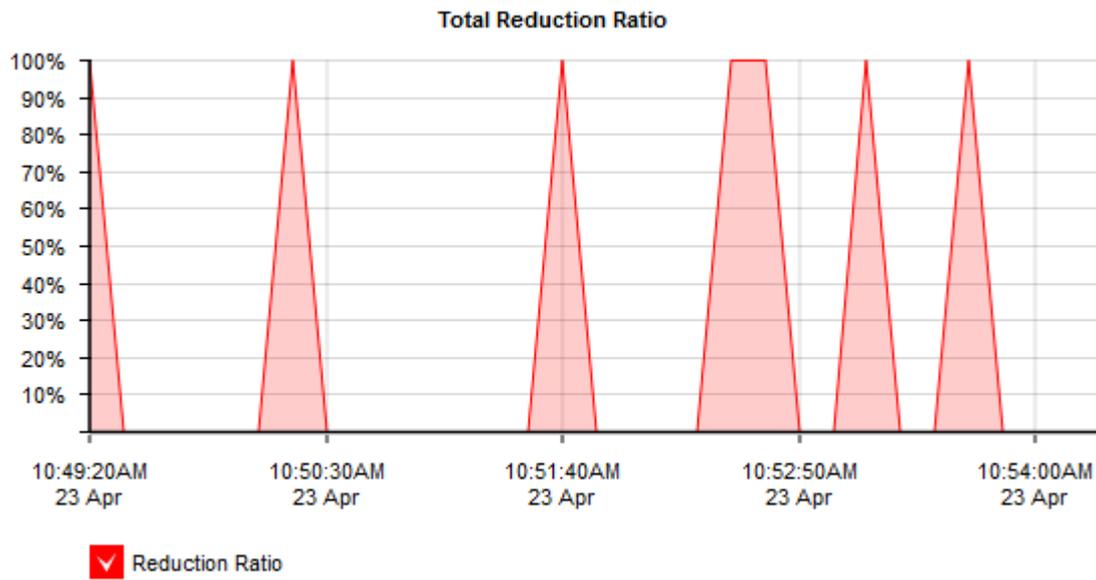
When reduction statistics are displayed as throughput, there is one time series plotted for LAN throughput and one for WAN throughput.



Screenshot 126: The Total Reduction Throughput graph displays reduction statistics over time.

#### NOTE

When reduction statistics are displayed as percentage reduction, one line graph represents the reduction ratio, which compares the amount of data transferred after or before Exinda:



Screenshot 127: The Total Reduction Ratio graph displays reduction ratios over time.

The table below the graphs show reduction statistics broken down by a remote Exinda Appliance peer and by application.

| Reduction Statistics by Peer |               |                  |                     |  |              |
|------------------------------|---------------|------------------|---------------------|--|--------------|
| Peer                         | LAN Data (MB) | WAN Data (MB)    | Reduction Ratio (%) |  |              |
| 1                            | will          | 374.01           | 286.44              |  | 23.41        |
| 2                            | war           | 2,762.56         | 872.18              |  | 68.43        |
| 3                            | bhl           | 230.23           | 171.05              |  | 25.70        |
| 4                            | wbri          | 374.16           | 222.65              |  | 40.49        |
| 5                            | man           | 76,821.56        | 54,609.32           |  | 28.91        |
| 6                            | tops          | 247.85           | 137.24              |  | 44.63        |
| 7                            | bed           | 186.72           | 136.68              |  | 26.80        |
| 8                            | bos           | 44.7             | 32.6                |  | 27.07        |
| 9                            | hol           | 420.65           | 189.31              |  | 55.00        |
| 10                           | rh            | 3,764.34         | 1,888.74            |  | 49.83        |
| 11                           | wor           | 106.46           | 89.48               |  | 15.95        |
| 12                           | wilt          | 906.41           | 478.35              |  | 47.23        |
| 13                           | wes           | 2,591.56         | 843.86              |  | 67.44        |
| <b>Total</b>                 |               | <b>88,831.21</b> | <b>59,957.9</b>     |  | <b>32.50</b> |

| Reduction Statistics by Application |                  |               |                     |  |       |
|-------------------------------------|------------------|---------------|---------------------|--|-------|
| Application                         | LAN Data (MB)    | WAN Data (MB) | Reduction Ratio (%) |  |       |
| 1                                   | Discovered Ports | 0.11          | 0.08                |  | 27.27 |
| 2                                   | URL              | 93.68         | 28.3                |  | 69.79 |
| 3                                   | URL 1            | 1.1           | 0.67                |  | 39.09 |
| 4                                   | http file        | 24.07         | 17.67               |  | 26.59 |
| 5                                   | EMC Replication  | 75,027.27     | 53,135.1            |  | 29.18 |
| 6                                   | ZCM              | 0.01          | 0.01                |  | 0.00  |
| 7                                   | HTTP             | 6,777.73      | 1,984.65            |  | 70.72 |
| 8                                   | CIFS             | 4,978.02      | 3,399.93            |  | 31.70 |
| 9                                   | LotusNotes       | 1,681.82      | 1,360.8             |  | 19.09 |
| 10                                  | MS-SQL           | 247.01        | 30.46               |  | 87.67 |
| 11                                  | Oracle           | 0.02          | 0.02                |  | 0.00  |
| 12                                  | NFS              | 0.09          | 0.05                |  | 44.44 |

Screenshot 128: The Reduction Statistics charts compares peer and application reduction.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > Optimization > Reduction**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How is the reduction ratio calculated?

Reduction Ratio = (Data Transfer Size Before Exinda - Data Transfer Size After

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### Monitoring edge cache reduction

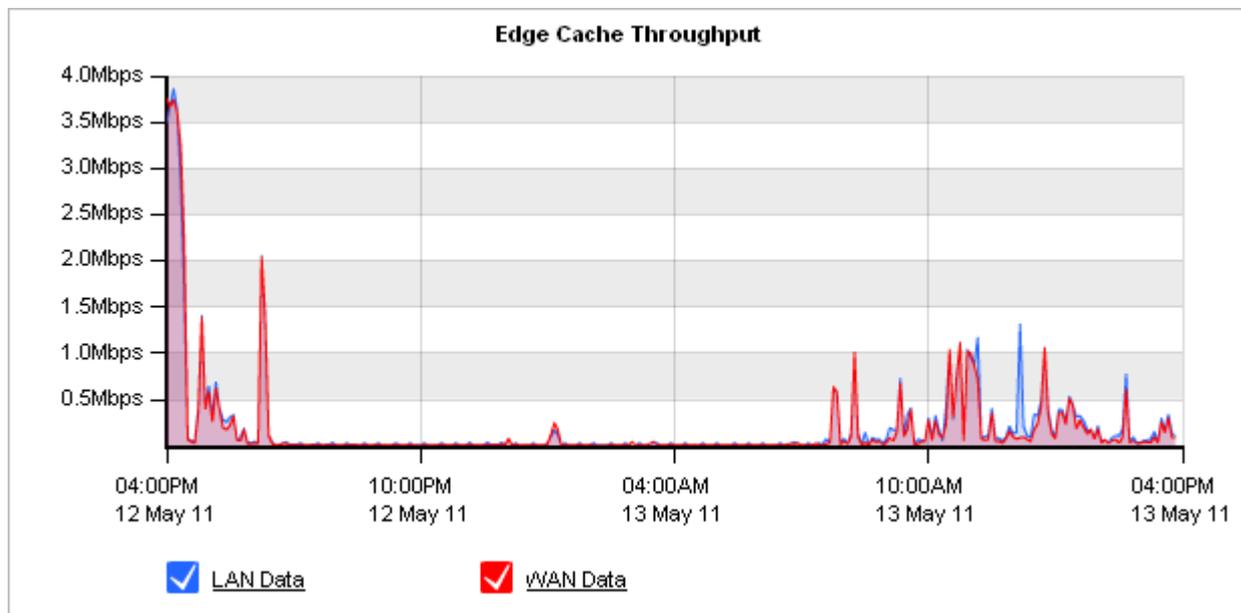
The Edge Cache report shows the reduction in the amount of traffic achieved due to caching techniques and the number of requests and the number of hits from the cache.

This report answers questions such as:

- » How much traffic reduction am I getting due to Edge Cache?
- » How cacheable is the network data and how frequently is the cache being accessed?

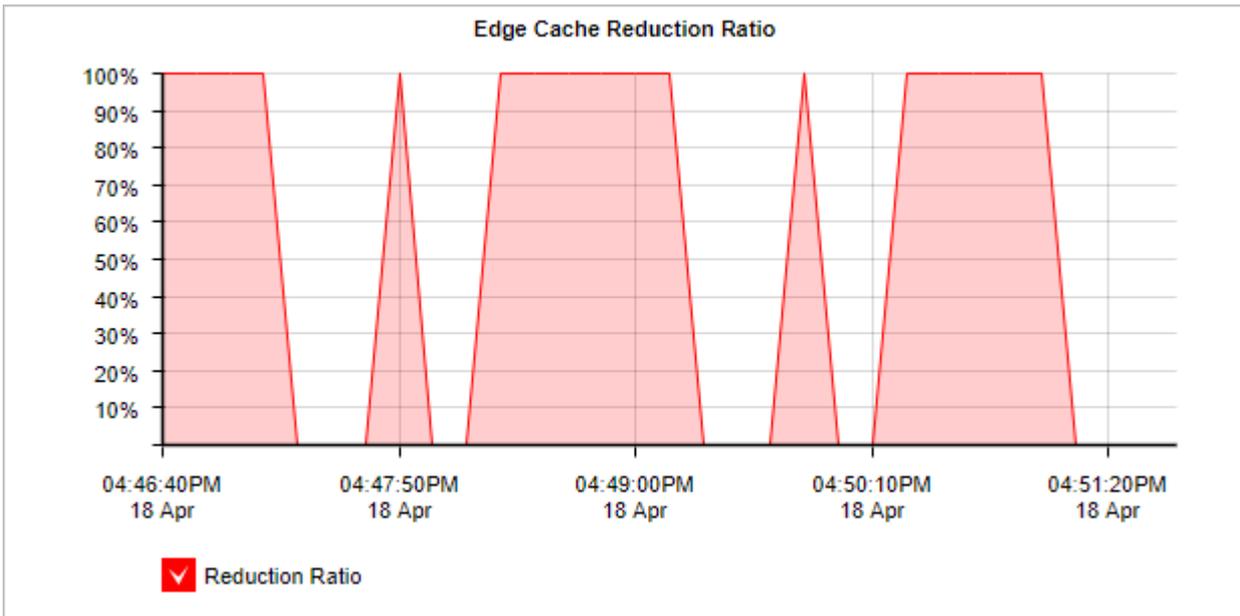
The chart shows the traffic reduction over time as either throughput or percentage reduction. The reported LAN throughput is the amount of traffic that was served to the client whether it was served from Edge Cache or not. This represents all traffic that was handled by Edge Cache policies. The reported WAN throughput is the amount of traffic that was not available in Edge Cache and needed to be retrieved from the application server. Therefore, the difference between WAN and LAN is the amount of traffic that could be served from Edge Cache.

When reduction statistics are displayed as throughput, there is one time series plotted for LAN throughput and one for WAN throughput. You should expect the LAN throughput to be higher than WAN throughput. When the LAN throughput is greater than the WAN throughput, the more traffic was served up from Edge Cache than was needed to be retrieved from the WAN.



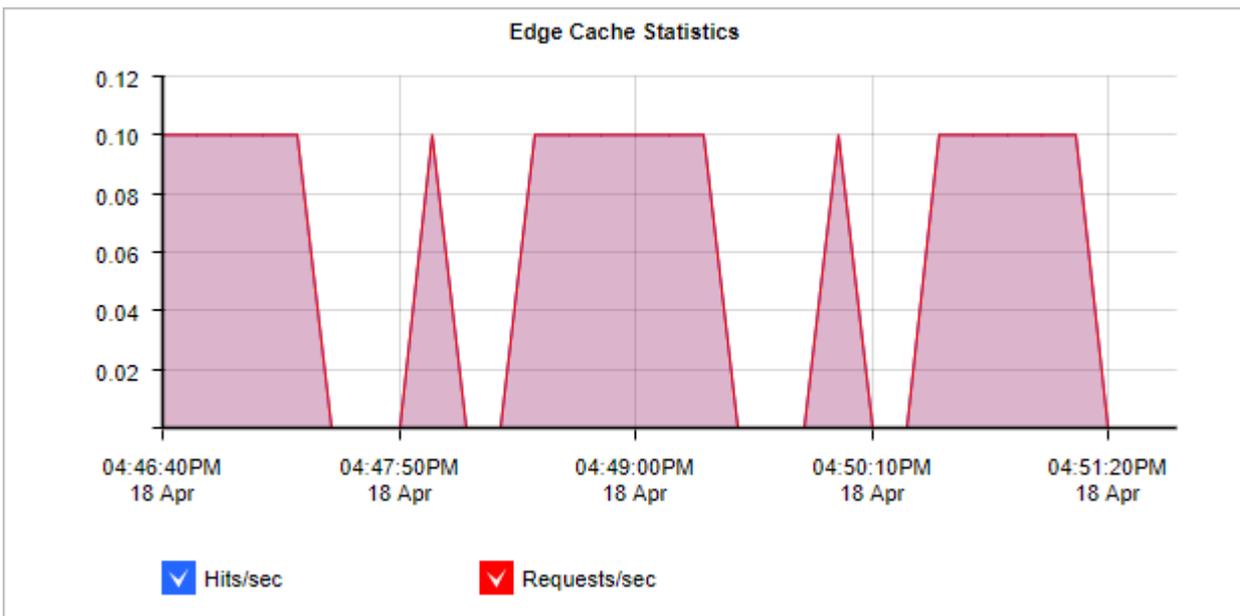
Screenshot 129: The Edge Cache Throughput graph.

When displayed as percentage reduction, it displays one line graph to represent the percentage of data transferred that was sent from Edge Cache instead of from the application server.



Screenshot 130: The Edge Cache Reduction Ration graph.

The Edge Cache Statistics chart display the number of requests per second and the number of hits per second. A request occurs when Edge Cache is checked for particular data. A hit occurs when a request is satisfied by an object already stored in the Edge Cache.



Screenshot 131: The Edge Cache Statistics graph.

The table shows a summary of Edge Cache reduction for the selected time period.

| LAN (MB) | WAN (MB) | Reduction Ratio (%)                                                                      | Requests | Hits   | Hit Ratio (%)                                                                              |
|----------|----------|------------------------------------------------------------------------------------------|----------|--------|--------------------------------------------------------------------------------------------|
| 2138.61  | 1930.90  | <div style="width: 9.71%; background-color: green; border: 1px solid black;"></div> 9.71 | 204030   | 120030 | <div style="width: 58.83%; background-color: green; border: 1px solid black;"></div> 58.83 |

Where do I find this report?

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.

#### 4. Go to **Monitor > Optimization > Edge Cache**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How to change the Edge Cache Throughput chart to a Edge Cache Reduction Ratio chart?

Select the desired type of chart from the Edge Cache Graph Type selector below the chart.

How is the Edge Cache Reduction Ratio calculated?

Reduction Ratio = (Data Transfer Size Before Exinda - Data Transfer Size After Exinda) / Data Transfer Size Before Exinda

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### 0.11.14 Monitoring Exinda Appliance system performance

Learn about the reports that provide feedback on the performance of your Exinda Appliance. The reports cover aspects of operational performance like number of concurrent connections, CPU utilization, CPU temperature, memory usage, disk IO and swap space usage.

#### Monitoring connections to an Exinda Appliance

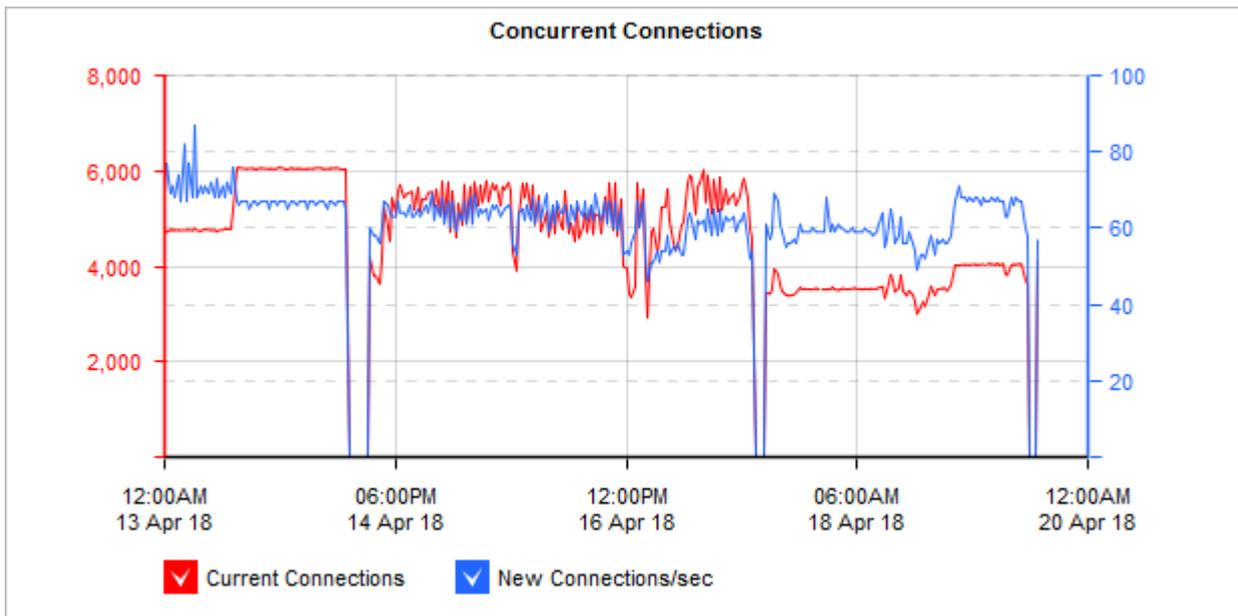
The Connections report shows the number of concurrent connections as well as the connection establishment rate over time for the selected time period.

This report answers questions such as:

- » Is there an unusual number of connections or is the connection rate particularly high?
- » Could I be experiencing some form of denial of service attack or network problem?"

#### NOTE

Systems reporting unusually high spikes in the number of connections or rate of connections may be experiencing a denial of service attack or network problem.



Screenshot 132: The Concurrent Connections graph displays connection statistics over time.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > Connections**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### Monitoring accelerated connections

The Accelerated Connections report shows the number of concurrent accelerated connections as well as the connection establishment rate over time for the selected time period. It also shows the number of connections for each application acceleration type (SSL, SMB1, SMB2, NCP). This chart can answer questions such as:

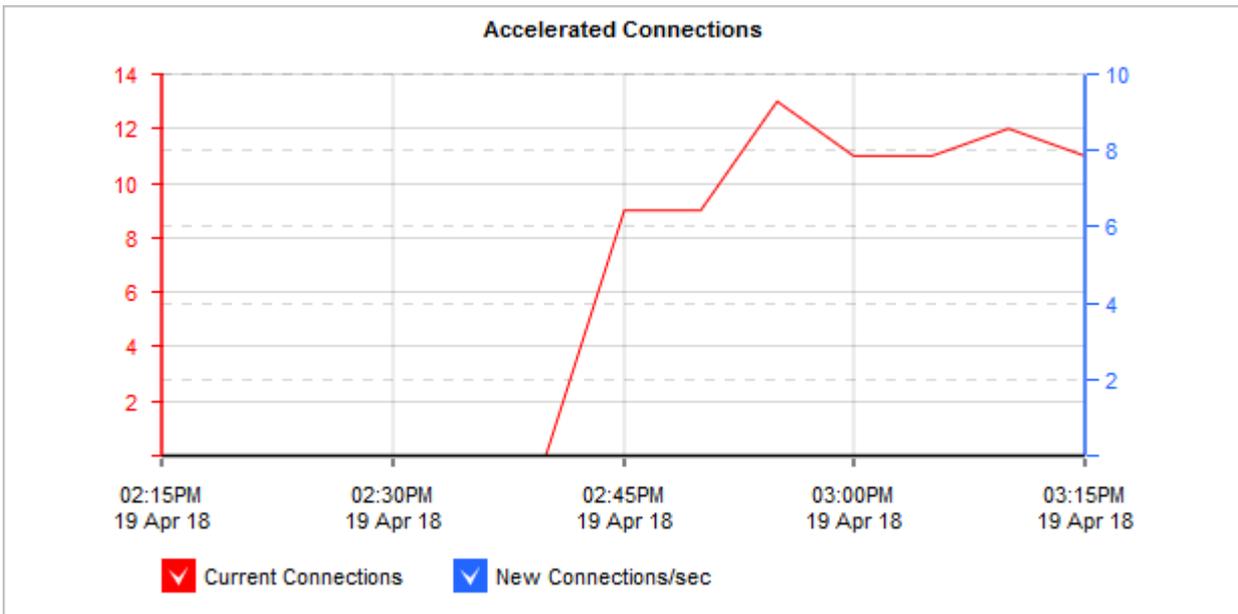
- » Is there an unusual number of accelerated connections or is the connection rate particularly high or low?
- » Is my traffic being accelerated as I expect?
- » Am I close to or have I exceeded my licensed maximum number of accelerated connections?"

NOTE

Connections over the licensed limit pass through the appliance without acceleration. If you are concerned about exceeding your licensed maximum, you can set an alert so that you will be emailed or so that the appliance will send an SNMP trap when the number of connections exceed your licensed limit.

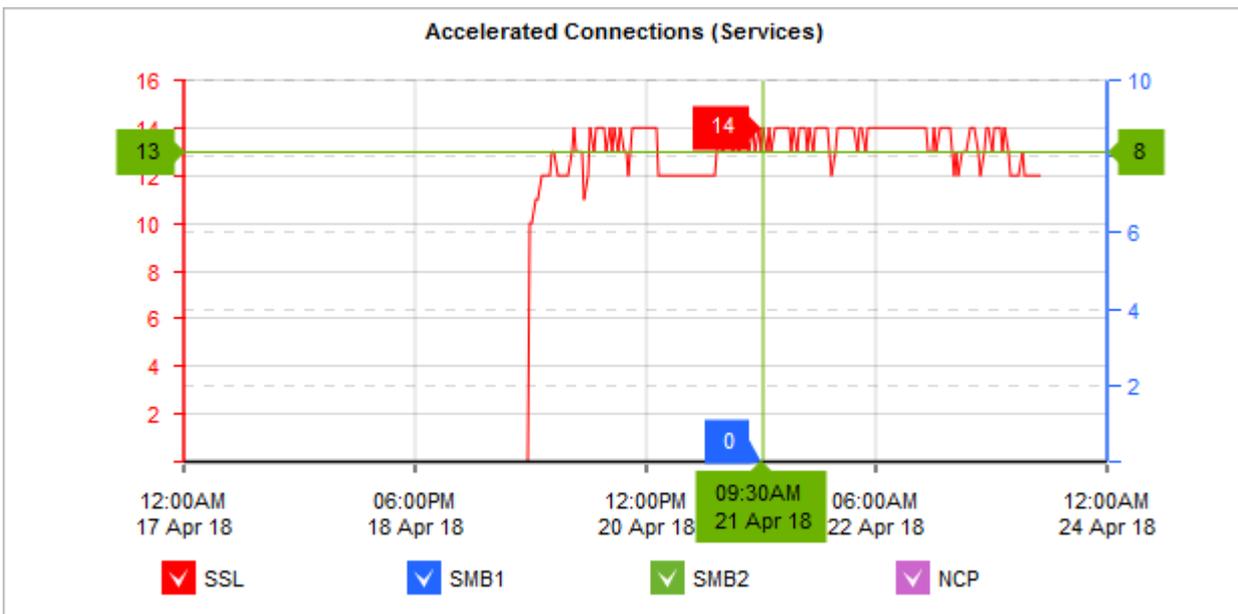
To find the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > Accelerated Connections**.



Screenshot 133: The Accelerated Connections graph displays accelerated connections over time.

The second chart shows accelerated connections for each type of accelerated traffic: SSL, SMB1, SMB2, and NCP.



Screenshot 134: The Accelerated Connections graph shows accelerated traffic by type over time.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

### Finding the licensed accelerated connections limit

To view the details of your license:

1. Go to **Configuration > System > Setup > License**.
2. The **Max AA Connections** field in the current system license status reports your licensed limit.

### Setting an alert or sending an SNMP trap when the number of accelerated connections exceeds licensed limit

1. Go to **Configuration > System > Setup > Alerts**.
2. Ensure the appropriate check boxes are selected for **Max Accelerated Connections Exceeded**.

#### NOTE

The appliance must already be configured for email or SNMP.

#### Related Topics

[Alerts](#)

[Email Configuration](#)

#### Using interactive time graphs

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### Monitoring Exinda Appliance CPU usage

The CPU Usage report shows how hard the CPU is working over time.

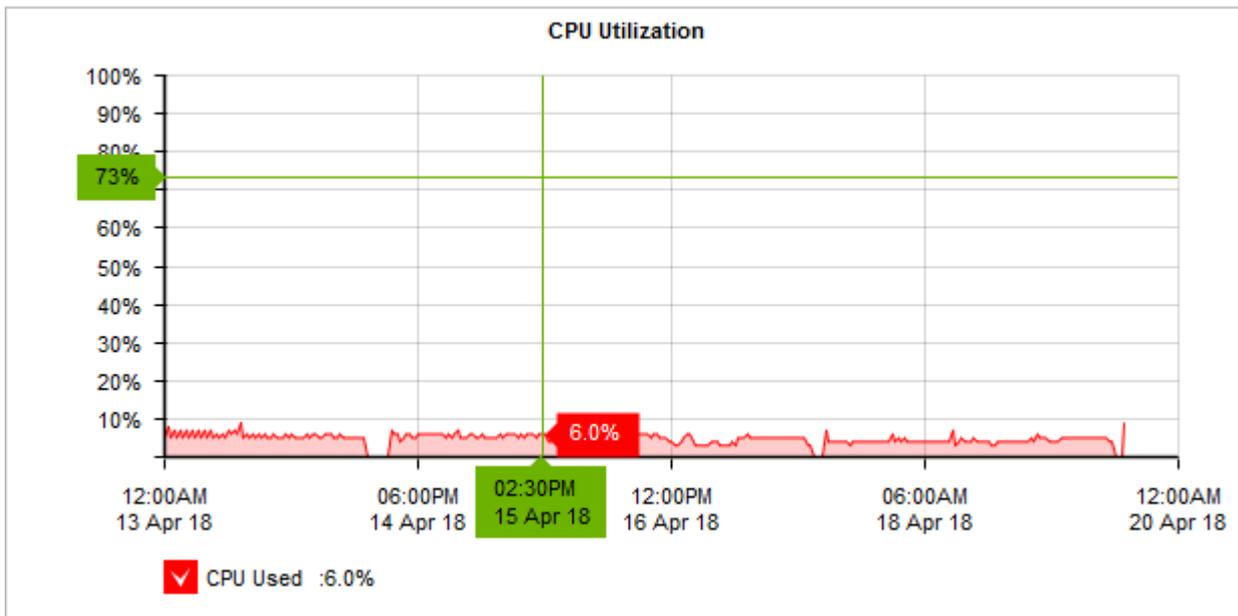
This report answers questions such as:

- » Are some of the other issues I'm seeing with my traffic due to overworking the appliance?
- » I see the appliance's CPU is highly utilized. What traffic issues could cause this to be happening?

High CPU usage may be due to a variety of processing intensive traffic characteristics:

- » The number of new connections per second is high
- » The number of accelerated connections is high
- » The appliance is encountering more accelerated traffic than it can handle. If this is the case, latency can be introduced when the appliance queues the packets for acceleration and cannot process them fast enough. Using virtual circuits, you can limit the amount of accelerated traffic to process.
- » The appliance is encountering more analysis-intensive traffic than it can handle. For instance, VoIP traffic is CPU intensive due to the processing required to compute the metrics such as rFactor, MOS, jitter, etc.

To diagnose a CPU usage problem, for each period where the CPU usage is high, compare with the Connections report, the Accelerated Connections report, the Reduction report, and the VoIP Solution report.



Screenshot 135: The CPU utilization graph shows how hard the Exinda Appliance works over time.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > CPU Usage**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

Where do I find the other reports for diagnosing the reason for high CPU usage?

- » The report for number of new connections can be found at **Monitor > System > Connections**. For more information, refer to [Monitoring connections to an Exinda Appliance](#) (page 312).
- » The report for the number of accelerated connections can be found at **Monitor > System > Accelerated Connections**. For more information, refer to [Monitoring accelerated connections](#) (page 313).
- » The report for the amount of accelerated traffic can be found at **Monitor > Optimization > Reduction**. For more information, refer to [Monitoring traffic reduction](#) (page 307).
- » The report for VoIP traffic can be found in the Solution Center (**Solution Center > Show Solution Center**). For more information, refer to [Using the Application Performance Monitor VoIP report](#) (page 335).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## Monitoring Exinda Appliance CPU temperature

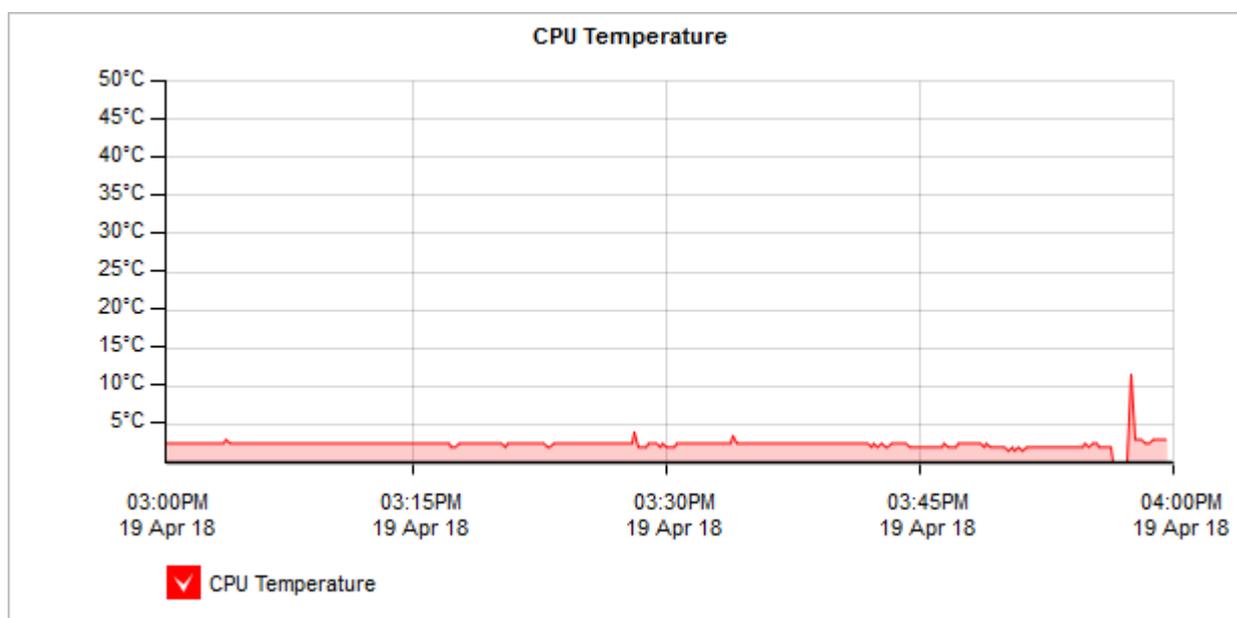
The CPU Temperature report shows the temperature in degrees Celsius of the appliance CPU over time for the selected time period.

This report answers questions such as:

- » Are some of the other issues I'm seeing with my traffic due to overworking the appliance?
- » I see the appliance's CPU temperature is high. Is it due to high CPU usage or is the ambient temperature around the Exinda Appliance too warm?

You should expect the CPU temperature to be considerably lower than 80 degrees Celsius, usually between 35-50 degrees. Systems running at very high temperatures may be experiencing a problem and system performance may be affected. Once the temperature gets too high (80-90 degrees) the appliance will throttle its processing speed to reduce heat emissions.

See the **CPU Usage** report to see if the temperature correlates with the processing activity on the appliance.



Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > CPU Temperature**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

Where do I find the CPU Usage report?

The report for CPU usage can be found at **Monitor > System > CPU Usage**. For more information, refer to [Monitoring Exinda Appliance CPU usage](#) (page 315).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using](#)

## Interactive Time Graphs.

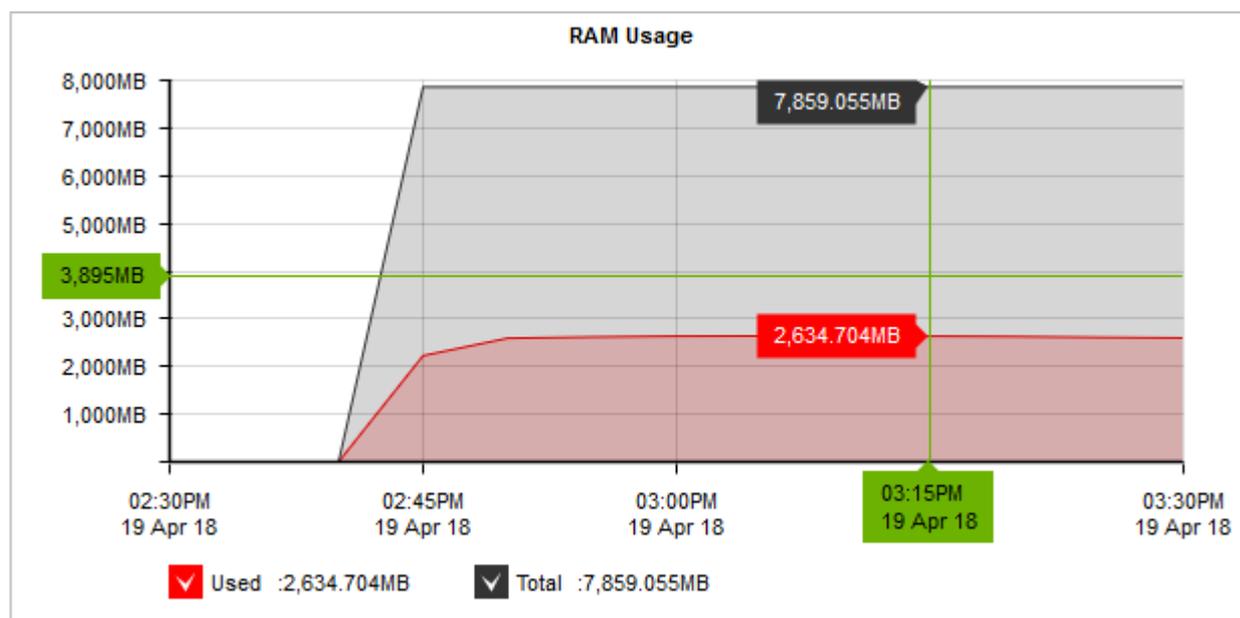
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## Monitoring Exinda Appliance RAM usage

The RAM Usage report shows how much memory the appliance is using relative to the amount of memory available for the selected time period.

This report answers questions such as:

- » Could the performance of my appliance be affected by insufficient RAM?



Screenshot 136: The RAM Usage chart displays memory consumption over time.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > RAM Usage**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

## Monitoring Exinda Appliance Disk IO

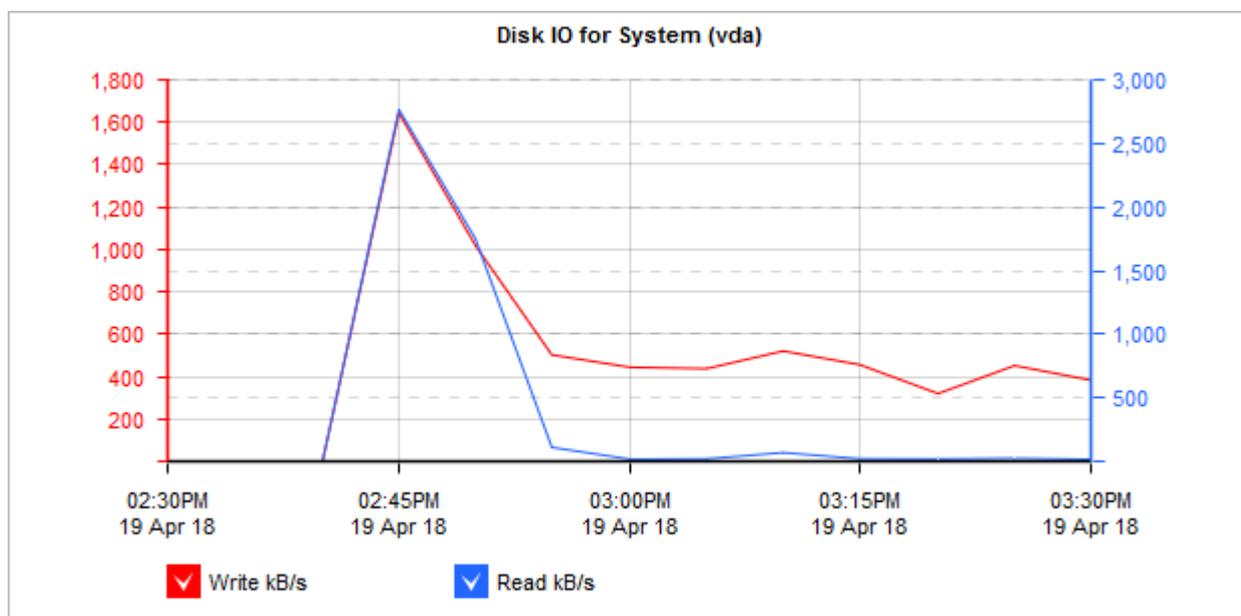
The Disk IO report shows read and write disk usage for each service in kB/s over time for the selected time period.

This report answers questions such as:

- » Has my disk I/O usage suddenly increased or over time? If so, which subsystem is responsible for the increased disk I/O usage?
- » If WAN memory acceleration, or CIFS acceleration, or edge cache performance is suffering, was there a decrease in its I/O load?
- » Was that decreased I/O load due to another subsystem's increased I/O load?
- » I swapped an appliance and loaded the same configuration and it seems slower. If all the I/O rates look lower, then maybe this is a physical disk issue.

The disk usage for each of the following services can be shown by selecting the desired service from the Service selector.

- » **System (vda)** – Total disk usage for all services combined for the single disk; Note that there may be two disks.
- » **monitor** – Disk usage required for storing the monitoring data
- » **swap** – Disk usage required for swapping/paging
- » **users** – Disk usage required to store the username information (i.e. data sent by the AD connector, manually configured users and groups, details of dynamic network objects)
- » **wan** – memory - Disk usage required for WAN memory acceleration techniques
- » **edge-cache** – Disk usage required for storing cached content for Edge Cache
- » **cifs** – Disk usage required for CIFS acceleration techniques



Screenshot 137: The Disk IO graph displays IO used by edge cache.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > Disk IO**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

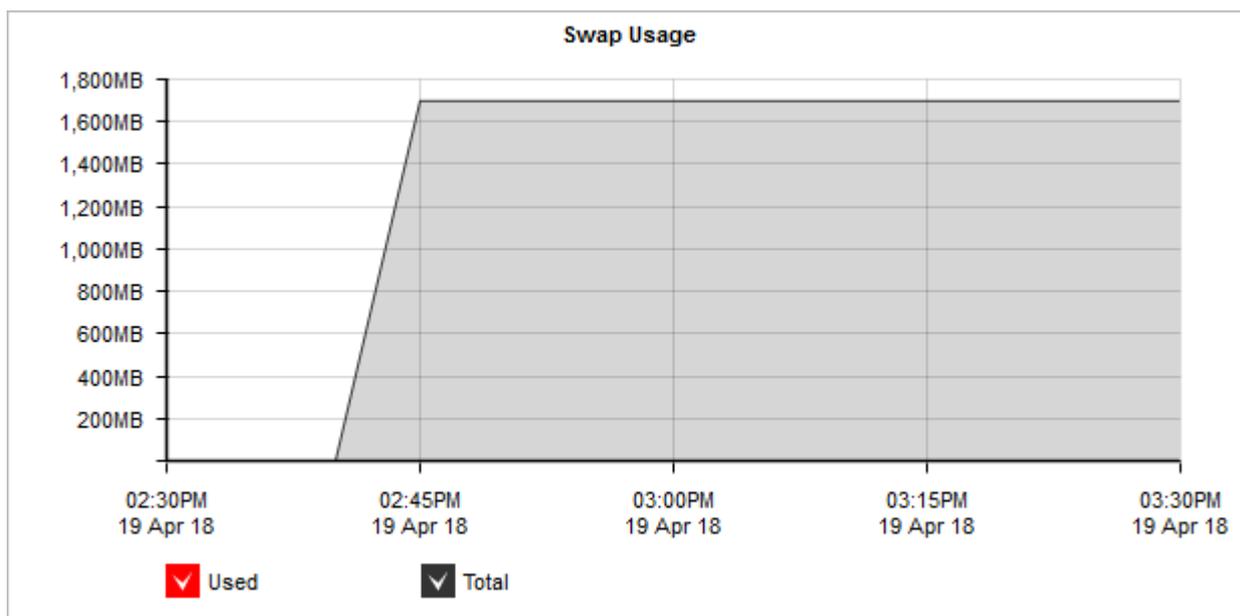
- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### Monitoring Exinda Appliance swap space usage

The Swap Usage report shows how much the appliance is swapping over time for the selected time period.

This report answers questions such as:

- » Could excessive swapping be affecting the performance of my appliance?



Screenshot 138: The Swap Usage graph displays system swap space utilization over time.

Where do I find this report?

To access the report:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Monitor > System > Swap Usage**.

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. For more information, refer to [Exporting, printing and scheduling reports](#) (page 325).

How do I interact with the interactive flash time graphs?

- » To understand how to get a better look at traffic patterns and to remove clutter on the time graph, see [Using Interactive Time Graphs](#).
- » To understand how to set the desired time range for a chart, see [Setting the Time Range](#).
- » To understand how to print the report or schedule the report, see [Printing and Scheduling Reports](#).

### 0.11.15 Viewing monitoring statistics

Your Exinda Appliance provides several ways to view the statistics gathered from monitoring your network. This section provides information about accessing those statistics and interpreting the information provided.

#### Understanding traffic direction

Monitored statistics reference the direction of the flow and the side of the appliance where the data is collected. It is assumed that the appliance is somewhat close to the edge of your network. Therefore, when installed inline, the appliance is cabled to the LAN and the WAN.

The LAN-side of the appliance is considered internal to your network and the WAN-side of the appliance is considered external to your network. When the traffic flows from the WAN-side to the LAN-side, the traffic is considered inbound. When the traffic is flowing from the LAN-side to the WAN-side, the traffic is considered outbound.

When traffic statistics are collected, the information about the flows are summarized. The data granularity that is stored is dependent on the time period shown.

| If the chart shows traffic for this duration... | The data granularity (the period that is summarized) is .. |
|-------------------------------------------------|------------------------------------------------------------|
| Real-time                                       | 1 seconds                                                  |
| 1 hour                                          | 5 minutes                                                  |
| 1 day                                           | 5 minutes                                                  |
| 1 week                                          | 30 minutes                                                 |
| 1 month                                         | 60 minutes                                                 |
| 1 year                                          | 24 hours                                                   |

#### Understanding the relationships between charts and data

On the new time series monitoring screens, there will be a pie chart and a bar chart showing the top data by volume for the specified time period and a time chart showing the same top times. There will be one set of these charts for inbound traffic and one for outbound traffic. In some instances, another chart that shows the top three items of another application type for each of the top data elements. For example, if looking at the virtual circuits chart, you will see the top virtual circuits as a bar chart. Then beside the bar chart, there will be a stacked horizontal bar chart which shows the top three applications for each of the top virtual circuits.



Screenshot 139: Virtual circuit details

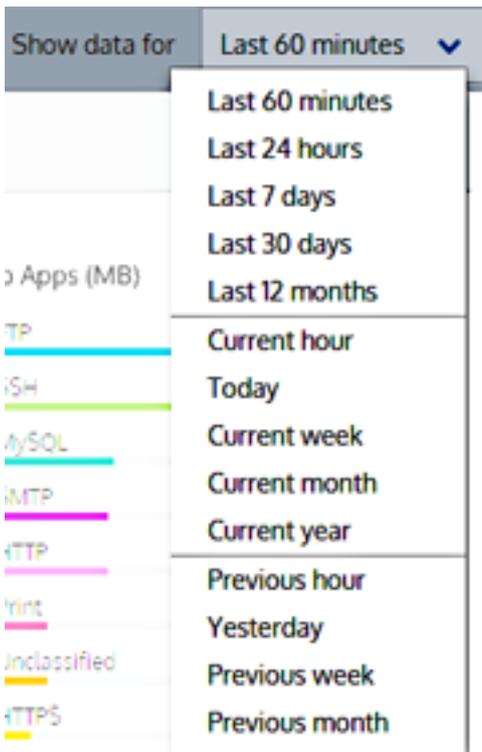
When you brush over an item in any of these charts, the item is highlighted in all the charts.

There are toggle buttons at the top of the chart which you can use to turn on and off different chart elements.

- » **Inbound** button – Toggles on or off all reporting for inbound traffic including all the charts and the data tables below the charts.
- » **Outbound** button – Toggles on or off all reporting for outbound traffic including all the charts and the data table below the charts.
- » **Stacked Chart** or **Line Chart** button – Toggles on or off the time series charts or selects a different view of the data. Press the down arrow next to the Stacked Chart button to select Line Chart to switch to the line chart. Conversely, press the down arrow next to the Line Chart button to select Stacked Chart and switch to the stacked area chart. The line chart shows the virtual circuits against the common zero baseline so that the throughput of the virtual circuits can be compared with one another and the pattern of a specific virtual circuit is clearer. You can look for particular patterns such as spikes or flat tops.
- » **Pie** button – Toggles on or off the data volume pie charts.
- » **Remaining Traffic** button – Toggles on or off the rest of the data that is not represented in the top virtual circuits. When toggled on, a gray chart series will appear in all charts (throughput, pie, top virtual circuits) representing all the virtual circuits in your network that are not explicitly represented in the top virtual circuits. If the remaining virtual circuits show vastly more data volume than the top virtual circuits, then the top virtual circuits may look insignificant relative to the total, and so you may need to toggle off the remaining traffic category to see the relative differences and usage patterns of the top virtual circuits.
- » **Data Details** button – Toggles on or off the tables of data below the charts.
- » **Top Apps Per X**(bar chart row) button –Toggles on or off extra charts showing the top 3 apps for each row in the bar chart and the number of apps for each row in the bar chart.

## Zooming into a time Interval on the time graphs

To view data within a time range, you can use the **Show data for** drop-down list to narrow the range. If does not provide the detail you need, you can narrow the search further by clicking and dragging within the graph or by using the zoom control below the graph. These methods allow you to define a custom time range.

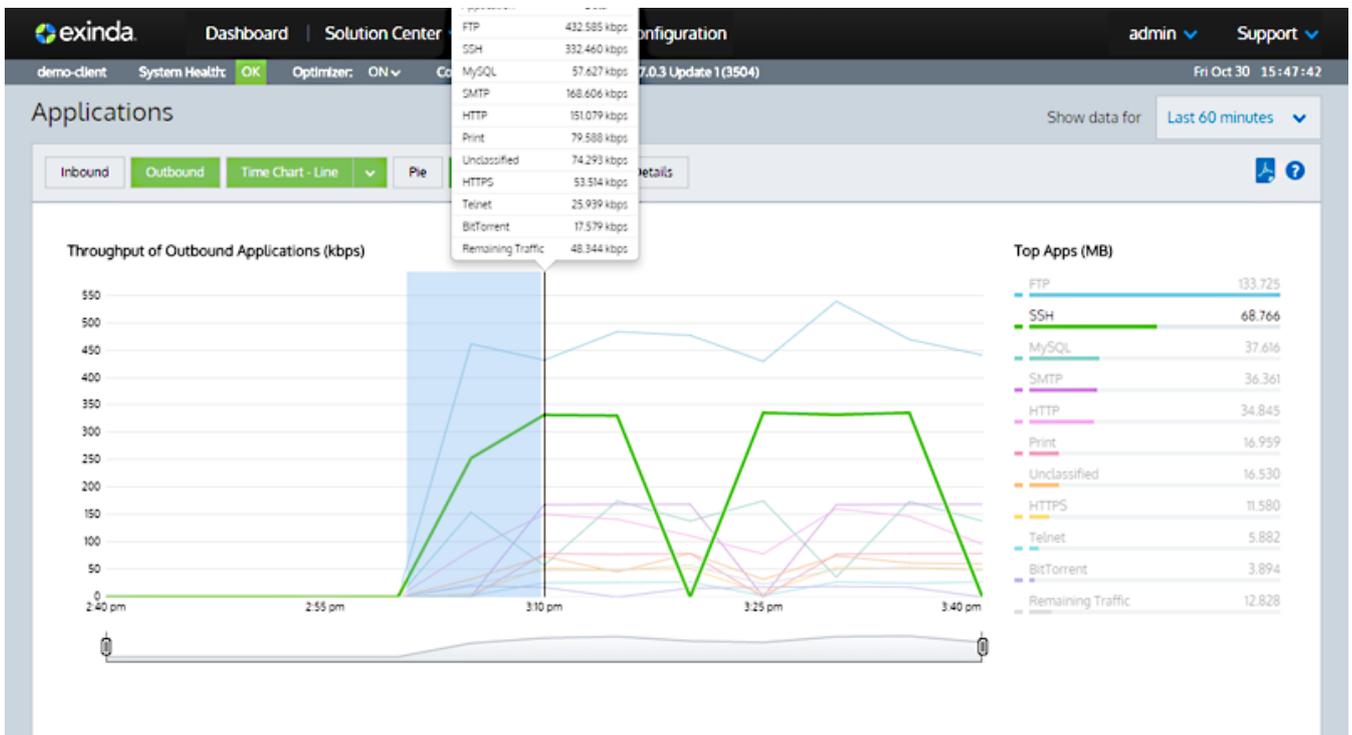


Screenshot 140: The "Show data for" drop-down list

Click and drag your mouse on the chart to select the desired time range.

Drag the handles on the zoom controller to modify the time range. As you drag, the area between the drag points becomes shaded. When you release, the shaded area expands to occupy the entire graph. To return to the initial time range, click on the Zoom Out

button to the left of the zoom controller.



You can also use the zoom controls that appear below the graph. Drag the handles in from the left and/or the right to isolate the data you need. The graph is dynamic, so you can immediately view the data. The handles remain in the



positions you left them, so the scope of the initial report remains evident.. When finished click on the Zoom Out button.

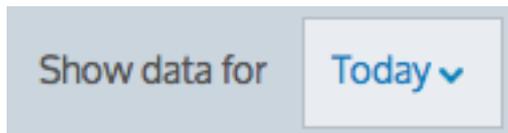


Screenshot 141: The Zoom controls

### Setting time ranges for charts and graphs

For each chart, you can set the time range that is reported in the chart.

At the upper-right of the report, select the desired date range from the drop down list. Custom time ranges are not supported.



After the date range is selected, the graphs and charts are immediately updated.

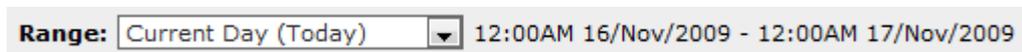
### Setting the time range (pre-7 series WebUI)

#### VERSION INFO

This topic relates to pre-7 series WebUI.

For each chart, you can set the time range that is considered in the chart.

Beside the title of the report, select the desired date range from the drop down list.



To specify a custom date range, in the drop down list select **Custom**. Select the start and end date and time to include in the report.



After the date range is select, the graphs and charts are immediately updated.

### Drilling into the chart data

Charts that show the data, such as applications, application groups, users, hosts, URLs, subnets, virtual circuits, allow you drill into a particular item to explore the details filtered by that item.

To drill-down:

- » For charts, click on an item in the table below the charts.
- » For bar charts, click on an item in the bar chart.

You can drill-down into the following details for each of the application types:

- » Application groups > applications > hosts
- » Applications > hosts
- » URLs > hosts

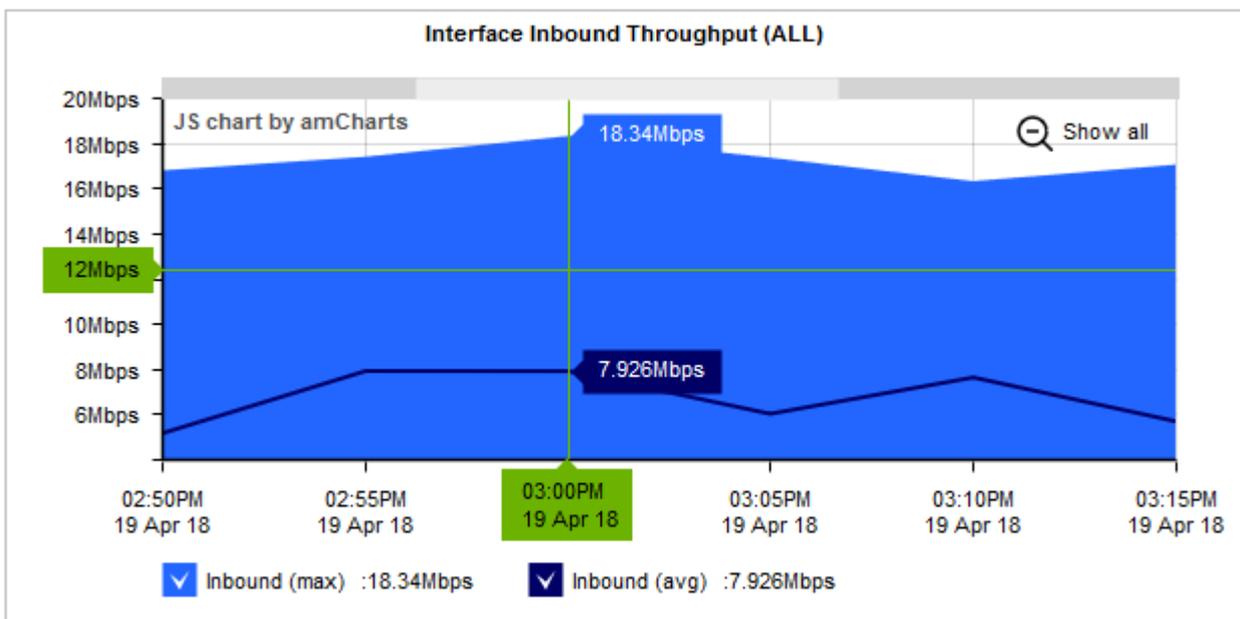
- » Users > applications or conversations or URLs or hosts
- » Hosts > applications or conversations or URLs
- » Subnets > applications > hosts
- » Subnets > hosts > applications or conversations or URLs
- » Subnets > users > applications or hosts or conversations or URLs
- » Subnets > conversations
- » Subnets > URLs
- » Virtual Circuits > applications

Drilling past these levels widens the filter. For example, drilling from users to applications filters the applications for that user, however, if you then drill into one of the applications for that users, it shows all hosts using that application.

### Using interactive time graphs

If you want a better look at a traffic pattern or if the chart is too cluttered, you can zoom in to a custom time range and remove time series lines that you are not interested in on the time graphs.

To zoom into a custom time range, click and drag your mouse on the chart to select the desired time range. To return to the initial time range click the 'Show all' magnifying glass icon. Any data displayed below these interactive graphs will automatically be updated with the data for the selected time range.



To remove a time series line, click on the check in the graph legend or in some cases the table below the chart to toggle off the display of that line.

#### NOTE

The interactive feature is only applicable to Flash generated graphs. To change the graph display option navigate to **Configuration > System > Setup > Monitoring**.

### Exporting, printing and scheduling reports

Monitoring reports can be exported as a PDF document, saved as a scheduled report, or can be printed directly from the Web UI. The following icons appear on the top-right of the interface:



» **Print:** Clicking on the Printer icon will open a new browser window and format the current report suitable for printing. It will then prompt you to select a printer.

#### NOTE

The print option is not available from the new application, subnet, and virtual circuit monitoring pages.

- » **Schedule PDF:** Clicking on the schedule icon will save the report configuration to the scheduled reports. It will prompt you for a report name, the scheduled frequency, the email addresses to send it to, and optionally a password if you choose to password protect the PDF.
- » **PDF:** Clicking on the PDF icon will render the current report as a PDF document and prompt you to save or open the PDF file once complete.

#### NOTE

Printed report and PDF reports may appear slightly different from the reports displayed on the Web UI.

### Generating PDF reports

PDF reports can be generated and downloaded on demand or generated and emailed at scheduled intervals. The content of the PDF reports can be configured in two ways:

- » exploring the data in the monitor screens and requesting a report
- » going to the Report page to configure the details of the PDF report

The following PDF report generation scenarios are supported:

- » Explore the data in the monitor screens and generate an ad hoc PDF report of what is shown on the screen.
- » Explore the data in the monitor screens and schedule a PDF report to be generated using the configuration and filters shown on the screen.
- » Configure a PDF report using the Report page .
- » Configure a PDF report using the Report page and request an on-demand generation of the PDF report.

Scheduled reports can be emailed to one or more email addresses by comma separating or semi-colon separating the email addresses in the appropriate field.

Scheduled reports can be generated hourly, daily, weekly, or monthly. The time range included in the report matches the frequency, that is, daily reports report on a day's worth of data and is generated once a day.

#### NOTE

- » Hourly scheduled reports are emailed to users at 22 minutes past the hour.
- » Hourly reports cannot be generated on-demand.
- » Daily scheduled PDF Reports are generated every morning at 1 a.m.

On-demand reports from the monitor pages can include any time range available to the monitoring screens, including custom time ranges.

Scheduled PDF reports can be branded by uploading your logo to be displayed on the title page of the reports.

Reports scheduled from the report page, can contain one or more charts in the PDF by selecting any number of charts.

To generate an on-demand PDF report from a monitor screen

1. Go to any monitor screen (except the Real Time screen) and configure it according to the available controls, such as the date range selector, the Internal or External selector for hosts and user charts and subnets, drilling into the data by selecting the links in the data tables under the charts, and so on.

2. Click on the Adobe PDF icon in the upper-right of the screen. 

3. The system will generate and present a PDF report that corresponds to what you see on the screen.

To schedule a PDF report from a monitor screen

1. Go to any monitor screen (except the Real Time screen) and configure it according to the available controls, such as the date range selector, the Internal or External selector for hosts and user charts and subnets, drilling into the data by selecting the links in the data tables under the charts, and so on.

#### NOTE

When you use this method to generate a report, the time range that appears on the front page of the report is updated to reflect the scheduled time range: "last 60 minutes" is updated to "last hour", "last 24 hours" is updated to "yesterday", etc. For example, the time range may be mapped to the following: 'Report Time Range 04:55PM 07 April 2015 to 05:55PM 07 April 2015'.

1. Click on the schedule PDF icon in the upper-right of the screen. 

2. Optionally protect PDF documents by specifying a password.

**PDF Security Option**

PDF Password Protected

Enter Password:

Re-enter Password:

3. On the Report Details page, specify the report name, the report frequency, and email addresses to send the report to.

- **Report Name**— a meaningful name for the new PDF Report.
- **Report Frequency**— This option is disabled when you click the Schedule button from a Monitor screen because the system assumes you want to use the timerange obtained from the Monitoring screen. If you need to change the time range, click the Add New PDF Report link at the top of the page.
- **Email Addresses**— one or more email addresses for scheduled PDF Reports. Email addresses are optional for on-demand PDF Reports. To specify multiple email addresses, comma or semicolon separate the addresses.

4. The system will add this scheduled report to the Reporting page (**Monitor > Schedule Reports**)

#### NOTE

PDF reports that were scheduled from a monitoring page cannot be edited. Ensure that you specify all the email addresses that you need it emailed to.

To schedule a new PDF report from the Reporting page

1. Go to **Monitor > Schedule Reports > PDF Reports**.

2. Click on the **Add New PDF Report** link at the top of the page.

3. Select the various reports you wish to include in the PDF report. Many of the reports available from the Web UI are available as PDF reports.

- **Interface Throughput Summary**— can select a specific interface(s), WCCP, or all WAN interfaces
- **Bridge PPS (Packets per Second) Summary** — can select specific bridge(s), WCCP, or all bridges
- **Network Summary**
- **Subnets Summary**
- **Detailed Subnet Reports**— can select specific subnet(s) and specific details for each subnet (i.e. application detail, conversation detail, host detail, URL detail, and user detail)
- **APS**
- **SLA**
- **TCP Health**
- **TCP Efficiency**
- **VoIP**
- **Virtual Circuit**— can select specific virtual circuit(s) and specific statistics (i.e. Optimization Policy Throughput Statistics and/or Discard Statistics)
- **Optimization Prioritization Ratio Statistics**
- **Optimization Reduction Statistics**
- **Optimization Edge Cache Statistics**
- **Appliance Statistics**— can select specific appliance system statistics (i.e. Concurrent Connections, Accelerated Connections, CPU Usage, CPU Temperature, RAM Usage, SWAP Usage, Disk IO)

4. Optionally protect PDF documents by specifying a password.



**PDF Security Option**

PDF Password Protected

Enter Password:

Re-enter Password:

5. On the Report Details page, specify the report name, the report frequency, and email addresses to send the report to.

- **Report Name**— a meaningful name for the new PDF Report.
- **Report Frequency**— the time range of the report and the frequency that it is sent. For example, daily frequency presents a day's worth of data and is emailed once a day.
- **Email Addresses**— one or more email addresses for scheduled PDF Reports. Email addresses are optional for on-demand PDF Reports. To specify multiple email addresses, comma or semicolon separate the addresses.

6. The system will add this scheduled report to scheduled report list.

To view a scheduled report on demand or edit or delete a report

1. Go to **Monitor > Schedule Reports > PDF Reports**.

2. The scheduled PDF reports are listed with a description of the charts that will be included in the report and the list of email addresses it will be sent to.

| PDF Reports                                 |                                                                                                                                                                                               |          |                                                                                                                                                                         |                                     |                                       |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|---------------------------------------|
| Name                                        | Exported Data                                                                                                                                                                                 | Email(s) | On-Demand                                                                                                                                                               | Edit                                | Delete                                |
| <b>Desktop_b_b</b><br>(Last 60 Minutes)     | <b>Virtual Circuit Detailed (Desktop Networks):</b><br>Peak vs Average Throughput Report<br>Optimization Policy Throughput Statistics                                                         |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
| <b>hosts_yesterday</b><br>(Scheduled Daily) | <b>Custom Selection</b>                                                                                                                                                                       |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
| <b>Test</b><br>(Last 60 Minutes)            | <b>Subnet Detailed (Desktops):</b><br>Applications<br>Conversations<br>Hosts<br>URLs<br>Users<br><br><b>Virtual Circuit Detailed (Desktop Networks):</b><br>Peak vs Average Throughput Report |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
| <b>test</b><br>(Scheduled Daily)            | <b>Custom Selection</b>                                                                                                                                                                       |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
| <b>VC-report</b><br>(Last 7 Days)           | <b>Virtual Circuit Detailed (Engineering Servers):</b><br>Peak vs Average Throughput Report<br>Optimization Policy Throughput Statistics                                                      |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |

- To view the report, click the Adobe PDF icon.
- To email the report to email recipients on demand, click the mail icon.
- To edit or delete a configured PDF report, click on the appropriate button next to the report in the table.

#### NOTE

PDF reports that were scheduled from a monitoring page cannot be edited. Ensure that you specify all the email addresses that you need it emailed to.

PDF reports can only be emailed on-demand if the report was configured with one or more email addresses.

To add a custom logo to the cover of the scheduled reports

- Go to **Monitor > Schedule Reports > Custom Logo**
- Upload your custom logo.
- The system will insert the logo on the cover page of any scheduled PDF report.

**Custom Logo**

Upload New Custom Logo:

#### NOTE

Files should be no more than 300px wide by 300px high and must be in PNG format with maximum file size of 3MB.

### CSV Reporting

CSV Reporting allows you to configure the export of raw CSV data to be emailed or downloaded either on demand or at scheduled intervals. Exported data can be sent to multiple recipients by comma or semicolon separating email addresses.

## NOTE

To configure a CSV Report, navigate to Report | CSV Reports on the Web UI, advanced mode.

CSV Reports are listed in the table on this page. CSV Report can be generated and either emailed or downloaded on-demand by clicking either the ZIP icon (to generate and download) or the envelope icon (to generate and email). CSV Reports can only be emailed on-demand if the report was configured with one or more email addresses.

You can also Edit or Delete a configured CSV Report by clicking on the appropriate button next to the report in the table.

| CSV Reports                          |                                  |          |                                                                                                                                                                         |                                     |                                       |
|--------------------------------------|----------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|---------------------------------------|
| Name                                 | Exported Data                    | Email(s) | On-Demand                                                                                                                                                               | Edit                                | Delete                                |
| <b>currentweek</b><br>(Current Week) | <b>Summary Reports:</b><br>flows |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |
| <b>currentday</b><br>(Today)         | <b>Summary Reports:</b><br>flows |          |   | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |

New CSV Reports can be added by using the form at the top of the page.

### Report Details

Report Name:

Report Frequency:

Email Addresses:

[Email Addresses is Optional for On-Demand Report]

Screenshot 142: Report Details

| Property         | Description                                                                                                                                                                                          |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Report Name      | Specify a meaningful name for the new CSV Report.                                                                                                                                                    |
| Report Frequency | Specify a time range for this CSV Report. Scheduled reports can be generated Daily, Weekly or Monthly. On-demand reports can include any time range available to the Exinda appliance.               |
| Email Addresses  | Specify 1 or more email addresses for scheduled CSV Reports. Email addresses are optional for on-demand CSV Reports. To specify multiple email addresses, comma or semicolon separate the addresses. |

## NOTE

Daily scheduled CSV Reports are generated every morning at 1am.

For information about the schema used in CSV Reports, consult the SQL Access using ODBC How to Guide.

To find this functionality, go to **Monitor > Schedule Reports > CSV Reports**.

## 0.12 Monitoring applications with the Exinda Solution Center

The Exinda Solution Center provides a series of predefined monitors you can run to generate network performance reports for applications like FTP, SSH, Salesforce.com, Office365 VoIP and many more.

The generated reports answer questions, such as:

- » How is salesforce.com performing for network users?
- » How are critical applications performing on the network?
- » How can I best mitigate data center disasters?

Each solution description indicates which Exinda OS version is required to run the solution, shown both in the solution list and in each solution description. You may need to upgrade your Exinda OS version to take advantage of the desired solutions. Some solutions may not yet be available and are shown as 'Coming soon'.

**Exinda Solution Center**

Suggest a new solution

Is there anything missing from this list that would be of value to you? [Let us know.](#)

Application Performance

- [Salesforce.com Performance](#) 7.0.0
- [Office 365 Performance](#) 7.0.0
- [Adobe Creative Cloud Performance](#) 7.0.0
- [Custom Application Performance](#) 7.0.0
- [VoIP Performance](#) 7.0.0

Network Governance

- [Recreational Traffic](#) 7.0.0
- [RIAA Notice Prevention](#)

Project Readiness

- [Data Center Continuity](#)

WAN Planning

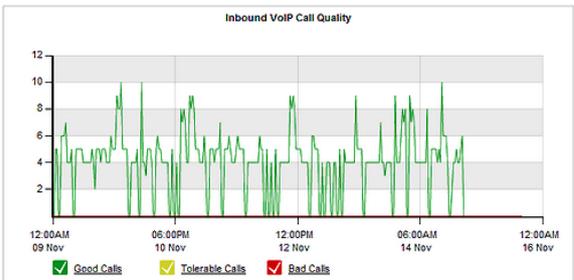
- [Bandwidth Usage Summary](#)
- [Bandwidth Usage - Top Apps](#)

### VoIP Performance

VoIP is business-critical for many companies. To prevent voice delays and jitter, VoIP performance monitoring is key to ensure predictable call quality.

With the VoIP report, you can see the number of good, tolerable and bad conversations and gain valuable insight into what is causing poor voice quality.

When VoIP performance is poor, Exinda makes it easy to pinpoint where the issue is to remediate it quickly and restore high quality call service. Ensure your employees can always communicate reliably with Exinda's VoIP report.



**Worst 30 Inbound VoIP Conversations**

| Internal Host | External Host | Delay (ms) | Jitter (ms) | Loss (%) | MOS  | rFactor |
|---------------|---------------|------------|-------------|----------|------|---------|
| 253.7.254.1   | 253.11.254.1  | 0          | 0.00        | 0.00     | 4.28 | 87.90   |
| 173.253.253.1 | 173.5.254.1   | 0          | 0.00        | 0.00     | 4.28 | 87.90   |

This solution is available on Exinda OS 7.0.0

Screenshot 143: The Exinda Solution Center

Performance monitors are divided into four solution categories: Application Performance, Network Governance, Project Readiness and WAN Planning.

Each monitor has a description you can display by clicking the its link in the left panel. Descriptions detail usage information and which Exinda OS version is required to run the monitor. Some monitors also have short video descriptions.

### 0.12.1 How performance reports work

An Exinda Appliance continuously collects network traffic data. Performance reports in the Exinda Solution Center provide insight into that data by grouping it in meaningful ways and displaying it in charts, tables and graphs.

The process starts by analyzing traffic and computing initial threshold values to create a baseline. A baseline requires an hour's worth of network traffic data. If no traffic is observed for an application during a baselining period, the process continues until enough data is collected.

The baseline process may not take an hour. If an Exinda Appliance has observed and stored traffic for the application within the hour the baseline process starts, the baseline process uses that stored information and only waits enough time for a total hour of data to be collected.

Exinda Network Orchestrator

| 331

For example, if you create an application monitor with ten minutes left in an hour and the Exinda captured network traffic for the application in that hour, the baseline process analyzes the previous fifty minutes of collected traffic data and completes the baseline period with data collected in the remaining ten minutes.

---

---

## 0.12.2 Using Application Performance reports

Application Performance monitors generate reports that display information about application users, application performance, application bandwidth consumption, and the amount of reduction achieved (if applicable).

Application Performance solutions provide a predefined set of application monitors. Except for **VoIP**, application monitors generate similar reports.

You can choose a monitor from the main Exinda Solution Center screen or click the **Custom Application Performance** link to bring up a list of applications to choose from.

### NOTE

The report description lists the minimum version of Exinda OS required to run the report. If your Exinda OS does not meet the requirement, the **Run** button is not available.

### Running an Application Performance report

The Exinda Solution Center lists a set of predefined reports on the Solution Center main screen. You can choose to run one of those or you can choose to run one of the dozens of other reports by clicking **Custom Application Performance**.

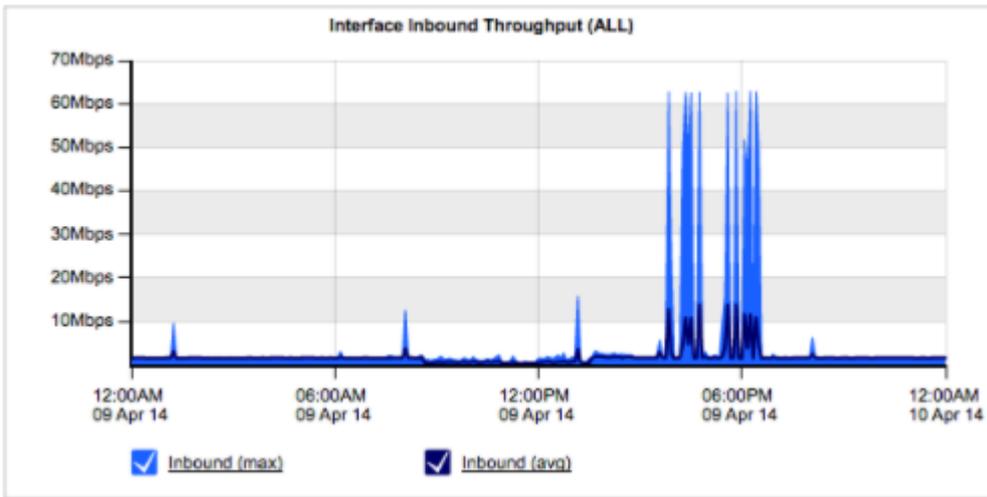
1. Go to **Solution Center > Show Solution Center**.
2. Under **Application Performance**, click the name of the report to run
3. Click **Run**. A confirmation screen opens.
4. Click **Ok**.

### NOTE

After the initial run, you can access the report by clicking **Solution Center**, hovering over **Applications** and clicking the report name.

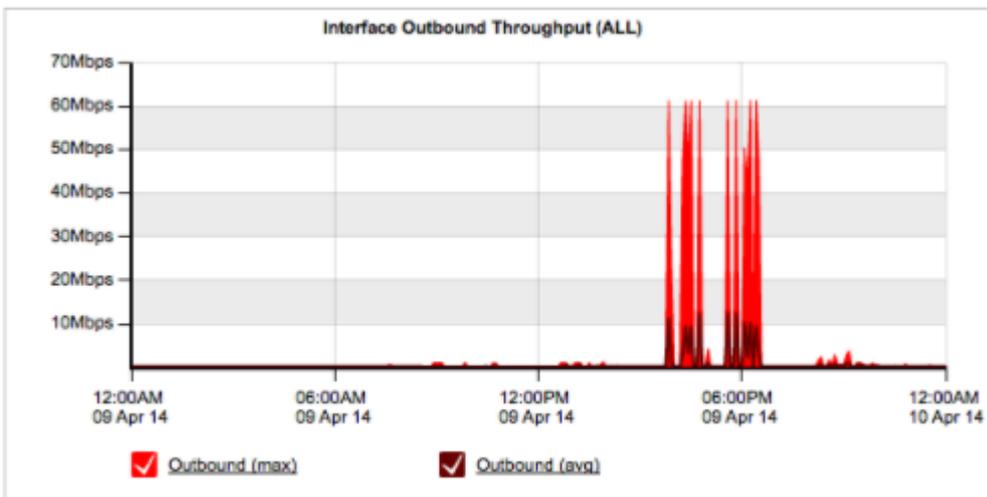
### Understanding the data displayed in an Application Performance report

An Application Performance report shows the network user experience of an application through a series of charts, tables and graphs.

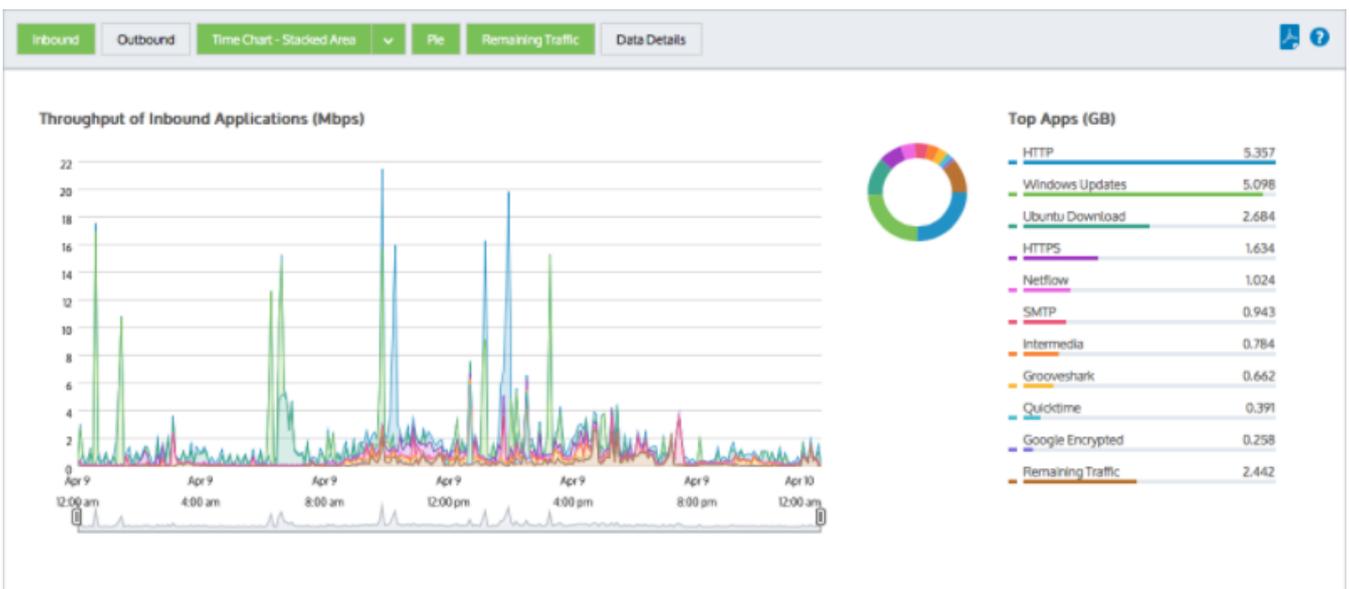


Screenshot 144: The Inbound Bandwidth chart shows data measured in the WAN-side before accelerated traffic is decompressed.

The Inbound and Outbound Bandwidth charts show how much bandwidth an application uses. Chart lines typically show spikes instead of raised flat tops. Flat tops may indicate traffic limitations imposed by policy rules.



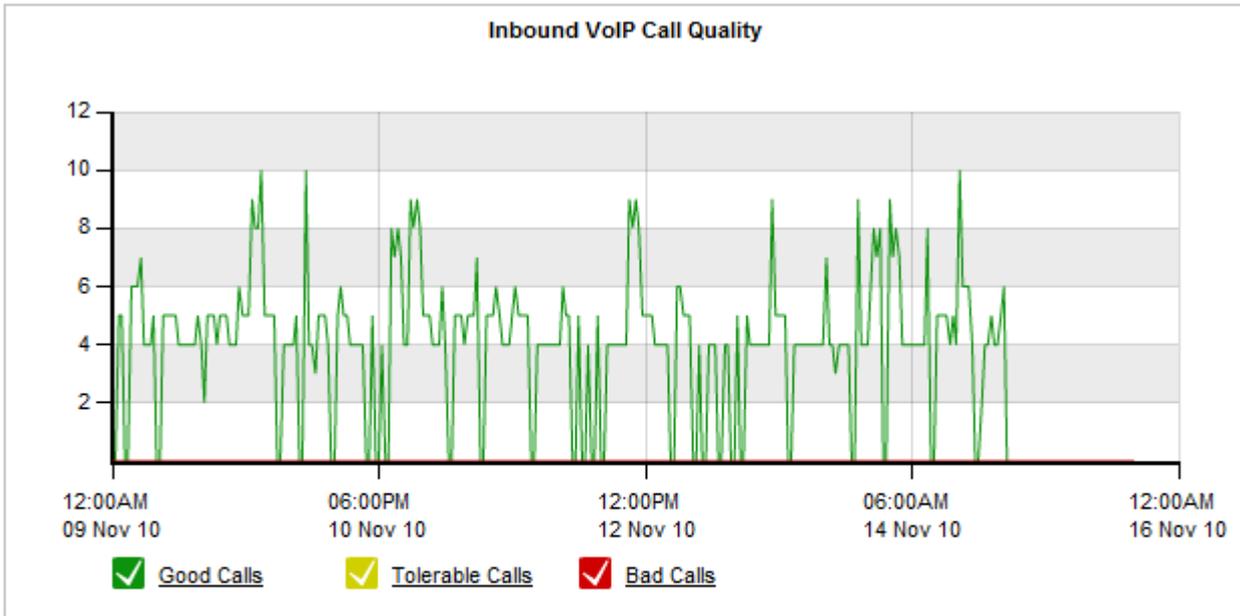
Screenshot 145: The Outbound Bandwidth chart shows data measured after acceleration and traffic shaping policies have been applied.



Screenshot 146: Users and hosts bar charts.

Users and Hosts bar charts display bandwidth volume by top listeners and talkers. Multi-user applications typically show an even distribution among top users or hosts. If a user or host displays more bandwidth volume than other users and hosts, that situation may warrant investigation.

You can choose to show internal endpoints (LAN-side of an Exinda Appliance), external endpoints (WAN-side of an Exinda Appliance), users only or hosts only. For more information, refer to [Monitoring real time application response](#) (page 252).



| Worst 30 Inbound VoIP Conversations |               |            |             |          |      |         |
|-------------------------------------|---------------|------------|-------------|----------|------|---------|
| Internal Host                       | External Host | Delay (ms) | Jitter (ms) | Loss (%) | MOS  | rFactor |
| 253.7.254.1                         | 253.11.254.1  | 0          | 0.00        | 0.00     | 4.28 | 87.90   |
| 173.253.253.1                       | 173.5.254.1   | 0          | 0.00        | 0.00     | 4.28 | 87.90   |

Screenshot 147: The Application Performance Scores and Metrics table.

Application Performance Scores and Metrics displays the APS score for the application. A good score is between 8.5 and 10.0. A score less than 7.0 may warrant an investigation.

### 0.12.3 Bandwidth usage

#### bandwidth usage - top apps

Knowing how much bandwidth your top applications consume can provide insight into whether controlling particular apps could help effectively reduce your throughput.



Screenshot 148: Top App usage graph

To see how much bandwidth your top application are using go to **Monitor > Applications**. For more information, refer to [Viewing application traffic volume](#) (page 275).

#### bandwidth usage - top summary

Knowing how much bandwidth your users are consuming is crucial to managing your network. If your link is congested, you need to know whether to plan for a bandwidth upgrade, or if policy-based shaping can effectively reduce your throughput instead.

To see how much of your network's bandwidth is in use, go to **Monitor > Interfaces**. For more information, refer to [Monitoring interface throughput](#) (page 256).

### 0.12.4 Using the Application Performance Monitor VoIP report

The VoIP report monitors and reports on the quality of VoIP transactions in a network. It displays data using telecommunication industry standard measures like MOS and rFactor.

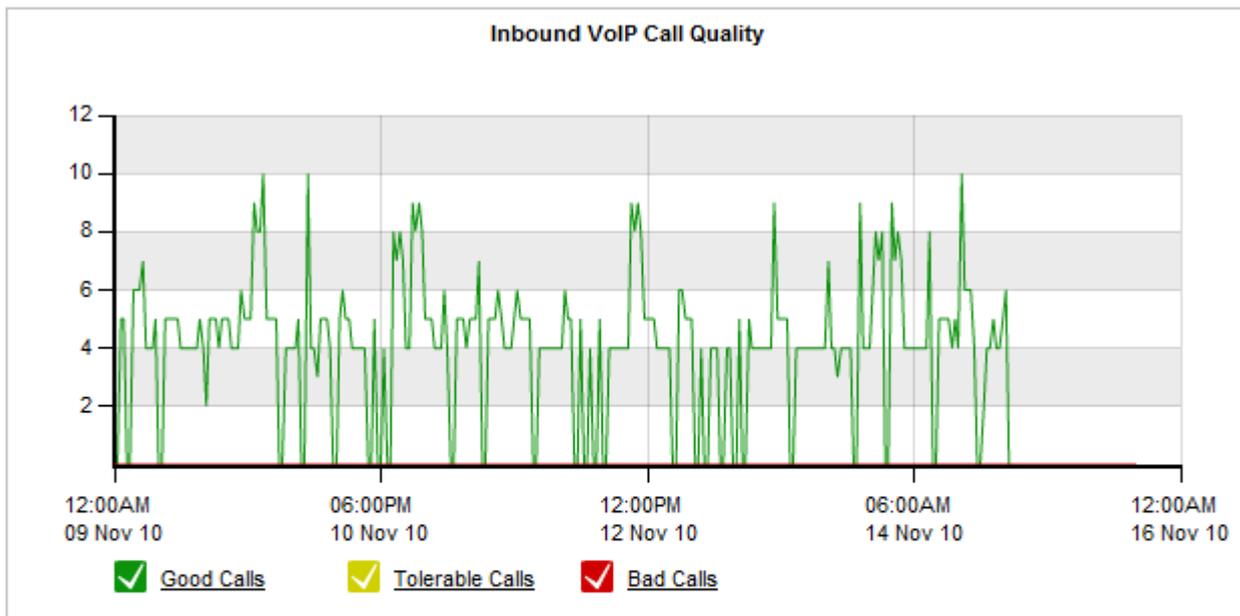
#### Running the Application Performance VoIP report

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Solution Center > Show Solution Center**.
5. Under **Application Performance**, click **VoIP Performance > Run**. A confirmation screen opens.
6. Click **Ok**. The report opens.

#### NOTE

After the initial run, you can access the report by clicking **Solution Center > VoIP Performance**.

## Understanding the data displayed in an Application Performance VoIP report



| Worst 30 Inbound VoIP Conversations |               |            |             |          |      |         |
|-------------------------------------|---------------|------------|-------------|----------|------|---------|
| Internal Host                       | External Host | Delay (ms) | Jitter (ms) | Loss (%) | MOS  | rFactor |
| 253.7.254.1                         | 253.11.254.1  | 0          | 0.00        | 0.00     | 4.28 | 87.90   |
| 173.253.253.1                       | 173.5.254.1   | 0          | 0.00        | 0.00     | 4.28 | 87.90   |

The graph shows three series, representing the number of "Good", "Tolerable" and "Bad" calls over time. The table below the chart lists the worst quality inbound and outbound VoIP calls for the specified time period.

The meanings of the colors:

- » Good (green) - MOS greater than 4.
- » Tolerable (yellow) - MOS between 2 and 4.
- » Bad (red) - MOS less than 2.

### What is MOS?

MOS, or Mean Opinion Score, is a measure of all quality. Historically, users would rate their call experience on a five point scale.

An Exinda Appliance automates MOS ratings, taking into account network dependency conditions. The scores have the following meanings.

- 5 - Perfect, like face-to-face conversations or radio reception.
- 4 - Fair, imperfections perceived, but the sound is clear. Cell phones calls are typically rated fair.
- 3 - Annoying.
- 2 - Very annoying, nearly impossible to communicate.
- 1 - Impossible to communicate.

### What is rFactor?

rFactor is a measure of the call quality in IP networks taking into account network delay and impairments. rFactor ranges

from 0 (extremely poor quality) to 100 (high quality). Any rFactor less than 50 is not acceptable.

### 0.12.5 Recreational Traffic

The Recreational Traffic Report shows the usage of recreational application groups over time for the specified time period. It shows information for games, instant messaging, peer-to-peer, social networking, and streaming. This report can answer questions such as:

- » How much data is going over my network for recreational applications?
- » How many hosts are involved?
- » How much time is spent transferring the data over the network?

Having visibility into key recreational applications is the first step in being able to manage them. These applications are generally undesirable because they can impact the performance of key business applications, negatively impact customer experience, reduce the productivity of users, introduce viruses to the network, and enable downloading of illegal or copyrighted material.

#### How to set up this report?

The report for recreational traffic can be created on the Exinda Solution Center

1. Go to **Solution Center > Show Solution Center**.
2. Under **Network Governance**, click **Recreational Traffic > Run**. A confirmation screen opens.
3. Click **Ok**.

#### NOTE

Once the report has been set up, you can access it by clicking **Solution Center > Recreational Traffic**.

### 0.12.6 Data Center Continuity

When a data center disaster happens, the company is in crisis. Depending on what caused the issue, you may be dealing with floods, fires, building integrity issues, or power and hardware failure.

To proactively prepare for an emergency, you need to have a network risk mitigation plan in place. When an emergency strikes, you don't want to be tweaking policies, reading help manuals, or trying to figure out how to control your network – you want to enact your risk mitigation plan.

Create your risk mitigation plan with Exinda's Data Center Continuity Solution.

There are several phases to your risk mitigation plan.

#### Phase One: Normal day-to-day operations

To ensure that your backups have adequate network bandwidth to complete on time, yet are not interfering with your day-to-day business activities, create a policy that allocates adequate bandwidth to your backup file transfer.

Go to the Optimizer and configure your policies for everyday use

#### Phase Two: Data center disaster strikes

To ensure that your remaining data center(s) can handle the traffic, before disaster strikes, create a separate policy configuration that will be suitable if a data center goes down. When the data center actually goes down, all you need to do is enact the disaster policy set.

#### Pre-disaster:

1. Go to **Optimizer** and configure your policies as you need them to behave immediately upon a data center disaster.
2. Go to **System > Maintenance** and save your disaster configuration as a named file. (e.g., DC-down-policies).

#### As disaster strikes:

1. Go to **System > Maintenance** and apply the previously saved configuration.
2. Check the checkbox for the desired configuration and click the **Switch To**.
3. Restart the Optimizer, by selecting **Restart from the Optimizer** menu on the **status bar**.

#### Phase Three: Data center recovery

To ensure that your remaining data center(s) can handle the traffic while your affected data center is recovering, create a separate policy configuration that will be suitable during recovery. When the data center starts the recovery process, all you need to do is enact the recovery policy set.

#### Pre-disaster:

1. Go to **Optimizer** and configure your policies as you need them to behave during a data center recovery phase.
2. Go to **System > Maintenance** and save your disaster recovery configuration as a named file. (e.g., DC-recovery-policies)

#### As DC recovers:

Go to **System > Maintenance** and apply the previously saved configuration.

Check the checkbox for the desired configuration and click the **Switch To**.

Restart the Optimizer, by selecting **Restart from the Optimizer** menu on the **status bar**.

## 0.12.7 RIAA Notice Prevention

The Recording Industry Association of America (RIAA) and the Motion Picture Association of America (MPAA) work to protect the intellectual property of artists, record labels, actors, and film studios. When users download copyrighted material on a campus network, educational institutions receive RIAA/MPAA notices from their Internet service providers, and the organization must take action to prevent further illegal downloads.

To reduce the amount of copyright violations on your network, you can set policies to reduce the ability to use encrypted P2P applications.

Optimizer | Policies | Wizard

Policies define the traffic to match as well as the action to take on that traffic.

**Add New VC Policy**

Policy Name: P2P - Choke 1%-3%

VC Policy Number: 210

Schedule: ALWAYS

Action: Optimize

Policy Enabled:

Filter Rules:

| VLAN | Source | Direction | Destination | ToS/DSCP | Application |
|------|--------|-----------|-------------|----------|-------------|
| ALL  | ALL    | Both      | ALL         | ALL      | ALL         |
|      |        | Both      |             |          |             |
|      |        | Both      |             |          |             |
|      |        | Both      |             |          |             |

NOTE: Filters cannot be deleted from this page. Please go to the [Optimizer|Policies](#) to delete filters.

Add New Policy | Cancel

Screenshot 149: Policies can be used to control traffic

Ensure your organization stays compliant with digital copyright laws with Exinda's RIAA Notice Prevention Solution by creating a policy that will significantly reduce that amount of P2P traffic that is allowed on the network, say to 1-3% of the available bandwidth with a burst priority of 10.

To proactively reduce the amount of RIAA copyright compliance violations:

1. Go to **Configuration > Traffic Policies > Optimizer**.
2. Click **Create New Policy...**
3. Set the name to an appropriate name. E.g. P2P - Choke 1%-3%
4. Set the fields as follows:

| Field                     | Value    |
|---------------------------|----------|
| Schedule                  | Always   |
| Action                    | Optimize |
| Policy Enabled            | Checked  |
| Guaranteed Bandwidth      | 1%       |
| Burst (Maximum Bandwidth) | 3%       |
| Burst Priority            | 10 (Low) |
| Filter Rules Application  | P2P      |

5. Click **Add New Policy**.

To allow legitimate uses of P2P create one or more policies to provide exceptions to the rule for particular users, departments, or subnets. Ensure that this policy is ordered higher in the Optimizer tree than the P2P - Choke policy.

To create an exception policy:

1. Go to **Configuration > Traffic Policies > Optimizer**.
2. Click **Create New Policy...**

3. Set the name to an appropriate name. E.g. P2P – Allow legitimate use 5%-100%

4. Set the fields as follows:

| Field                     | Value    |
|---------------------------|----------|
| Schedule                  | Always   |
| Action                    | Optimize |
| Policy Enabled            | Checked  |
| Guaranteed Bandwidth      | 5%       |
| Burst (Maximum Bandwidth) | 100%     |
| Burst Priority            | 6 (Low)  |

5. Click **Add New Policy**.

### 0.12.8 Using Network Governance reports

Network Governance reports provide data that allow you to manage your network resources according to ethical boundaries set by your organization. Solution categories include Recreational Traffic and RIAA Notice Prevention.

#### Understanding the data displayed in the Recreational Traffic report

The Recreational Traffic Report shows the bandwidth consumption of recreational applications for a specified time period. It shows network traffic data for games, instant messaging, peer-to-peer, social networking, and streaming.

Recreational applications are generally deemed undesirable on business networks because they impact the performance of key business applications, negatively impact customer experience, reduce user productivity, introduce viruses to the network, and enable downloading of illegal or copyrighted material.

| Recreational - -                  |       |        |      | <a href="#">More details</a> ? |
|-----------------------------------|-------|--------|------|--------------------------------|
| Application                       | Hosts | Time   | Data |                                |
|                                   | 2     | 11m    | 2MB  |                                |
| <a href="#">Games</a>             | 1     | 10s    | 0MB  |                                |
| <a href="#">Instant Messaging</a> | 0     | 0s     | 0MB  |                                |
| <a href="#">P2P</a>               | 0     | 0s     | 0MB  |                                |
| <a href="#">Social Networking</a> | 2     | 4m 30s | 2MB  |                                |
| <a href="#">Streaming</a>         | 1     | 6m 20s | 0MB  |                                |

Screenshot 150: The Recreational Traffic report.

#### Running the Recreational Traffic report

1. Go to **Solution Center > Show Solution Center**.
2. Under **Network Governance**, click **Recreational Traffic**.
3. Specify any details that the wizard requires.
4. Click **Ok**.

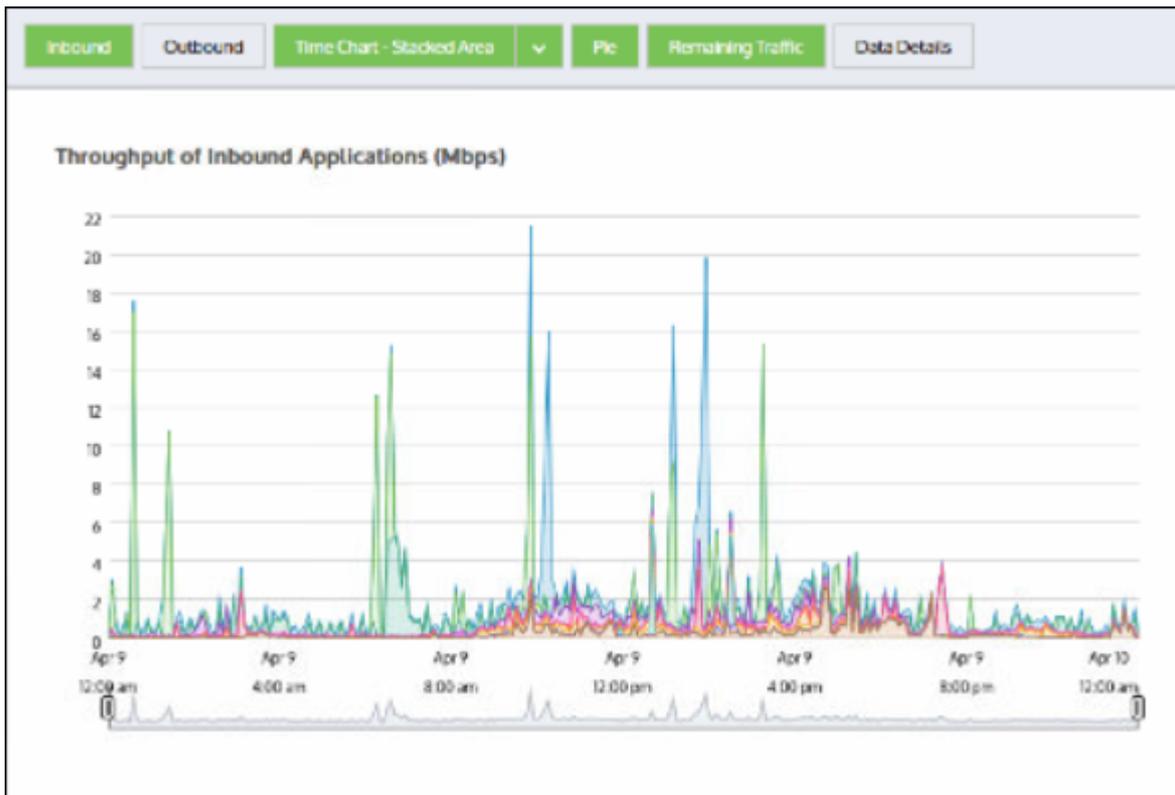
## NOTE

After the initial run, you can access the report by clicking **Solution Center > Recreational Traffic**.

### 0.12.9 Answers to common questions about Solution Center Application Performance

What type of data is available in an Application Performance report?

An Application Performance report shows the network user experience of an application through a series of charts, tables and graphs.



Screenshot 151: Inbound bandwidth application

The Inbound and Outbound Bandwidth charts show how much bandwidth the application is using. You should expect the bandwidth to show spikes instead of raised flat tops. Flat tops in the graph often indicate traffic may be limited by [policy rules](#).

The chart shows data measured on the WAN-side of the appliance before accelerated traffic is decompressed for inbound traffic and after acceleration and traffic shaping policies have been applied for outbound traffic.

You can overlay the data measured on the LAN-side of the appliance to show the amount of reduction achieved due to acceleration and traffic shaping.

The users and hosts bar charts show the WAN-side data volumes consumed by the top users and hosts for the application. Typically, applications are used by multiple users or hosts and the traffic distribution is fairly even amongst the top users or hosts.

If one user or host shows considerably more data volume than the other users, it may be reasonable behavior or it may indicate a problem worthy of further investigation. Also, you can choose to show just internal endpoints, that is, hosts and users on the LAN-side of your appliance, or just external endpoints, that is, hosts and users on the WAN-side of your appliance. You can also choose to show just users, just hosts, or both.

Application Performance Score Metrics You should expect a good score (between 8.5 and 10.0). If the score is less than 7.0, you may want to investigate.

### What is Application Performance report baselining?

An Application Performance monitor requires a baseline understanding of observed traffic for an application in your network. The process of collecting data and setting a baseline is called baselining.

Once you create a monitor, the baselining operation starts automatically analyzing traffic and begins the process of computing initial threshold values. This process requires an hour's worth of network traffic data.

If no traffic is observed for an application during a baselining period, the baselining process repeats until traffic is observed and thresholds are calculated.

The baselining process may not always take an hour if an Exinda Appliance has observed and stored traffic for the application within the hour the baselining process starts.

For example, if you create an application monitor with ten minutes left in an hour and the Exinda captured network traffic for the application in that hour, the baselining process analyzes the previous fifty minutes of collected traffic data. It then completes the baseline period in the remaining 10 minutes of the hour.

### What if the Solution Center indicates there are no solutions?

The solution descriptions are served up from an Exinda hosted server. If your Exinda Solution Center indicates there are no solutions, check for Internet connectivity and connectivity to the Exinda hosted server. If the Exinda hosted server is down, previously instantiated solutions will still be available in your Solution Center.

### What if a solution requires a higher Exinda OS version?

The solution's **Run** button won't be available until you upgrade your Exinda OS to the appropriate version or higher.

### Can I run a solution more than once?

Yes. You can run a solution multiple times if the solution takes configuration parameters. For instance, you can create multiple Custom Application Performance monitors where each report monitors a different application. For solutions without configuration parameters, such as VoIP Performance, you cannot create the solution more than once.

## 0.12.10 Adding and deleting Solutions

Use the instructions that follow to add Exinda Solutions to your configuration , and later, if necessary, delete them. When defined, the solutions provide access to reports that focus on the specified applications.

### To add a solution

The Exinda Solution Center includes several predefined solutions, but you can also define you own.

1. Go to **Solution Center > Show Solution Center**. The solutions, broken into categories, are accessible through the various links on the left.
2. Select the desired solution from the list.
3. Click the **Run** button.
4. Specify any details that the wizard requires. The final page of the wizard specifies where to find the report.
5. Clicking **Ok** will take you to your report.

#### NOTE

Once a report has been set up, a link to it is available from the main task bar at the top of the page. Click **Solution Center > Report Name**.

### To delete a solution

The only way to delete a solution is through the command line. However, for some solutions you need to determine the solution ID from the Web UI before you can remove the solution.

1. Go to **Configuration > Objects > Service Levels > Application Performance Scores**.
2. Find the application in the **APS Name** column.

#### NOTE

The formatting of the name includes "Solution Center" and the ID. For example the CIFS APS object would be called CIFS Solution Center (208).

3. Make a note of the solution ID.
4. Open the CLI.
5. At the prompt, type `no solutionc <id>`. Examples:

- `no solutionc 208`
- `no solutionc VoIPPerformance`
- `no solutionc RecreationalTraffic`

### 0.12.11 Setting a new baseline

Use the following instructions to set a new baseline for an application performance score. If you need to set a new baseline, you should do this when you expect the application to perform reasonably well.

1. Go to **Configuration > Objects > Service Levels > Application Performance Score (APS)**.
2. Find and select the APS object.

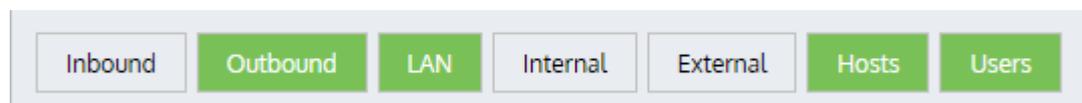
#### NOTE

The object name includes "Solution Center" and is suffixed with an application ID.

3. Select **Auto Baseline Period** and click **Start Baseline**.

### 0.12.12 Working with Application Performance charts

You can filter the data displayed on the page by toggling on and off the various charts. Click the buttons at the top of the page to switch between the views. If a button is green, the data appears on the page.



### Determining throughput values for specific points in time in the throughput chart

Hover your cursor over the chart. A data brush will appear showing average throughput for the specific point in time.

## Investigating reduction due to acceleration, edge cache, and bandwidth shaping

Toggle on the LAN-side reporting to overlay the LAN-side values on the WAN-side values. The difference between the LAN-side and WAN-side throughput lines indicates the amount of reduction that was achieved. Note that the total reduction in data volume as a percentage is shown under the chart.

## Investigating application usage by the top internal users and top internal hosts

Ensure the **Internal** button is toggled on and the **External** button is toggled off. The top hosts and top users (if configured) found on the LAN-side of the appliance are shown.

## Investigating application usage by the top external users and top external hosts

Ensure the **External** button is toggled on and the **Internal** button is toggled off. The top external hosts and top external users (if configured) found on the WAN-side of the appliance are shown.

### 0.12.13 Investigating a poor application performance score (APS)

From an application performance chart, click **Show details**. A new screen charts the measures contributing to the APS:

- » Network delay and normalized network delay – the amount of time for data to traverse the network (on the wire)
- » Server delay and normalized server delay – the amount of time for a server to respond to a request
- » Round trip time – the amount of time for data to travel from a device across a network and return
- » Jitter – a measure of the variability of Network Delay. We define it as one standard deviation of the Network Delay.
- » Inbound loss and outbound loss – the amount of data retransmitted

Inspect the charts to determine which attribute caused the poor APS score. For example, if the server delay measures are good but the network delay measures are bad, then you know that the network is to blame and perhaps you can do something about it. If the network delay measures are good but the server delay measures are bad, then you should have someone investigate why the server is performing poorly.

Note that if the baselining period was not typical, then the calculated thresholds may be overly high or low. For example, if the application was baselined on a weekend when there was very little traffic, then the thresholds may be much lower than would be expected when the network is in a typical use scenario. Similarly, if the application was base-lined during an extremely busy time, such as when most employees are watching an online CEO webcast, then the thresholds may be much higher than would be expected when the network is in a typical use scenario.

### 0.12.14 Investigating unusual performance

Zoom into the area of interest by opening the **Show data for** drop-down list and selecting one of the defined time spans. If necessary, use the slider controls below the chart to further refine the period. All the time series charts on this screen (inbound throughput, outbound throughput) will synchronize so that you can look for correlations in the data.

If any flat tops correspond to drops in application performance score, you likely have an issue that relates to a policy (or policies) that control the application. The policy environment may not guarantee the application sufficient bandwidth, or that less important apps, such as recreational apps, have too much bandwidth. To determine your best course of action you may need to look at the charts for other applications or application groups to see if the allowed bandwidth is appropriate.

### 0.12.15 Deleting an Application Performance report

The only way to delete a solution is through the command line. However, for some solutions you need to determine the solution ID from the Web UI before you can remove the solution.

1. Go to **Configuration > Objects > Service Levels > Application Performance Scores**.
2. Find the application in the **APS Name** column.

**NOTE**

The formatting of the name includes "Solution Center" and the ID. For example the CIFS APS object would be called CIFS Solution Center (208).

3. Make a note of the solution ID.
4. Open the CLI.
5. At the prompt, type `no solutionc <id>`. Examples:

- `no solutionc 208`
- `no solutionc VoIPPerformance`
- `no solutionc RecreationalTraffic`

## 0.13 Managing network traffic

As you monitor your traffic, you may recognize patterns of activity that you need to manage. The Exinda Appliance provides the ability to define policies, based on several criteria, that you can use to control, by either limiting or giving priority to, specified classes of traffic on your network.

- » You can create network objects to isolate particular parts of the network. Using network objects, you can create policies that apply to only specific branches, departments, classes of devices, and so on. A static network object can include one or more subnets and one or more IP addresses. For more information, refer to [Adding network objects](#) (page 178).
- » You can configure Active Directory to identify users and user groups on the network. By identifying users and user groups, you can create policies that apply to only specific users or user groups. For more information, refer to [Integrate with Active Directory](#) (page 553). To make specific network users or network user groups available to the policy configuration, they must be accordingly flagged to be used in the policies. See [Configure Network User Objects](#) to create a network object based on a network user. See [Configure Network User Group Objects](#) to create a network object based on a network user group.
- » You can configure VLANs to isolate particular parts of the network that is unrelated to their physical location. Using VLANs, you can create policies that apply to particular functions within the network. For more information, refer to [Configuring VLAN objects](#) (page 192).
- » You can create an application object if a new or custom application is not identified by the system. Configuring your own application objects allow you to create policies that apply to these specific applications. For more information, refer to [Adding application objects](#) (page 197).
- » You can create schedules to define time periods that more closely match your business. Configuring your own schedules allow you to create policies that change based on time periods that are important to your business. For more information, refer to [Configuring schedule objects](#) (page 215).
- » You can create adaptive response limit rules to automatically restrict a user's bandwidth once a set transfer limit has been reached within a specified period of time. For more information, refer to [Configuring adaptive response limits](#) (page 217).

---

|                                     |     |
|-------------------------------------|-----|
| 0.13.1 The Exinda policy tree ..... | 346 |
|-------------------------------------|-----|

---

---

|                                                      |     |
|------------------------------------------------------|-----|
| 0.13.2 Circuits .....                                | 349 |
| 0.13.3 Virtual Circuits .....                        | 356 |
| 0.13.4 Policies overview .....                       | 371 |
| 0.13.5 The optimizer wizard .....                    | 389 |
| 0.13.6 Calculating network performance metrics ..... | 393 |

---

## 0.13.1 The Exinda policy tree

All network behavior that you want to modify by using the Exinda Appliance is specified by policies in the optimizer. This includes traffic shaping, prioritization, acceleration, and packet marking. These policies are arranged hierarchically in a tree so that you can assign different policy rules to different types of traffic on your network. The hierarchy consists of circuits, virtual circuits, and policies.

### Role of circuits, virtual circuits, and policies

The policy tree is arranged as a three-level hierarchy. The top level of the hierarchy consists of circuits that partition the traffic by bridge. The second level of the hierarchy consists of virtual circuits that logically partitions each circuit. The third level of the hierarchy consists of policies that define which actions to take on the traffic.

Circuits relate to physical connections to the Exinda Appliance. For appliances with more than one bridge configured, you can create a separate circuit for each bridge or you can treat all bridges as one combined circuit. Typically, one circuit would be created for each physical link.

A circuit can contain one or more virtual circuits for the purpose of partitioning the traffic that falls in that circuit. The virtual circuit defines what traffic will be processed in this partition and how much bandwidth it is allowed. Each virtual circuit will have its own set of policy rules. The following are common use cases for virtual circuits.

- » A circuit needs to be partitioned into traffic destined for a remote office versus traffic destined for the Internet.
- » A circuit needs to be partitioned into traffic destined for particular branch offices.
- » You are a service provider and your circuit needs to be partitioned into traffic destined for each of your customers.
- » You have an asymmetric circuit and you want to control absolute bandwidth limits (e.g. To guarantee 100 kbps for VoIP traffic inbound and outbound).
- » You want to control and monitor different subnets in your network (e.g. One virtual circuit to control and monitor your wifi network, and another to control and monitor your servers, and another to control and monitor your back-office computers.)
- » You want to enforce fair sharing amongst the network hosts in the virtual circuit.
- » You want a certain number of hosts or connections to get preferential treatment such as having access to a high-bandwidth virtual circuit.

Each virtual circuit contains a set of policies that define what actions to perform on each type of traffic. Policies can shape the traffic by throttling traffic, blocking traffic, or prioritizing important traffic higher than unimportant traffic. Policies can also accelerate the traffic by using acceleration techniques or caching techniques. Policies can also merely monitor the traffic. The following are common uses for policies.

- » Policies to throttle or block non-critical traffic, such as P2P traffic.
- » Policies to protect business-critical traffic, such as Mail or Database traffic.
- » Policies to protect and prioritize business-critical latency sensitive traffic, such as voice traffic.

- » Policies to throttle a particular user or host, such as throttling John's access to streaming videos.
- » Policies to protect a particular subnet, such as a server or set of servers.

Circuits, virtual circuits, and policies can be managed by selecting various options from the drop-down to the right of the respective item.

To learn more about circuits, virtual circuits, or policies, see [Circuits](#), [Virtual Circuits](#), and [Policies](#).

## How traffic is evaluated against the policy tree

Traffic is evaluated against the hierarchical policy tree in a top-down manner. That is, the traffic is evaluated against the circuits in top-down order to determine which circuit will be handling the traffic. Once the appropriate circuit is determined, the traffic is evaluated against that circuit's virtual circuits in top-down order to determine which virtual circuit will be handling the traffic. Once the appropriate virtual circuit is determined, the traffic is evaluated against that virtual circuit's policies in top-down order to determine which policy will be handling the traffic. Any given packet will only be handled by one circuit, one virtual circuit, and one policy.

### NOTE

Special care is needed when creating virtual circuits based on limiting the number of active connections, or limiting the number of active hosts, or providing fair sharing between the active hosts. When a connection or host limit is reached, it will no longer match any incoming traffic. Therefore, connections or hosts that arrive later will be evaluated against the remaining virtual circuits in the circuit. You should ensure that the overflow connections or overflow hosts are handled according to your business needs by creating a virtual circuit immediately after the virtual circuit based on limits that deals with that traffic appropriately.

Later, once some of the active connections or active hosts in the virtual circuit terminate, the virtual circuit will be used to match new traffic again.

### EXAMPLE

Consider the scenario where the first circuit is defined to match bridge br10, the first virtual circuit within that circuit is defined to match the subnet for a particular branch site Springfield, and the first policy within that virtual circuit is defined to match P2P traffic.

Traffic arrives. If the traffic arrives on br10, then it is evaluated against that circuit's first virtual circuit. If the traffic matches against the Springfield subnet, then it is evaluated against its first policy. If the traffic matches P2P, then the action within the P2P policy is taken.

In each case, if the traffic did not match a given circuit, then the traffic would be evaluated against the next circuit. If the traffic did not match a given virtual circuit in the circuit that was matched, then the traffic would be evaluated against the next virtual circuit. If the traffic did not match a given policy in the virtual circuit that was matched, then the traffic would be evaluated against the next policy.

|                                                                 |                      |                                                                                                                  | Operations            |
|-----------------------------------------------------------------|----------------------|------------------------------------------------------------------------------------------------------------------|-----------------------|
| <b>Circuit 10 - Default (10000 kbps)</b>                        |                      |                                                                                                                  | --Actions--           |
| <b>Virtual Circuit 10 - WAN inbound (10000 kbps from 'ALL')</b> |                      |                                                                                                                  | --Actions--           |
| <input checked="" type="checkbox"/>                             | 10                   | P2P - Choke 1%-3% (Optimize 1% - 3%, Priority 10)                                                                | --Actions--           |
| <input checked="" type="checkbox"/>                             | 20                   | Recreational - Limit Low 2%-10% (Optimize 2% - 10%, Priority 10)                                                 | --Actions--           |
| <input checked="" type="checkbox"/>                             | 30                   | Software Updates - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)                                        | --Actions--           |
| <input checked="" type="checkbox"/>                             | 40                   | Voice - Guarantee Critical 15%-100% (Optimize 15% - 100%, Priority 1)                                            | --Actions--           |
| <input checked="" type="checkbox"/>                             | 50                   | Interactive and Secure - Guarantee High 10%-100% (Optimize 10% - 100%, Priority 3)                               | --Actions--           |
| <input checked="" type="checkbox"/>                             | 60                   | Thin Client - Guarantee High 10%-100% (Optimize 10% - 100%, Priority 3)                                          | --Actions--           |
| <input checked="" type="checkbox"/>                             | 70                   | Files - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)                                                   | --Actions--           |
| <input checked="" type="checkbox"/>                             | 80                   | Web - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)                                                     | --Actions--           |
| <input checked="" type="checkbox"/>                             | 90                   | Mail - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)                                                    | --Actions--           |
| <input checked="" type="checkbox"/>                             | 100                  | Database - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)                                                | --Actions--           |
| <input checked="" type="checkbox"/>                             | 200                  | ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)                                                     | --Actions--           |
| Order:                                                          | <input type="text"/> | Policy: ALL - Accelerate                                                                                         | Add To 'WAN inbound'  |
| <a href="#">Create New Policy...</a>                            |                      |                                                                                                                  |                       |
| <b>Virtual Circuit 15 - WAN outbound (10000 kbps to 'ALL')</b>  |                      |                                                                                                                  | --Actions--           |
| <input checked="" type="checkbox"/>                             | 10                   | P2P - Choke 1%-3% (Optimize 1% - 3%, Priority 10)                                                                | --Actions--           |
| <input checked="" type="checkbox"/>                             | 20                   | Recreational - Limit Low 2%-10% (Optimize 2% - 10%, Priority 10)                                                 | --Actions--           |
| <input checked="" type="checkbox"/>                             | 30                   | Software Updates - Guarantee Low 5%-100% - Accelerate (Optimize 5% - 100%, Priority 6, Application Acceleration) | --Actions--           |
| <input checked="" type="checkbox"/>                             | 40                   | Voice - Guarantee Critical 15%-100% (Optimize 15% - 100%, Priority 1)                                            | --Actions--           |
| <input checked="" type="checkbox"/>                             | 50                   | Interactive and Secure - Guarantee High 10%-100% (Optimize 10% - 100%, Priority 3)                               | --Actions--           |
| <input checked="" type="checkbox"/>                             | 60                   | Thin Client - Guarantee High 10%-100% (Optimize 10% - 100%, Priority 3)                                          | --Actions--           |
| <input checked="" type="checkbox"/>                             | 70                   | Files - Guarantee Med 8%-100% - Accelerate (Optimize 8% - 100%, Priority 4, Application Acceleration)            | --Actions--           |
| <input checked="" type="checkbox"/>                             | 80                   | Web - Guarantee Med 8%-100% - Accelerate (Optimize 8% - 100%, Priority 4, Application Acceleration)              | --Actions--           |
| <input checked="" type="checkbox"/>                             | 90                   | Mail - Guarantee Low 5%-100% - Accelerate (Optimize 5% - 100%, Priority 6, Application Acceleration)             | --Actions--           |
| <input checked="" type="checkbox"/>                             | 100                  | Database - Guarantee Med 8%-100% - Accelerate (Optimize 8% - 100%, Priority 4, Application Acceleration)         | --Actions--           |
| <input checked="" type="checkbox"/>                             | 200                  | ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)                                                     | --Actions--           |
| Order:                                                          | <input type="text"/> | Policy: ALL - Accelerate                                                                                         | Add To 'WAN outbound' |
| <a href="#">Create New Policy...</a>                            |                      |                                                                                                                  |                       |
| <a href="#">Create New Virtual Circuit...</a>                   |                      |                                                                                                                  |                       |
| <a href="#">Create New Circuit...</a>                           |                      |                                                                                                                  |                       |

Screenshot 152: The policy tree

## When traffic is not caught by the policy tree

There are Auto-Catch All circuits, Auto-Catch All virtual circuits and Auto-Catch All policies for any traffic that does not get caught by your configured circuits, virtual circuits, and policies. Auto-Catch All do not have any policies applied to it, but they are monitored and will be shown on the virtual circuit graph.

Consider a configuration where you have two bridges on the appliance. If you define a circuit for the first bridge but not the second, then any traffic that comes in on the second bridge will fall into the Auto Catch-All circuit and will not have any of your policies applied to it.

Similarly consider a configuration where you have a circuit that will capture traffic for your entire office. You define multiple virtual circuits for each subnet that you are interested in, but your virtual circuits do not cover your entire office. For example, you have virtual circuits for your WiFi network, your servers, your staff desktop network, but perhaps you forgot your printers. The system will automatically define a virtual circuit for the remaining traffic called Auto Catch-all. Any traffic that falls into a circuit, but not into any of the circuit's virtual circuits, will fall into the Auto Catch-all virtual circuit and will not have any of your policies applied to it.

Finally consider a configuration where you have a virtual circuit that captures all the traffic for a particular virtual circuit, such as your WiFi network. You define policies for each type of traffic that you want to control or protect or accelerate, but do not include all the traffic on your WiFi network. For example, you have policies for P2P, streaming, and web, but you do not have policies for software updates. The system will automatically define a policy for the remaining traffic called Auto

Catch-all. Any traffic that falls into a virtual circuit, but not into any of policies, will fall into the Auto Catch-all policy and will not have any of your policies applied to it.

Consider the opposite scenario where you have more than one circuit that could capture the same traffic. For example, you have two circuits that could both capture br10 traffic. The traffic will be evaluated against the virtual circuits and policies within the first circuit. If the traffic does not match any virtual circuits and policies in the first circuit, then it will attempt to match the traffic to the virtual circuits and policies in the second circuit. If the traffic does not match any of the virtual circuits and policies in any of the matching circuits, then the traffic will fall into the Auto Catch-All virtual circuit and Auto Catch-All policy in the last matching circuit and will not have any of your policies applied to it.

### BEST PRACTICE

It is a best practice to explicitly create your own "catch-all" policy so that you can apply the appropriate control and prioritization.

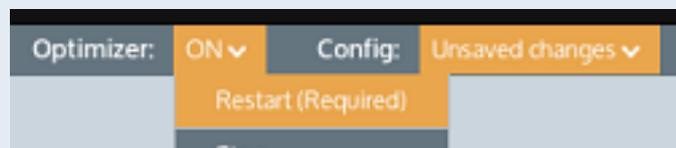
If traffic is not caught by a circuit or by a virtual circuit or by a policy, it will show up in the real-time monitor as an **Auto Catch-all** circuit and an **Auto Catch-all** virtual circuit and an **Auto Catch-all** policy. The Auto Catch-all circuits and virtual circuits will also be shown on the Virtual Circuit monitor report.

## Enabling policies

Policies that are enabled are shown with a green checkmark in the policy tree. Policies that are disabled are shown with a red x. You can toggle the enable state by clicking the checkmark or x.

### NOTE

Whenever the Optimizer is modified, it must be restarted before the changes will take effect on your traffic. Restart the Optimizer, by selecting the Optimizer menu in the status bar at the top of the screen.



Screenshot 153: Web interface menu

## Related Topics

- » [Adding a Policy to the Policy Tree](#)
- » [Policy Ordering Matters](#)
- » [Re-using Policies from the Library](#)

## 0.13.2 Circuits

Circuits define physical connections to the WAN/Internet. For appliances with more than one bridge configured, you can bind each bridge to a separate circuit or you can treat all bridges as one combined circuit. Typically, when the appliance is placed in line, one circuit would be created for each physical link.

When defining circuits, ensure each bridge is bound to a circuit; otherwise network traffic on any bridge that is not bound to a circuit, will not be affected by policy and will be monitored in a catch-all circuit.

**NOTE**

- » For version 7.0.2 and earlier: A circuit can only be bound to one bridge, or all bridges (but not greater than 1 bridge and less than all bridges) or one WCCP interface.
- » For version 7.0.2 Update 1 and later: A circuit can be bound to any number of bridges or one WCCP interface.

| Add New Circuit                        |                                      |
|----------------------------------------|--------------------------------------|
| Circuit Number                         | <input type="text" value="20"/>      |
| Circuit Name                           | <input type="text"/>                 |
| Inbound Bandwidth                      | <input type="text"/> kbps            |
| Outbound Bandwidth                     | <input type="text"/> kbps            |
| Attach to Bridge/Out-of-path Interface | <input type="button" value="ALL"/> ▾ |

Screenshot 154: Circuit form for versions 7.0.2 and earlier

**Add New Circuit**

|                                        |                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Circuit Number                         | <input style="width: 80%;" type="text" value="30"/>                                                                                                                                                                                                                                                                                                                    |
| Circuit Name                           | <input style="width: 100%;" type="text"/>                                                                                                                                                                                                                                                                                                                              |
| Inbound Bandwidth                      | <input style="width: 80%;" type="text"/> kbps                                                                                                                                                                                                                                                                                                                          |
| Outbound Bandwidth                     | <input style="width: 80%;" type="text"/> kbps                                                                                                                                                                                                                                                                                                                          |
| Attach to Bridge/Out-of-path Interface | <div style="background-color: #e0e0e0; padding: 5px; border: 1px solid #ccc;"> <b>Select bridge(s) to associate with this circuit</b> <ul style="list-style-type: none"> <li><input type="checkbox"/> ALL</li> <li><input checked="" type="checkbox"/> br10</li> <li><input checked="" type="checkbox"/> br12</li> <li><input type="checkbox"/> WCCP</li> </ul> </div> |



Screenshot 155: Circuit form for versions 7.0.2 Update 1 and later

Circuits are part of the policy tree. To learn how circuits, virtual circuits, and policies work together, see [Policy Tree](#).

### Can a bridge be bound to more than one circuit?

A bridge can be assigned to more than one circuit. Once a type of traffic matches a virtual circuit and policy, it will not be processed by later circuits, virtual circuits, and policies. If the type of traffic does not match any policy in the first circuit, then it will attempt to match the policy in the second circuit.

For example, consider where circuits, virtual circuit, and policy are defined as:

- » Circuit A = br10
  - Virtual Circuit
    - Policy for P2P
- » Circuit B = All bridges
  - Virtual Circuit
    - Policy for Mail

In this case, br10 is assigned to both circuit A & B. However, br10 is bound to circuit A first.

- » If the traffic is P2P, then that traffic type will be processed by Circuit A and will not be processed by Circuit B.
- » If the traffic is mail, then that traffic type will not match Circuit A. It will attempt to match policy in Circuit B. In this case, it will match the policy for mail.

» If the traffic is not P2P or mail, then that traffic type will not match Circuit A or Circuit B and therefore will fall into an auto-catch-all circuit, auto-catch-all virtual circuit, and auto-catch-all policy.

## Creating a circuit in the Exinda Web UI

Use the following instructions to create a new circuit.

1. Go to **Configuration > Traffic Policies > Optimizer > Optimizer**, then click the **Create New Circuit...** link at the bottom of the page or edit an existing circuit.
2. In the **Circuit Number** field, type the ranking order of the circuit, relative to other circuits.
3. Enter a name for the circuit.
4. Identify the available **Inbound** and **Outbound** bandwidth of the circuit.
  - If the circuit is symmetric, the inbound and outbound bandwidth values should be the same.
  - If the circuit is asymmetric, the inbound and outbound bandwidth values will be different.
5. Specify the bridge or out-of-path interface to which to bind the circuit. All bridges, individual bridge names, policy-based routing interfaces, and WCCP interfaces are available.
6. Click **Add New Circuit**.

## Creating a circuit in the EMC

Circuits define physical connections to the WAN or the Internet. A circuit defines the inbound and outbound bandwidth and the named circuit type.

On an Exinda Appliance, a circuit specifies the named bridge (or bridges) to which it is bound. In the Exinda Management Center, the binding to bridges is through a designated Circuit Type. Circuit Types represent the intended use of a circuit. So you can configure a circuit for multiple appliances without requiring the bridges on the appliances to have the same name. This is favorable where the number of bridges or names of bridges or the cabling of the bridges is not consistent across the appliances.

If multiple bridges on an appliance are mapped to the same Circuit Type, then all those bridges will be bound to the single circuit in the Policy Tree that is configured with that Circuit Type. To learn how Circuit Types are used to determine which circuits are sent to the appliances, read [Circuit Types](#).

Circuits can be created in the Configuration Library directly and then later assigned to an Optimizer Policy Tree, Circuits can also be created in an Optimizer Policy Tree, from where the configuration is saved to the Configuration Library so that it can be used elsewhere. The circuits within a tenant must have unique names. When a circuit is modified, all uses of it are modified.

Circuits are part of the Optimizer Policy Tree. To learn how circuits, Virtual Circuits, policy sets, and policy rules work together, see [Policy Tree](#).

### Where do I find circuits?

Circuit library items can be found in **Library > Circuits**. Also, the circuits that are applied to appliances can also be found in the policy trees for each appliance group. Go to the **(desired appliance group) > Optimizer Policy Tree**.

### To create a circuit in the Configuration Library

1. Go to **Library > Circuits**.
2. Click **Create new circuit in the library ...**

3. In the **Name** section, key-in the name of the circuit. The name must be unique in the tenant.
4. In the **Bandwidth** section, key-in the **Inbound Bandwidth** and the **Outbound Bandwidth**. The bandwidths can be specified in kbps, Mbps, or Gbps.
5. In the **Bind to Circuit Type** section, select the **Circuit Type**.

#### NOTE

The circuit type represents the purpose of the circuit; it is a user created object. If the desired circuit type does not exist, you can click **Create new circuit type in the library** to create it. For more information, refer to [Circuit types in the EMC](#) (page 423).

6. Click **Create**. The circuit appears in the library list.

#### To create a new circuit directly in the Policy Tree

1. Go to the **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Create new circuit**.
3. In the **Name** section, key-in the name of the circuit. The name must be unique in the tenant.
4. In the **Bandwidth** section, key-in the **Inbound Bandwidth** and the **Outbound Bandwidth**. The bandwidths can be specified in kbps, Mbps, or Gbps.
5. In the **Bind to Circuit Type** section, select the **Circuit Type**.

#### NOTE

The circuit type represents the purpose of the circuit; it is a user created object. If the desired circuit type does not exist, you can click **Create new circuit type in the library** to create it. For more information, refer to [Circuit types in the EMC](#) (page 423).

6. Click **Save**. The circuit is added to the Policy Tree and is also saved to the configuration library.

#### To add a circuit to a Policy Tree

1. Go to the **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add circuit from library**.
3. Select the desired circuit from the drop-down list.

## Circuits

Circuits define physical connections to the WAN/Internet. For appliances with more than one bridge configured, you can bind each bridge to a separate circuit or you can treat all bridges as one combined circuit. Typically, when the appliance is placed in line, one circuit would be created for each physical link.

When defining circuits, ensure each bridge is bound to a circuit; otherwise network traffic on any bridge that is not bound to a circuit, will not be affected by policy and will be monitored in a catch-all circuit.

#### NOTE

- » For version 7.0.2 and earlier: A circuit can only be bound to one bridge, or all bridges (but not greater than 1 bridge and less than all bridges) or one WCCP interface.
- » For version 7.0.2 Update 1 and later: A circuit can be bound to any number of bridges or one WCCP interface.

**Add New Circuit**

|                                        |                                                     |
|----------------------------------------|-----------------------------------------------------|
| Circuit Number                         | <input style="width: 60%;" type="text" value="20"/> |
| Circuit Name                           | <input style="width: 60%;" type="text"/>            |
| Inbound Bandwidth                      | <input style="width: 60%;" type="text"/> kbps       |
| Outbound Bandwidth                     | <input style="width: 60%;" type="text"/> kbps       |
| Attach to Bridge/Out-of-path Interface | <input type="button" value="ALL"/>                  |



Screenshot 156: Circuit form for versions 7.0.2 and earlier

**Add New Circuit**

|                                        |                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Circuit Number                         | <input style="width: 60%;" type="text" value="30"/>                                                                                                                                                                                                                                                                                                                                                     |
| Circuit Name                           | <input style="width: 60%;" type="text"/>                                                                                                                                                                                                                                                                                                                                                                |
| Inbound Bandwidth                      | <input style="width: 60%;" type="text"/> kbps                                                                                                                                                                                                                                                                                                                                                           |
| Outbound Bandwidth                     | <input style="width: 60%;" type="text"/> kbps                                                                                                                                                                                                                                                                                                                                                           |
| Attach to Bridge/Out-of-path Interface | <div style="border: 1px solid #ccc; padding: 5px;"> <div style="background-color: #e0e0e0; padding: 2px;"><b>Select bridge(s) to associate with this circuit</b></div> <ul style="list-style-type: none"> <li><input type="checkbox"/> ALL</li> <li><input checked="" type="checkbox"/> br10</li> <li><input checked="" type="checkbox"/> br12</li> <li><input type="checkbox"/> WCCP</li> </ul> </div> |



Screenshot 157: Circuit form for versions 7.0.2 Update 1 and later

Circuits are part of the policy tree. To learn how circuits, virtual circuits, and policies work together, see [Policy Tree](#).

## Can a bridge be bound to more than one circuit?

A bridge can be assigned to more than one circuit. Once a type of traffic matches a virtual circuit and policy, it will not be processed by later circuits, virtual circuits, and policies. If the type of traffic does not match any policy in the first circuit, then it will attempt to match the policy in the second circuit.

For example, consider where circuits, virtual circuit, and policy are defined as:

- » Circuit A = br10
  - Virtual Circuit
    - Policy for P2P
- » Circuit B = All bridges
  - Virtual Circuit
    - Policy for Mail

In this case, br10 is assigned to both circuit A & B. However, br10 is bound to circuit A first.

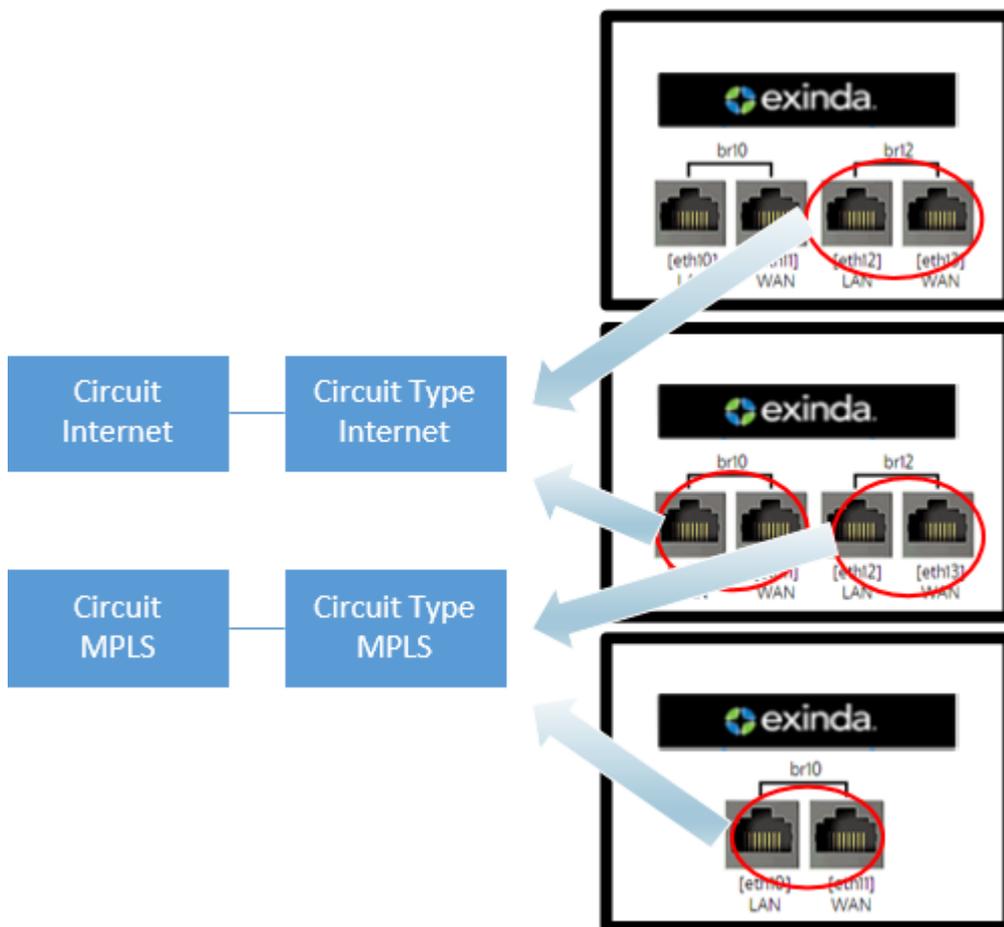
- » If the traffic is P2P, then that traffic type will be processed by Circuit A and will not be processed by Circuit B.
- » If the traffic is mail, then that traffic type will not match Circuit A. It will attempt to match policy in Circuit B. In this case, it will match the policy for mail.
- » If the traffic is not P2P or mail, then that traffic type will not match Circuit A or Circuit B and therefore will fall into an auto-catch-all circuit, auto-catch-all virtual circuit, and auto-catch-all policy.

## Circuit types in the EMC

Circuit Types are used to identify the purpose of the circuit and the appliance bridges. When the purpose of the circuit and an appliance bridge align, then the circuit is bound to that bridge in the configuration that is sent to the appliances. Note that Circuit Types do not exist on appliances. Circuit Types abstract the binding between the circuits and the appliance bridges. Therefore, the appliances can be treated similarly even when they do not have the same number of bridges and when they are not connected in the same way.

Circuit Types are defined in the Configuration Library and are used by circuits and appliance bridge-to-circuit type mappings.

The following example will further clarify the concept of circuit types:



Screenshot 158: Circuit types example

Consider the following in the example above:

- » Bridge **br12** from first appliance and bridge **br10** on second appliance are both bound to circuit type "Internet". Since, this circuit type "Internet" is tied to the circuit "Internet", all the policies within the 'Internet' circuit are applied to bridge br12 on the first appliance and br10 on the second appliance.
- » Bridge **br12** on the second appliance and bridge **br10** on third appliance are bound to circuit type "MPLS". Since, this circuit type 'MPLS' is tied to the circuit "MPLS", all the policies within the 'MPLS' circuit are applied to bridge br12 on the second appliance and br10 on the third appliance.

The **Circuit Type** library items can be found in **Library > Circuit Types**.

To assign the Circuit Type to a circuit, see [Circuits](#).

To assign Circuit Types to appliance bridges, see [Bridge/Circuit Type Mapping](#).

### 0.13.3 Virtual Circuits

Virtual circuits are created within circuits in the policy tree and are used to logically divide or partition the circuit. The virtual circuit defines what traffic will be processed in this partition and how much bandwidth it is allowed. The virtual circuit can enforce fair sharing amongst the network hosts. Traffic is evaluated against the definition of the virtual circuit. Traffic that does not fall within the virtual circuit is evaluated by the next virtual circuit and so on. Each virtual circuit will have its own set of policy rules.

The following are common use cases for virtual circuits.

- » Virtual circuits are configured to allocate bandwidth or specify policy rules for each branch office or service provider's customer.
- » Separate virtual circuits are configured for WAN data and for Internet data.
- » Separate virtual circuits are configured for inbound traffic and for outbound traffic. Virtual circuits cannot be asymmetric, so this technique will accommodate the need for asymmetric partitioning of the circuit.
- » Separate virtual circuits are configured for different aspects of your network for monitoring purposes. For instance, one virtual circuit could be configured to control and monitor your wifi network, and another to control and monitor your servers, and another to control and monitor your back-office computers.
- » A virtual circuit can enforce fair sharing amongst the network hosts in the virtual circuit.
- » A virtual circuit can ensure that a limited number of hosts or connections get preferential treatment.

A virtual circuit can partition the circuit by filtering the traffic by time of day, by VLAN range, by subnets or hosts, by particular application or application group, by traffic direction, and by capping the number of active connections or capping the number of active hosts within the virtual circuit. Any combination of these filters can be applied. For example, you can create a virtual circuit such that a particular branch (or subnet) is allowed a certain set of policies for inbound traffic (as direction) during off work hours (using a schedule).

A virtual circuit specifies its desired bandwidth either as kbps or as a percentage of its parent circuit. When the sum of the desired bandwidths for all the virtual circuits within a circuit exceeds the circuit's bandwidth, the circuit is oversubscribed. Each virtual circuit specifies how it would like to deal with oversubscription. That is, either let the system automatically share the bandwidth amongst the virtual circuits, or specify a minimum bandwidth that is required.

Additionally, a virtual circuit can enable fair sharing amongst hosts in the virtual circuit. When fair sharing is enabled, the virtual circuit is called a dynamic virtual circuit and a further level of traffic shaping is introduced. Traffic is first shaped at the host level, then at the policy level. The bandwidth allocated will be the minimum of the two levels.

A virtual circuit can provide preferential treatment to a limited number of active hosts or to a limited number of active connections.

#### **NOTE**

When configuring a dynamic virtual circuit, the system will not allow the per host bandwidth to be less than 10 kbps, in which case the number of allowed hosts is calculated to be the virtual circuit bandwidth/10 kbps. Any hosts beyond the limit are then evaluated against later virtual circuits in the policy tree.

There is a system limit of 32,500 hosts that can fall into each dynamic virtual circuit. This may occur if the virtual circuit has more than 300 Mbps of bandwidth. When this limit is exceeded, hosts fall into the next applicable virtual circuit.

#### **BEST PRACTICE**

It is a best practice to create an overflow virtual circuit immediately after a virtual circuit with a connection limit or a host limit, to capture the connections or hosts that were excluded.

| Add New Virtual Circuit                                                                      |                                                                            |
|----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number                                                                       | 10 . 35                                                                    |
| Virtual Circuit Name                                                                         |                                                                            |
| Schedule                                                                                     | ALWAYS                                                                     |
| Bandwidth Options                                                                            |                                                                            |
| Virtual Circuit Bandwidth                                                                    | <input type="text"/> %                                                     |
| Oversubscription                                                                             | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit                                                                      | <input type="checkbox"/>                                                   |
| Connection Options                                                                           |                                                                            |
| Connection Limit                                                                             | <input type="text"/>                                                       |
| Filter Options                                                                               |                                                                            |
| VLAN Object                                                                                  | ALL                                                                        |
| Network Object                                                                               | ---Network Objects---                                                      |
| Application                                                                                  | ALL                                                                        |
| Direction                                                                                    | Both                                                                       |
| <input type="button" value="Add New Virtual Circuit"/> <input type="button" value="Cancel"/> |                                                                            |

Screenshot 159: Setting options for a virtual circuit.

Network objects are typically used when virtual circuits are created for specific branch office locations, or other subsets of the network, or user groups. Each branch office location or user group would be represented by a static network object or a dynamic network object (such as an Active Directory group). A default network object, private net, exists which defines all non-routable subnets. This can be used to create a virtual circuit for all WAN data.

The direction is used to ensure that the virtual circuit only captures traffic in a certain direction. This is useful for asymmetric circuits, as these generally require that at least two virtual circuits are defined - one for the inbound bandwidth and one for the outbound bandwidth.

Virtual circuits are part of the policy tree. To learn how circuits, virtual circuits, and policies work together, see [Policy Tree](#).

### Related Topics

- » [Creating a new Virtual Circuit](#)
- » [How to handle oversubscription in a Virtual Circuit](#)
- » [Configuring a Virtual Circuit as a Dynamic Virtual Circuit](#)
- » [Sharing bandwidth equally](#)
- » [Ensuring minimum bandwidth when sharing equally](#)
- » [Capping bandwidth usage per host](#)
- » [Capping bandwidth usage per host with minimum bandwidth](#)

- » [Limiting the number of hosts sharing bandwidth](#)
- » [Troubleshooting Virtual Circuits](#)

## Creating a virtual circuit

Use the following instruction to create a new virtual circuit (VC). You use VCs to manage bandwidth through the Exinda Appliance.

### Before you begin...

If you need to manage oversubscription to the VC, see the following:

- » [How to handle oversubscription in a Virtual Circuit](#)

As you create the new VC, you may want to employ the Dynamic Options. Depending on the configuration, the following can give some insight into the best way to achieve your goals:

- » [Configuring a Virtual Circuit as a Dynamic Virtual Circuit](#)
- » [Sharing bandwidth equally among as many hosts as possible](#)
- » [Ensuring minimum bandwidth when sharing equally among as many hosts as possible](#)
- » [Capping bandwidth usage per host](#)
- » [Capping bandwidth usage per host with minimum bandwidth allowed](#)
- » [Limiting the number of hosts sharing bandwidth](#)

### Creating a virtual circuit in the Exinda Web UI

1. Go to **Configuration > Traffic Policies > Optimizer**.
2. If not already open, select the **Optimizer** tab.
3. Click the **Create New Virtual Circuit...** link at the bottom of each circuit or edit an existing virtual circuit.
4. In the **Virtual Circuit Number** field, type a number that will sort the virtual circuit in the policy tree.
5. In the Virtual Circuit Name field, type a suitable name for the VC.
6. If necessary, from the **Schedule** drop-down, select a schedule that defines a particular time period, the default is ALWAYS.
7. Enter the desired **Virtual Circuit Bandwidth** as kbps or as a percentage of it's parent circuit.

#### NOTE

The desired bandwidth for a single virtual circuit must not exceed it's parent circuit's bandwidth in either direction.

8. Specify how to handle **Oversubscription** . See the "Before you begin..." section above for a link to more information about managing oversubscriptions.

- **Automatic**– The VC shares bandwidth proportionally with other "Automatic" VCs.
- **Manual**– The VC is guaranteed a specified amount of bandwidth.

### IMPORTANT

Any VCs with manually assigned bandwidth take priority over VCs with automatically assigned bandwidth. This means that "Automatic" VCs share the bandwidth that remains after all "Manual" VCs have received their guaranteed bandwidth. If no bandwidth remains, the "Automatic" VCs receive nothing.

9. If necessary, select the **Dynamic Virtual Circuit** option. The "Dynamic Options" open. See the "Before you begin..." above for links to topics that explain how to use the options.

10. If necessary, specify a **Connection Limit**. This limits the number of active connections in the virtual circuit. Leave this field blank if you do not want to cap the number of connections.

### NOTE

Once a virtual circuit connection limit has been reached, the VC will no longer match any incoming traffic. Therefore, connections that arrive later will be evaluated against the remaining virtual circuits in the Policy Tree. As active connections terminate, the VC then evaluates other connections.

### BEST PRACTICE

To prevent traffic being entirely excluded, it is a best practice to create a further virtual circuit to capture traffic rejected by the VC.

11. Complete the following **Filter Options**, as needed.

- Specify a **VLAN Object** to filter which traffic falls into the virtual circuit. The default is **ALL**.
- Specify a **Network Object** to filter which subnet, hosts, users, or user groups falls into the virtual circuit.
- Specify an **Application** to filter which application or application group falls into the virtual circuit. The default is **ALL**.

12. Specify the **Direction** of traffic to capture in the virtual circuit. The direction options are:

### NOTE

The direction is relative to the LAN. Consider an example where a network object and a direction is specified:

- **Both**– both inbound and outbound traffic
- **Inbound**– inbound from the WAN only
- **Outbound**– outbound to the WAN only

| Network Object                           | Direction | Captured Traffic                                                                               |
|------------------------------------------|-----------|------------------------------------------------------------------------------------------------|
| 'Internal' or 'External' network objects | Both      | Only inbound and outbound traffic to and from the subnets defined by the network object.       |
| 'Internal' network object                | Inbound   | Only inbound traffic to the subnets defined as 'internal' by the network object.               |
| 'External' network object                | Inbound   | Only inbound traffic to the LAN from the subnets defined as 'external' by the network object.  |
| 'Internal' network object                | Outbound  | Only outbound traffic from the subnets defined as 'internal' by the Network Object.            |
| 'External' network object                | Outbound  | Only outbound traffic from the LAN to the subnets defined as 'external' by the Network Object. |

13. Click **Add New Circuit**. The virtual circuit appears in the policy tree.
14. If necessary, in the policy tree, change the ranking of the VC. This ensures the VC is handled in the correct order.

### To delete a policy from the system

To delete a policy from the system, it must be deleted from the policy library. Go to **Configuration > Traffic Policies > Optimizer > Policies**.

#### Related Topic

### [Troubleshooting Virtual Circuits](#)

### [Creating a virtual circuit in the EMC](#)

You have the option of creating a virtual circuit within the Configuration Library first, and then later assigning it to a circuit in the Optimizer Policy Tree, or you can create the virtual circuit directly within the Optimizer Policy Tree, which also saves it as a Library item. The virtual circuits within a tenant must have unique names. When a virtual circuit is modified, all instances of its use are modified.

To learn how circuits, virtual circuits, policy sets, and policy rules work together, see [Policy Tree](#).

### [Where do I find Virtual Circuits?](#)

Virtual Circuit library items can be found in **Library > Virtual Circuits**. Also, the virtual circuits that will be sent to appliances are found in the policy trees for each appliance group. Go to the **(desired appliance group) > Optimizer Policy Tree**.

### [To create a Virtual Circuit in the Configuration Library](#)

1. Go to **Library > Virtual Circuits**.
2. Click **Create new virtual circuit**.
3. In the **Name** section, key-in a name for the virtual circuit. The name must be unique within the tenant.

#### NOTE

If you want, you can leave the EMC to define a name for you, It does this based on the configuration of the virtual circuit.

4. In the **Filter** section, select the combination of filters to apply to the virtual circuit. Optionally, type a value to limit the number of connections at one time on this virtual circuit.

#### NOTE

The virtual circuit can partition the circuit by filtering the traffic based on these filters. You can apply any combination of these filters. Defined network object library items appear in the [Network Object](#) list, and you can also choose filters from the pre-defined application groups.

5. In the **Bandwidth** section, type the desired bandwidth for this virtual circuit.
6. Also specify how to share bandwidth with other virtual circuits when there is insufficient bandwidth due to over subscription..

**NOTE**

The bandwidth can be specified in %, Kbps, Mbps or Gbps.

7. In the **Dynamic Virtual Circuit** section, set the options that provide the control you need.

**NOTE**

See For more information, refer to [Dynamic Virtual Circuits](#) (page 428). for more information about configuring dynamic virtual circuits.

8. In the **Schedule** section, set the time values for when the virtual circuit will be enforced.

**NOTE**

Options in the list are determined by the library [Schedules](#) category

9. Click the **Create** button. The virtual circuit is added to the Virtual Circuits Library category.

### To create a new Virtual Circuit directly in the Optimizer Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Create new virtual circuit**.
3. In the **Name** section, key-in a name for the virtual circuit. The name must be unique within the tenant.

**NOTE**

If you want, you can leave the EMC to define a name for you, It does this based on the configuration of the virtual circuit.

4. In the **Filter** section, select the combination of filters to apply to the virtual circuit. Optionally, type a value to limit the number of connections at one time on this virtual circuit.

**NOTE**

The virtual circuit can partition the circuit by filtering the traffic based on these filters. You can apply any combination of these filters. Defined network object library items appear in the [Network Object](#) list, and you can also choose filters from the pre-defined application groups.

5. In the **Bandwidth** section, type the desired bandwidth for this virtual circuit.

6. Also specify how to share bandwidth with other virtual circuits when there is insufficient bandwidth due to over subscription..

**NOTE**

The bandwidth can be specified in %, Kbps, Mbps or Gbps.

7. In the **Dynamic Virtual Circuit** section, set the options that provide the control you need.

**NOTE**

See For more information, refer to [Dynamic Virtual Circuits](#) (page 428). for more information about configuring dynamic virtual circuits.

8. In the **Schedule** section, set the time values for when the virtual circuit will be enforced.

**NOTE**

Options in the list are determined by the library [Schedules](#) category

9. Click the **Create** button. The virtual circuit is added to the Circuit within the Optimizer Policy Tree and is also saved to the Configuration Library.

### To add a Virtual Circuit from the Library to a Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add Virtual Circuit from library**.
3. Select the desired virtual circuit from the drop-down list.

### Calculating virtual circuit oversubscription

Virtual circuits are oversubscribed when the sum of the virtual circuit bandwidths exceeds the bandwidth of the parent circuit. For example:

Circuit Bandwidth = 3Mbps

- » Virtual Circuit Bandwidth A = 2Mbps
- » Virtual Circuit Bandwidth B = 1Mbps
- » Virtual Circuit Bandwidth C = 1Mbps

This means, the sum of the three virtual circuits is 4Mbps, but the circuit bandwidth is only 3Mbps.

|                             |                                                                                                                                                                                                          |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VC A:<br>automatic          | Each VC gets:<br>desired bandwidth / sum of VC bandwidth * circuit bandwidth                                                                                                                             |
| VC B:<br>automatic          | » VC A = $2/4 * 3 = 1.5$ Mbps                                                                                                                                                                            |
| VC C:<br>automatic          | » VC B = $1/4 * 3 = 0.75$ Mbps                                                                                                                                                                           |
|                             | » VC C = $1/4 * 3 = 0.75$ Mbps                                                                                                                                                                           |
| VC A: manual<br>= 2 Mbps    | Each virtual circuit with manually set oversubscription bandwidth will get their guaranteed amount                                                                                                       |
| VC B: manual<br>= 0.5 Mbps  | » VC A = 2 Mbps                                                                                                                                                                                          |
| VC C: manual<br>= 0.5 Mbps  | » VC B = 0.5 Mbps                                                                                                                                                                                        |
|                             | » VC C = 0.5 Mbps                                                                                                                                                                                        |
| VC A:<br>automatic          | Each virtual circuit with manually set oversubscription bandwidth will get their guaranteed amount. Virtual circuits with automatic oversubscription calculations will share the remaining bandwidth as: |
| VC B:<br>automatic          | virtual circuit's desired bandwidth / total remaining bandwidth * (circuit's bandwidth - sum of manually specified oversubscription bandwidths)                                                          |
| VC C: manual<br>= 0.75 Mbps | » VC C = 0.75 Mbps                                                                                                                                                                                       |
|                             | » VC A = $2/3 * (3 - 0.75) = 1.5$ Mbps                                                                                                                                                                   |
|                             | » VC B = $1/3 * (3 - 0.75) = 0.75$ Mbps                                                                                                                                                                  |

## Related Topics

» [Creating a new Virtual Circuit](#)

### Configuring a virtual circuit as a dynamic virtual circuit

When defining a VC, you can configure it as a dynamic virtual circuit (DVC). DVCs have the flexibility to manage the available bandwidth based on the number of active hosts.

The screenshot shows the 'Add New Virtual Circuit' configuration form. The 'Dynamic Options' section is highlighted with a red border. The form includes the following fields and options:

- Add New Virtual Circuit**
  - Virtual Circuit Number: 10 . 30
  - Virtual Circuit Name: [Empty text box]
  - Schedule: ALWAYS [Dropdown arrow]
- Bandwidth Options**
  - Virtual Circuit Bandwidth: [Empty text box] % [Dropdown arrow]
  - Oversubscription:  Automatic,  Manual
  - Dynamic Virtual Circuit:
- Dynamic Options** (highlighted with a red border)
  - Per Host Bandwidth:  Automatically Share, [0] % [Dropdown arrow]
  - Per Host Max Bandwidth:  No Bursting Allowed, [0] % [Dropdown arrow]
  - Host Location: Internal [Dropdown arrow]
  - Max Hosts:  Auto, [0]
- Connection Options**
  - Connection Limit: [Empty text box]
- Filter Options**
  - VLAN Object: ALL [Dropdown arrow]
  - Network Object: ---Users & Groups--- [Dropdown arrow]
  - Application: ALL [Dropdown arrow]
  - Direction: Both [Dropdown arrow]

Buttons: Add New Virtual Circuit, Cancel

Screenshot 160: The Dynamic Options for a VC

DVCs can either be configured for fair sharing among the hosts, or they can be configured to limit the number of hosts to give those hosts preferential treatment.

For fair sharing, you fix the per host bandwidth and have the system calculate the number of allowed hosts. To limit the number of hosts, you fix the number of hosts and have the system calculate the per host bandwidth allowed to each.

Both methods can limit the number of hosts, but for when there is spare capacity, you can configure the VC to allow bursting when there are less than the allowed number of hosts. The DVC then allows each active host to gain more bandwidth. In the configuration, you set the VC to automatically calculate the per host bandwidth and the number of allowed hosts. The system then assigns bandwidth by dividing what is available to the virtual circuit by the number of active hosts.

#### Related Topics

- » [Sharing bandwidth equally among as many hosts as possible](#)
- » [Ensuring minimum bandwidth when sharing equally among as many hosts as possible](#)
- » [Capping bandwidth usage per host](#)
- » [Capping bandwidth usage per host with minimum bandwidth](#)
- » [Limiting the number of hosts sharing bandwidth](#)

#### Sharing bandwidth equally with a dynamic virtual circuit

You can use Dynamic Virtual Circuits to enforce fair sharing of bandwidth among the hosts, or to limit the number of hosts on the circuit to ensure that those hosts get preferential treatment.

- » For fair sharing, you must specify how you would like the bandwidth in the virtual circuit to be shared among the hosts. You can fix the per host bandwidth and have the system calculate the number of allowed hosts. Note that if there are less than the allowed hosts, each active host can burst to gain more bandwidth (if you have configured the virtual circuit to allow bursting).
- » For limiting the number of hosts, you can have the system calculate the amount of bandwidth that is then allowed to each host. You can specify an automatic calculation of the per host bandwidth and the number of allowed hosts. The system then divides the virtual circuit bandwidth by the number of active hosts.

#### Adding a dynamic virtual in the Exinda Web UI

On the **Add New Virtual Circuit** form, do the following:

1. Select the **Dynamic Virtual Circuit** checkbox.
2. Set the DVC settings to:
  - **Per Host Bandwidth:** Automatically Share
  - **Per User Max Bandwidth:** 100%
  - **Max Hosts:** Auto

These settings ensure that the all hosts, up to the maximum allowed, automatically receive an equal amount of bandwidth. When the number of hosts is less than the maximum, those hosts are able to share all bandwidth equally up to the maximum available bandwidth.

### EXAMPLE

Consider a 50 Mbps virtual circuit bandwidth.

- » If there is one host, it gets all 50 Mbps.
- » If there are two hosts, they each get 25 Mbps.
- » If there are 500 hosts, they each get 100 kbps.
- » If there are 5000 hosts, they each get 10 kbps (which is the minimum that the DVC will allow)
- » If there are more than 5000 hosts, the additional hosts will not match this virtual circuit.

### Related Topics

- » [Ensuring minimum bandwidth when sharing equally](#)
- » [Limiting the number of hosts sharing bandwidth](#)

### [Adding a dynamic virtual circuit in the EMC](#)

The options available in the Exinda Management Center appear in the following screenshot:

## Dynamic Virtual Circuit:

Allocate minimum and maximum bandwidths to each unique host that matches the virtual circuit.

Enable dynamic virtual circuit

Allocate a bandwidth for each host on the network

Manually define bandwidth usage to be allocated for each host

 % 

Automatically adjust and share bandwidth for each host

Define maximum bandwidth usage for each host

Manually define maximum bandwidth burst for each host

 % 

Do not allow maximum bandwidth burst for each host

Location of hosts to allocate bandwidth

Internal - LAN side

External - WAN side

Define number of maximum hosts to fall under this dynamic virtual circuit

Manually define number of hosts to be affected by this dynamic virtual circuit

Automatically determine number of hosts to be affected by this dynamic virtual circuit

Screenshot 161: Dynamic virtual circuit options

Consider the following:

- » While allocating bandwidth usage to each host on the network.
  - If you manually define the bandwidth usage for each host, you limit the number of hosts that can be accommodated on the dynamic virtual circuit. You can do this by either defining an actual bandwidth or by defining a percentage of the available bandwidth. By default, though, the EMC can set no less than 10 kbps for any one host, so a hard limit to the maximum number of hosts is the total available bandwidth divided by 10 kbps. This setting allows you set higher bandwidth quotas for a limited number of hosts.
  - By allowing the EMC to automatically adjust and share bandwidth, you are letting as many hosts as can be to accommodate within the available bandwidth on the dynamic virtual circuit. The minimum bandwidth that the EMC can provide is 10 kbps, the hard limit to the number of hosts is the total bandwidth divided by 10 kbps.

- » While defining a maximum bandwidth usage for each host, when spare capacity exists on the dynamic virtual circuit because few hosts are active, you can allow the active hosts have greater bandwidth (to burst). You can set the burst rate limit as an actual bandwidth measure (kbps, Mbps, etc.) or specify a percentage of the available bandwidth. You can also disallow bursting.
- » While specifying the location of the hosts as internal or external, remember that this setting allows you specify whether the hosts on the dynamic virtual circuit located within the LAN or outside the LAN.
- » While defining the maximum number of hosts, you can either manually define the maximum number of hosts that can be accommodated on the dynamic virtual circuit, or you can let the EMC control how many hosts are allowed on the circuit.

### Ensuring minimum bandwidth when sharing equally

Use the following instructions to ensure a minimum amount of bandwidth when sharing with as many hosts as possible.

On the **Add New Virtual Circuit** form, do the following:

1. Select the **Dynamic Virtual Circuit** checkbox.
2. If necessary, clear the **Automatically Share** checkbox.
3. Set the DVC settings to:
  - **Per Host Bandwidth:** – Type the minimum amount of bandwidth for the hosts to receive
  - **Per Host Max Bandwidth:** – Set to 100%
  - **Max Hosts:** – Select **Auto**

#### Related Topics

- » [Sharing bandwidth equally](#)
- » [Limiting the number of hosts sharing bandwidth](#)

### Capping bandwidth usage per host

Use the following instructions to cap bandwidth usage on a per host basis.

On the **Add New Virtual Circuit** form, do the following:

1. Select the **Dynamic Virtual Circuit** checkbox.
2. Set the DVC settings to:
  - **Per Host Bandwidth:** Automatically Share
  - **Per User Max Bandwidth:** No burst; set limit to your manual cap value
  - **Max Hosts:** Auto

This will ensure that the hosts will automatically share the bandwidth equally amongst the maximum number of hosts, however no host will use more than the specified cap. If there are many hosts they will share the bandwidth equally, even if it means very little bandwidth per host.

### EXAMPLE

Consider a 50 Mbps virtual circuit bandwidth and a 1 Mbps cap per host.

- » If there is one host, it gets 1 Mbps.
- » If there are 50 hosts, they each get 1 Mbps.
- » If there are 200 hosts, they each get 250 kbps.
- » If there are 5000 hosts, they each get 10 kbps.
- » If there are more than 5000 hosts, the additional hosts will not match this virtual circuit.

### Can I mistakenly prevent fully utilizing the bandwidth?

If you set a bandwidth cap per host by setting the Per User Max Bandwidth and you set the maximum number of hosts, it is possible that you prevent access to excess bandwidth.

$\text{bandwidth-cap-per-host} \times \text{number-of-hosts} = \text{allocated-bandwidth}$

If the allocated bandwidth is less than the virtual circuit bandwidth, then you are making some of the bandwidth inaccessible.

#### Related Topic

[Capping bandwidth usage per host with minimum bandwidth](#)

[Limiting the number of hosts sharing bandwidth](#)

[Sharing bandwidth equally](#)

### Capping bandwidth usage per host with minimum bandwidth

Use the following instructions to cap the bandwidth available to all hosts in a DVC but allow a minimum bandwidth.

#### To cap bandwidth but provide a minimum

On the **Add New Virtual Circuit** form, do the following:

1. Select the **Dynamic Virtual Circuit** checkbox.
2. Set the DVC settings to:
  - **Per Host Bandwidth:** – Type the minimum bandwidth that you want the hosts to have
  - **Per User Max Bandwidth:** – Select **No Bursting Allowed** and set the limit to your manual cap value
  - **Max Hosts:** – Select **Auto**

### Can I mistakenly prevent fully utilizing the bandwidth?

If you set a bandwidth cap per host by setting the Per User Max Bandwidth and you set the maximum number of hosts, it is possible that you prevent access to excess bandwidth.

$\text{bandwidth-cap-per-host} \times \text{number-of-hosts} = \text{allocated-bandwidth}$

If the allocated bandwidth is less than the virtual circuit bandwidth, then you are making some of the bandwidth inaccessible.

## Related Topics

- » [Capping bandwidth usage per host](#)
- » [Limiting the number of hosts sharing bandwidth](#)
- » [Sharing bandwidth equally](#)

## Limiting the number of hosts sharing bandwidth

Use the following instructions to limit the number of hosts sharing the available bandwidth.

On the **Add New Virtual Circuit** form, do the following:

1. Select the **Dynamic Virtual Circuit** checkbox.
2. Set the DVC settings to:
  - **Per Host Bandwidth:** – Select the **Automatically Share** checkbox
  - **Per User Max Bandwidth:** – Set to **100%**
  - **Max Hosts:** – Type a suitable number

This ensures that only the specified number of hosts are managed by this virtual circuit, and that all the hosts in the virtual circuit automatically share the bandwidth equally.

### EXAMPLE

Consider a 50 Mbps virtual circuit bandwidth and 100 hosts as a maximum.

- » If there is one host, it gets all 50 Mbps,
- » If there are ten hosts, they each get 5 Mbps.
- » If there are 100 hosts, they each get 500 kbps.
- » If there are more than 100 hosts, the additional hosts will not match this virtual circuit.

## Related Topics

- » [Sharing bandwidth equally](#)
- » [Ensuring minimum bandwidth when sharing equally](#)

## Troubleshooting virtual circuits

If you are unsure whether traffic is being processed properly by the virtual circuit or the policies within your virtual circuit, it is best to use real time monitoring to determine if traffic is hitting your virtual circuit as you would expect.

1. Go to **Monitor > Real Time > Conversations**.
2. Select the **Show Policies** checkbox at the top of the screen. The screen refreshes to group the conversations by virtual circuit and policy.
3. Look at the traffic falling under your virtual circuit to decide if the correct traffic is in there.
4. If there is traffic that you were expecting that is not there, look in other virtual circuit groups to see where you expected traffic is.
5. If your traffic is falling into the wrong virtual circuit, check the definition of the affected virtual circuits and ensure the most specific virtual circuit is higher in the policy tree.

## 0.13.4 Policies overview

Policies define what actions to perform on specific targeted traffic. The policies can specify whether:

- » to optimize the traffic by bandwidth shaping, acceleration, or marking the packets.
- » to block the traffic by discarding the packets.
- » to monitor the traffic by ignoring the packets.
- » or to redirect the traffic to a specific URL.

The policy managed traffic can then be filtered by:

- » Application or application group
- » Hosts or subnets
- » Hosts or subnets that are communicating with other specific hosts or subnets
- » VLAN
- » ToS/DSCP markings
- » Time of day

You can apply any combination of these filters. For example, the policy could be targeted to traffic between a particular branch and headquarters, which has particular ToS markings on a particular VLAN during work hours.

Furthermore, you can also add more than one filter. That is, the policy could target a particular branch site for Netflix and the same branch site for Silverlight.

When creating policies, they are added to the policy library. If any changes are made to a policy definition they will affect all other virtual circuits that use that policy.

**Add New VC Policy**

Policy Name:

VC Policy Number:

Schedule:

Action:

Policy Enabled:

Guaranteed Bandwidth:  %

Burst (Max) Bandwidth:  %

Burst Priority:

Acceleration:

WM Reduction Type:

ToS/DSCP Mark:

VLAN Rewrite: ID:    Priority:

Filter Rules:

| VLAN                                                              | Host                                                              | Direction                                | Host                                                              | ToS/DSCP                                                          | Application                                                       |
|-------------------------------------------------------------------|-------------------------------------------------------------------|------------------------------------------|-------------------------------------------------------------------|-------------------------------------------------------------------|-------------------------------------------------------------------|
| <input type="button" value="ALL"/>                                | <input type="button" value="ALL"/>                                | <input type="button" value="&lt;-&gt;"/> | <input type="button" value="ALL"/>                                | <input type="button" value="ALL"/>                                | <input type="button" value="File Services"/>                      |
| <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="&lt;-&gt;"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> |
| <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="&lt;-&gt;"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> |
| <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="&lt;-&gt;"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> | <input type="button" value="↑"/> <input type="button" value="↓"/> |

Screenshot 162: Adding a new VC policy.

Policies are part of the policy tree. To learn how circuits, virtual circuits, and policies work together, see [Policy Tree](#).

### Related Topics

- » [Creating a policy for a specific Virtual Circuit](#)
- » [Creating a policy for use with any Virtual Circuit](#)

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to shape bandwidth](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response object webpage](#)
- » [Configuring a policy to redirect HTTP traffic to a URL](#)
- » [Configuring a policy to mark packets](#)

## The policy library

The policy library stores all policy definitions. The policies can be created and edited in the policy tree or in the policy library. A single policy can be used by more than one virtual circuit. Not all policies in the policy library are necessarily used within a virtual circuit.

### NOTE

Any changes made to a policy definition will affect all other virtual circuits that use that policy.

The form at the top of the page allows you to create custom Policies. See [Policies](#) for more information regarding creating and editing policies. In the list of policies, you can edit, clone, or delete policies. You can also use this delete individual filter rules from within policies. Policies are part of the policy tree. To learn how circuits, virtual circuits, and policies work together, see [Policy Tree](#).

To access this functionality, go to **Configuration > Traffic Policies > Optimizer > Policies**.

## Adding a policy to the policy tree

Use the following instructions to add a policy to the policy tree that already exists in the policy library.

1. Go to **Configuration > Traffic Policies > Optimizer > Optimizer > Policies**.
2. Select the desired policy using the drop-down list at the bottom of a virtual circuit's policy list in the Optimizer.
3. Specify the rank order number so that it will be inserted in the desired location in the policy tree.
4. Press the **Add to 'name of virtual circuit'** button.

If you want to create a brand new policy

1. Click on the Create New Policy link at the bottom of the desired virtual circuit's policy list and fill out the policy details.
2. The system will insert the policy to the bottom of the policy list for the given virtual circuit. Reorder the policy if needed. For more information, refer to [Policy ordering matters](#) (page 372).

## Policy ordering matters

Given the top-down evaluation order, if you have multiple virtual circuits that could match the same traffic although one is more specific than the other, the more specific virtual circuit needs to be ordered higher in the circuit's list of virtual circuits as compared to the more general virtual circuit. Otherwise the traffic would be caught by the general virtual circuit and would never reach the more specific virtual circuit. This is true for policies as well. That is, more specific policies need to be ordered higher in the virtual circuit's set of policies as compared to more general virtual circuits.

## EXAMPLE

Consider the scenario where you throttle streaming traffic to 5% of your network and yet John is getting most of that 5%. You could create a policy that throttles John's streaming to 2%. John's streaming policy needs to be ordered higher in the policy tree than the general streaming policy or else his traffic will be handled by the general policy and his traffic will never be evaluated against the policy that was created specifically for his traffic.

### To reorder policies

For each policy that you would like moved, edit the ranking number to the order you would like, then select **Reorder** from the **Actions** menu on the right-hand-side. For example, if there are already policies with ranking order 10, 20, and 30, and you want to swap the order of the policies with ranking order 20 and 30, then either change the ranking order of 30 to be between 10 and 20, say 15, or change the ranking order of 20 to be after 30, say 35. In either case, select **Reorder** from the **Actions** menu on the right-hand-side of the policy that changed its ranking number to submit this change.

### To reorder virtual circuits

For each virtual circuit that you would like moved, edit the virtual circuit by selecting **Edit** from the **Actions** menu on the right-hand-side. Then modify the rank, which is the editable text field as part of the **Virtual Circuit Number**. For example, if there are already virtual circuits with ranking order 10, 20, and 30, and you want to swap the order of the virtual circuits with ranking order 20 and 30, then either change the ranking order of 30 to be between 10 and 20, say 15, or change the ranking order of 20 to be after 30, say 35. In either case, select **Reorder** from the **Actions** menu on the right-hand-side of the virtual circuit that changed its ranking number to submit this change.

### Re-using policies from the library

Policy definitions are created and saved to a policy library, which allows the same policy definition to be used in more than one virtual circuit. This means that if you edit a policy in one virtual circuit, the definition is modified wherever it is used. To modify a policy for a specific virtual circuit, you should clone the policy and then modify the details. Note that if a policy is removed from a virtual circuit, it is not deleted from the policy library.

### Creating a policy for use with any virtual circuit

Use the instructions that follow to create a policy for general use. Once created, it can be applied to any virtual circuit.

#### Before you begin

Some of the following steps require understanding that is not included in the instructions. You should consult the following, where needed:

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response webpage](#)
- » [Configuring a policy to redirect HTTP traffic to a URL](#)

#### To create a policy in the Exinda Web UI

1. Go to **Configuration > Optimizer**. The right side of the screen refreshes.
2. Select the **Policies** tab.
3. In the **Policy Name** field, type a suitable name for the policy.

4. Optionally, specify a **Schedule**.

**NOTE**

The default is "ALWAYS". Schedules that have been defined will appear in the drop-down list. If necessary, you can create schedules by going to **Configuration > Objects > Schedules**.

5. Specify the **Action**.

- **Optimize** – Perform traffic shaping, acceleration, or packet marking on the traffic
- **Discard** – Discard the packets effectively blocking the traffic
- **Ignore** – Perform no action on the packets allowing them to pass through the appliance unaffected. This setting is effective for a monitor only mode. Note that this option should not be used for policies within a dynamic virtual circuit.
- **Return HTML Response** – Present the webpage defined in the HTML response object to the source web browser.
- **HTTP Redirect** – Redirect the packets to a specified URL effectively presenting the URL to the network client.

**NOTE**

Each of the policy actions comes with its own set of options.

6. Complete the additional settings as required by the action.

7. Enter the **Filter Rules**. Any of the following fields may be used to specify how to filter the traffic.

| VLAN | Host | Direction | Host | ToS/DSCP | Application   |
|------|------|-----------|------|----------|---------------|
| ALL  | ALL  | <->       | ALL  | ALL      | File Services |
|      |      | <->       |      |          |               |
|      |      | <->       |      |          |               |
|      |      | <->       |      |          |               |

- **VLAN**– Select traffic based on 802.1Q VLAN ID and/or 802.iP VLAN priority tag using a pre-defined VLAN object.
- **Source/Direction/Destination**– Select traffic based on one end of the conversation belonging to a pre-defined network object (static or dynamic) or select traffic based on one way or two way conversations between two predefined network objects. For the first host, select a network object that filters for the initiation of a conversation. For the second host, select a network object that filters for the destination of the conversation. If hosts are not specified, ALL network objects are assumed. Traffic direction is relative to the Exinda appliance.
- **ToS/DSCP**– Select traffic based on particular ToS/DSCP marks in the IP header.
- **Application**– Select traffic based on a predefined application object or application group.

**NOTE**

- » By default, initially only four filter rules can be created per policy. If more are required, fill out the first four rules, save the Policy, then edit the Policy and four more lines will become available.
- » To delete individual filter rules, set all the fields for that filter rule to blank.

8. Click **Add New Policy** to add this policy to the policy tree. The policy will also be added to the policy library.
9. For each virtual circuit where the policy is required, on the bottom row, open the **Policy** drop-down list, select the new policy.
10. In the adjacent **Order** field, type a number that would place the policy in the correct position in the tree.

#### Related Topics

- » [Creating a policy for a specific Virtual Circuit](#)

#### VERSION INFO

Exinda Management Center 1.5.0 does not support policies for HTTP Redirect or HTTP Response.

When you create policies, they are added to the Policies Library. If you amend a policy definition, any changes made to it affect all Virtual Circuits that use that policy. To learn how circuits, virtual circuits, policy sets, and policy rules work together, see [Policy Tree](#).

#### Where do I find policy rules?

Policy library items can be found in **Library > Policies**.

Also, the policies linked to appliances can also be found in the policy trees for each appliance group. Go to the desired appliance group's **Optimizer Policy Tree**.

#### To create a policy in the EMC configuration Library

1. Got to **Library > Policies**
2. Click **Create new policy in the library**.
3. In the **Name** section, type a name for the policy. The name must be unique within the tenant.
4. In the **Action** section, specify what type of action the rule should take. Select *one* of the following:
  - Optimize**– Selecting optimize causes a new action to appear in the UI where you can specify whether you want to apply bandwidth shaping, prioritization, acceleration, or packet marking.
    - **Discard**– Select discard to specify that you want to block a particular type of application by discarding the packets.
    - **Ignore**– Select ignore to specify that you want to allow packets to pass through without manipulation, that is, traffic monitoring only .
4. In the **Filter** section, specify the type of traffic to which you want to apply the policy. Set any of the following traffic attributes.
  - **Application** – Select traffic based on a predefined application or application group from the list. Custom applications that you have created in the library will appear in this drop-down list.
  - **Source/Direction/Destination** – Select traffic based on one end of a conversation belonging to a pre-defined network object or select traffic based on one way or two way conversations between two predefined network objects. For the source, select a network object that filters for the initiation of a conversation. For the destination, select a network object that filters for the destination of the conversation. If hosts are not specified, ALL network objects are assumed. Traffic direction is relative to the Exinda appliance.
  - **ToS/DSCP** – Select traffic based on particular ToS/DSCP markings in the IP header. When multiple settings are specified, the intersection of those attributes determines whether the traffic is matched.

5. Click the **Create** button. The policy set will be added to the library list.

#### NOTES

- » While creating the policy set, if you create a new policy but cancel creating the policy set, the new policy is retained in the library.
- » Some applications may not be compatible with all appliance versions. When an application is not compatible, it does not appear in the drop-down list.

## Creating a policy for a specific virtual circuit in the Exinda Web UI

Use the following instructions to create a policy that relates to a specific virtual circuit (VC). While these instructions assume that the policy is intended for a specific VC, it can later be used with other VCs.

### Before you begin

Some of the following steps require understanding that is not included in the instructions. You should consult the following, where needed:

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response webpage](#)
- » [Configuring a policy to redirect HTTP traffic to a URL](#)

### To create the policy

1. Go to **Configuration > Optimizer**. The right side of the screen refreshes.

#### NOTE

You can create a policy on either the Optimizer or Policies tab. If you create the policy on the Optimizer tab, the policy must first be associated with a specific virtual circuit. The policy can be associated with other virtual circuits later. If you create the policy on the Policies tab, it is then available for use with any of the virtual circuits.

2. If not already open, select the **Optimizer** tab

3. Click the **Create New Policy** link adjacent to the virtual circuit that requires the new policy. The page refreshes to show the new policy configuration options.

4. In the **Policy Name** field, type a suitable name for the policy.

5. If necessary, edit the pre-populated policy number in the **VC Policy Number** field.

#### NOTE

The policy number determines the order in which the policy is enacted in the policy tree; lower numbers receive higher priority, so highly specific policies should receive lower numbers and more general policies should receive higher numbers.

6. Optionally, specify a **Schedule**.

**NOTE**

The default is "ALWAYS". Schedules that have been defined will appear in the drop-down list. If necessary, you can create schedules by going to **Configuration > Objects > Schedules**.

7. Specify the **Action**.

- **Optimize**– Perform traffic shaping, acceleration, or packet marking on the traffic.
- **Discard**– Discard the packets effectively blocking the traffic
- **Ignore**– Perform no action on the packets allowing them to pass through the appliance unaffected. This setting is effective for a monitor only mode. Note that this option should not be used for policies within a dynamic virtual circuit.
- **Return HTML Response**– Presents the webpage defined in the HTML response object to the source web browser.
- **HTTP Redirect**– Redirect the packets to a specified URL effectively presenting the URL to the network client.

**NOTE**

Each of the policy actions comes with its own set of options.

## 8. Complete the additional settings as required by the action.

9. Enter the **Filter Rules**. Any of the following fields may be used to specify how to filter the traffic.

| VLAN  | Host  | Direction | Host  | ToS/DSCP | Application     |
|-------|-------|-----------|-------|----------|-----------------|
| ALL ▾ | ALL ▾ | <-> ▾     | ALL ▾ | ALL ▾    | File Services ▾ |
| ▾     | ▾     | <-> ▾     | ▾     | ▾        | ▾               |
| ▾     | ▾     | <-> ▾     | ▾     | ▾        | ▾               |
| ▾     | ▾     | <-> ▾     | ▾     | ▾        | ▾               |

- **VLAN**– Select traffic based on 802.1Q VLAN ID and/or 802.iP VLAN priority tag using a pre-defined VLAN object.
- **Source/Direction/Destination**– Select traffic based on one end of the conversation belonging to a pre-defined network object (static or dynamic) or select traffic based on one way or two way conversations between two predefined network objects. For the first host, select a network object that filters for the initiation of a conversation. For the second host, select a network object that filters for the destination of the conversation. If hosts are not specified, ALL network objects are assumed. Traffic direction is relative to the Exinda appliance.
- **ToS/DSCP**– Select traffic based on particular ToS/DSCP marks in the IP header.
- **Application**– Select traffic based on a predefined application object or application group.

**NOTE**

- » By default, initially only four filter rules can be created per policy. If more are required, fill out the first four rules, save the Policy, then edit the Policy and four more lines will become available.
- » To delete individual filter rules, set all the fields for that filter rule to blank.

10. Click **Add New Policy** to add this policy to the policy tree. The policy will also be added to the policy library.

To delete a policy from the system

To delete a policy from the system, it must be deleted from the policy library. Go to **Configuration > Traffic Policies > Optimizer > Policies**.

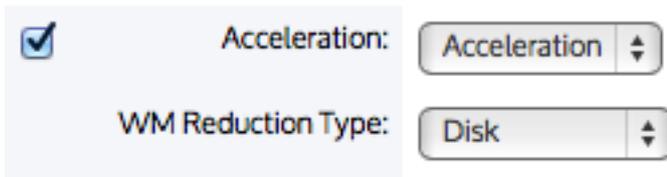
#### Related Topics

» [Creating a policy for use with any Virtual Circuit](#)

### Configuring a policy to accelerate traffic in the Exinda Web UI

Use the following instructions to configure a policy that enables acceleration to the traffic identified by the policy.

1. From the **Action** drop-down list select **Optimize**, and then select the **Acceleration** section checkbox to enable acceleration and reduction techniques.



Screenshot 163: Settings in the Optimizer>Policies tab.

2. Specify the **Acceleration** technique to use to make the traffic faster.

- **Acceleration**- Accelerate using TCP-based acceleration techniques. This is only available on x800 licensed appliances. Only outbound TCP traffic will be accelerated.
- **Edge Cache**- Cache particular types of traffic so that when the same traffic is requested again, it can be served up locally from the cache. Edge Cache can be used with a single appliance.

3. If acceleration was selected, specify the WAN memory (**WM**) **Reduction Type** to reduce the amount of traffic that needs to traverse the network.

- **None**- Do not attempt to reduce the traffic. The traffic will still be accelerated via TCP-based acceleration techniques.
- **Compression**- Compress the traffic using a network optimized LZ compression algorithm. The traffic will also be TCP-accelerated.

#### NOTE

If the compressed output is larger than the original, the appliance will send the original.

- **Disk**- De-duplicate the traffic. The appliance's hard disk drive is used to store the deduplication patterns. The traffic will also be compressed and will be TCP-accelerated.

4. Create the **Filter Rules**.

#### Related Topics

- » [Configuring a policy to shape bandwidth](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response object webpage](#)

- » [Configuring a policy to redirect HTTP traffic to a URL](#)
- » [Configuring a policy to mark packets](#)

## Configuring a policy to shape bandwidth in the Exinda Web UI

Use the following instruction to configure a policy that shapes bandwidth.

To configure a policy

1. From the **Action** drop-down list, select **Optimize** and select the Bandwidth section checkbox for traffic shaping.

The screenshot shows a configuration panel with three rows:

- Guaranteed Bandwidth:** A checked checkbox, an empty text input field, a percentage sign, and a spinner control.
- Burst (Max) Bandwidth:** An unchecked checkbox, an empty text input field, a percentage sign, and a spinner control.
- Burst Priority:** A dropdown menu showing '1 (High)' and a spinner control.

Screenshot 164: Settings in the Optimizer>Policies tab.

2. Specify the **Guaranteed Bandwidth** for the policy specified as kbps or percentage of the parent virtual circuit's bandwidth. This is the amount of bandwidth that the policy will have available if needed. This is not reserved bandwidth. If the policy does not use the entire guaranteed bandwidth, the excess bandwidth will be made available for use by other policies in the virtual circuit. The guaranteed bandwidth of a single policy must not exceed the parent virtual circuit bandwidth. The sum of all guaranteed bandwidths for each policy within a virtual circuit must not exceed the virtual circuit bandwidth.
3. Specify the **Burst (Maximum) Bandwidth** for the policy specified as kbps or percentage of the parent virtual circuit's bandwidth. This is the maximum amount of bandwidth that the policy can have access to if there is excess bandwidth available. The burst bandwidth must be greater than the guaranteed bandwidth, and less than or equal to the parent virtual circuit bandwidth.
4. Specify the **Burst Priority** ranging from **1 (High)** to **5 (Normal)** to **10 (Low)**. If excess bandwidth is available, the burst priority is used to decide how excess bandwidth is distributed. Policies with a higher burst priority will be preferred when allocating excess bandwidth.
5. Enter the **Filter Rules**.

### Related Topics

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response object webpage](#)
- » [Configuring a policy to redirect HTTP traffic to a URL](#)
- » [Configuring a policy to mark packets](#)

### How traffic-shaping queue modes work

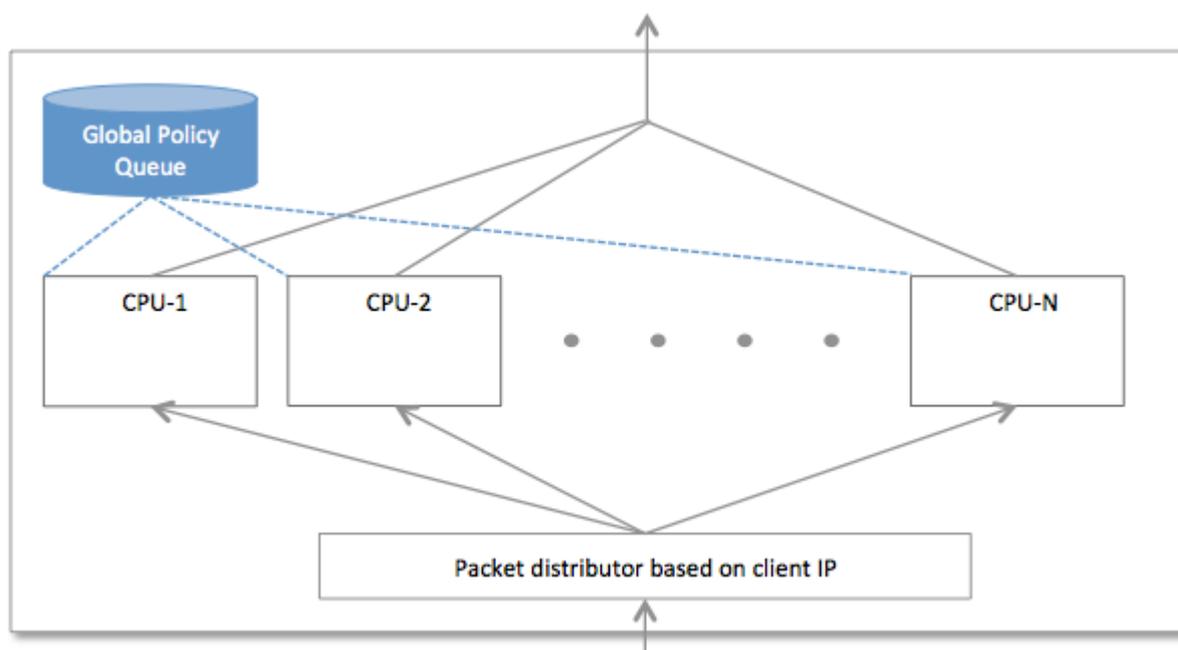
When shaping traffic as specified by the policies, the Exinda Appliance needs to maintain a queue of packets for each policy within their respective virtual circuits. There are three different queuing methods to address the needs of different use scenarios when using a multi-processor appliance.

- » Single queue mode is the default. It is good for environments that require < 1.5 Gbps of traffic shaping.
- » Multi queue mode is good for environments that require 10 Gbps of traffic shaping and have many flows per virtual circuit. Education institutions are ideal for this queuing mode.
- » Multi per VC queue mode is good for environments that require < 5 Gbps of traffic shaping and there are potentially fewer flows per virtual circuit, for example, when the bandwidth of a particular virtual circuit is tested using a single-flow generator, such as speedtest.net.

The mode can be modified via [CLI: Optimizer](#).

### Single Queue Mode

The single queue mode uses a single global policy queue in memory to handle the traffic shaping of all virtual circuits. The bandwidth limit is due to limitations of accessing the global memory with appropriate memory locks.



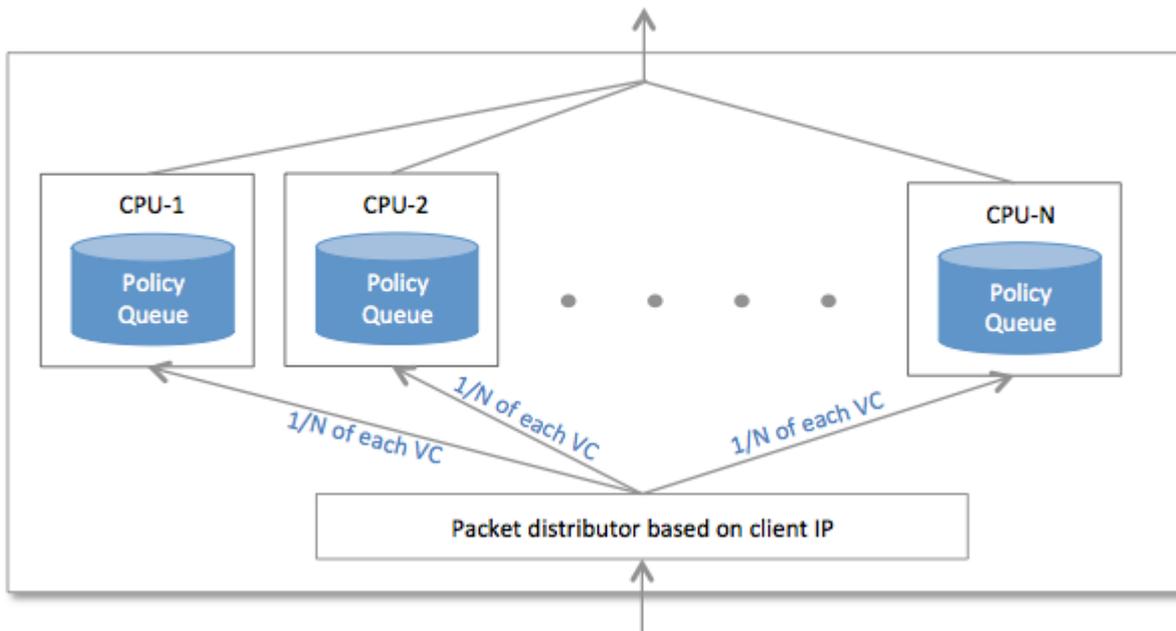
Screenshot 165: Single Queue Mode

### Multi-Queue Mode

The multi-queue mode uses one policy queue per CPU where the flows of a given virtual circuit are divided evenly among the policy queues. This eliminates the cross processor locking of global memory by distributing the policy queues among the processors. Each flow within each policy is handled by a single processor.

Each CPU processes  $1/N$  of each virtual circuit's traffic, where  $N$  is the number of CPUs. That is, if the virtual circuit is specified in the Optimizer as having 10 Mbps bandwidth, in an appliance with four CPUs, each CPU policy queue will be allowed  $10 \text{ Mbps} / 4 = 2.5 \text{ Mbps}$  per CPU. In order to have even distribution, it assumes multiple flows that can be distributed among the  $N$  CPUs. This queuing method is not good for environments where customers validate the amount of bandwidth they receive by sending a single long flow through their virtual circuit.

In this case, the flow is handled by a single CPU and the other CPUs are idle. It then appears that they are getting  $1/N$  of the amount of traffic that they are expecting even though in more realistic use of the network, where the flows can be distributed more evenly, they will get the appropriate amount of bandwidth.



Screenshot 166: Multi Queue Mode

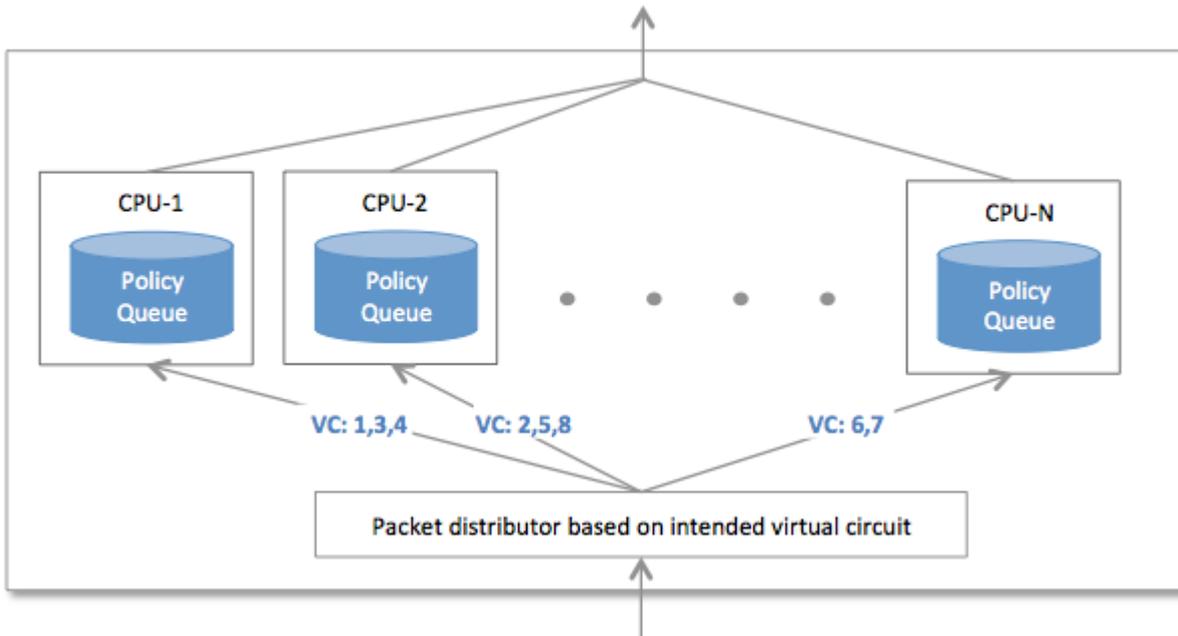
### Multi Per VC Queue Mode

The multi per VC queue mode uses one policy queue per CPU where each virtual circuit is assigned to a single policy queue. That is, the virtual circuits are distributed amongst the policy queues, but any given virtual circuit exists on a single policy queue. The flows for any given virtual circuit will be handled by a single CPU policy queue and therefore will not be limited to 1/N of the traffic when a single flow is tested.

#### CAUTION

Each virtual circuit is assigned to individual policy queues and any given virtual circuit cannot use a policy queue that it has not been assigned to. Therefore, the virtual circuits cannot be oversubscribed, that is the sum of the desired bandwidths for the virtual circuits cannot be higher than the specified bandwidth of the circuit. This is because there can be no automatic redistribution of minimum bandwidth among the virtual circuits when the virtual circuits are oversubscribed.

If the circuits are oversubscribed, then the shaping queuing mode will revert to the **multi-queue mode**.



Screenshot 167: Multi per VC Queue Mode

## Configuring a policy to discard (block) traffic in the Exinda Web UI

Use the following instructions to configure a policy that discards (blocks) identified traffic.

1. From the Action drop-down list, select **Discard** and optionally check the **Discard only the first packet of a connection** option. This option can be used in conjunction with a uni-directional virtual circuit to discard connections originating from a specific side (WAN or LAN) of the appliance. For example, when used with an inbound Virtual Circuit, the first (SYN) packet will be discarded - effectively blocking connection establishment from the WAN but allowing traffic from established connections.

Block Options:  Discard only the first packet of a connection

2. Enter the **Filter Rules**.

## Configuring a policy to mark packets in the Exinda Web UI

Use the following instructions to configure a policy such it marks packets

### To configure the policy

1. From the Action drop-down list, select **Optimize** and select the **Packet Marking** section checkbox to mark individual packets matching this policy.

**ToS/DSCP Mark:**

**VLAN Rewrite:** ID:  Priority:

2. Specify which **ToS/DSCP Mark** to put in the IP header of each packet. For more information, refer to [How Exinda Uses the ToS/DiffServ field](#) (page 386).

## NOTE

If you select the **Optimize** Action and then select the **Packet Marking** checkbox, the interface provides an opportunity to use a **VLAN Rewrite** feature by selecting an **ID** and a **Priority** from drop down lists. While the option is there, the feature has been disabled and will be removed in a future version of the software.

3. Enter the **Filter Rules**.

### Related Topics

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to shape bandwidth](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response object webpage](#)
- » [Configuring a policy to redirect HTTP traffic to a URL](#)

## Configuring a policy to redirect HTTP traffic to an HTTP response webpage

Use the following instructions to configure a policy to redirect HTTP traffic to an HTTP response page.

### To configure the policy

1. From the **Action** drop-down list, select the **Return HTTP Response** and select the **HTML Response** object checkbox. The HTML Response object allows you to "host" a webpage on the Exinda appliance. You create the HTML Response webpage at **Configuration > Objects > HTML Response**. For more information, refer to [Creating an HTML response object](#) (page 213). HTTP traffic matching this policy is presented with HTTP webpager. This option is useful to notify users when they are no longer allowed to use the network for HTTP, HTTP-ALT, or HTTPS traffic.

2. Enter the **Filter Rules**. Similar to other policy configurations, you can specify VLAN, Source, Direction, Destination, ToS/DSCP, or Application. However, the only allowable applications are HTTP, HTTP-ALT, and HTTPS.

### Related Topics

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to shape bandwidth](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to a URL](#)
- » [Configuring a policy to mark packets](#)

## Configuring a policy to redirect HTTP traffic to a URL

Use the following instructions to configure a policy such that identified HTTP traffic is redirected to a specific URL.

### To configure the policy

1. From the **Action** drop-down, select the **HTTP Redirect** action and specify the **Redirect URL**. HTTP traffic matching this policy is forwarded to the specified URL. This option is useful to redirect a particular set of traffic to another URL, such as redirecting unauthenticated traffic to a login page, or redirecting traffic from a host that has exceeded the bandwidth allotment to a URL that presents a message.

2. Enter the **Filter Rules**. Similar to other policy configurations, you can specify VLAN, Host/Direction/Host, ToS/DSCP, or Application. However, the only allowable applications are HTTP, HTTP-ALT, and HTTPS.

To learn more about common use cases for redirecting HTTP traffic, such as, redirecting unauthenticated users to a login page of your captive portal, see [Integrating Exinda with your Captive Portal](#), or such as, redirecting users who have exceeded their data usage quota, see [Quota Enforcement Scenarios](#).

#### Related Topics

- » [Configuring a policy to accelerate traffic](#)
- » [Configuring a policy to shape bandwidth](#)
- » [Configuring a policy to discard \(block\) traffic](#)
- » [Configuring a policy to redirect HTTP traffic to an HTTP Response object webpage](#)
- » [Configuring a policy to mark packets](#)

#### Reading and writing ToS and DiffServ packet marks

Exinda Appliances can read and write ToS/DSCP marks in packets, allowing fine-grained control and classification of applications that are marked with Tos/DSCP values as well as applying marking policies to ensure traffic is treated appropriately by onward network equipment.

Used in conjunction with Exinda classification techniques, including advanced layer 7 detection, users have complete control over how traffic is marked and subsequently treated in the WAN cloud.

#### Related topics

- » [How ToS and DiffServ Work](#)
- » [The ToS/DiffServ Field](#)

#### How ToS and DiffServ work

The ToS (type of service) or DiffServ (differentiated services) field in the IPv4 header, and the Traffic Class field in the IPv6 header are industry standard fields used to classify IP packets so that routers can make QoS (quality of service) decisions about what path packets should traverse across the network.

Exinda Appliances can read and write ToS/DSCP marks in packets, allowing users fine-grained control and classification of applications that are marked with ToS/DSCP values as well as applying marking policies to ensure traffic is treated appropriately by onward network equipment. For example, users may want to ensure that VoIP uses high quality, low latency (and expensive) links, or, they might want to ensure email or recreational traffic uses cheaper (but less reliable) links.

Used in conjunction with the Exinda classification techniques, including advanced layer 7 detection, users have complete control over how traffic is marked, and subsequently treated in the WAN cloud.

Previously, there were 5 different categories that users could use to classify their traffic with the IP ToS field (see [RFC 791](#)): Normal Service, Minimize Cost, Maximize Reliability, Maximize Throughput, and Minimize Delay.

These have since been replaced by a new set of values called DSCP (DiffServ Code Points, see [RFC 2474](#)). A DSCP is a 6-bit number; therefore, there are 64 possible DSCP combinations, of which, only a portion have been standardized. IPv6 contains an 8 bit Traffic Class field. The 6 most significant bits are treated the same as IPv4 DSCP. The least 2 significant bits are not modified by the appliance.

All Exinda products can read and write the ToS/DiffServ field, allowing users to:

- » Match packets with a ToS/DSCP value and apply optimizer policies to this traffic.
- » Mark the packets with a ToS/DSCP value based on source/destination host/subnet, source/destination port, layer 7 application, time of day, VLAN id, etc.

**How to create optimizer policies based on ToS/DSCP value?**

For more information, refer to [Policies overview](#) (page 371).

**How to mark packets with ToS/DSCP value?**

For more information, refer to [Policies overview](#) (page 371).

**The ToS/DiffServ field**

The ToS (type of service) or DiffServ (differentiated services) field in the IPv4 header, and the Traffic Class field in the IPv6 header are used to classify IP packets so that routers can make QoS (quality of service) decisions about what path packets should traverse across the network.

For example, users may want to ensure that VoIP utilizes high quality, low latency (and expensive) links, or, they might want to ensure email or recreational traffic uses cheaper (but less reliable) links.

Previously, there were 5 different categories that users could classify their traffic with using the IP ToS field (see RFC 791).

- » Normal Service
- » Minimize Cost
- » Maximize Reliability
- » Maximize Throughput
- » Minimize Delay

These have since been replaced by a new set of values called DSCP (DiffServ Code Points, see RFC 2474). A DSCP is a 6-bit number. This provides 64 possible DSCP combinations, of which, only a portion have been standardized and are listed below.

IPv6 contains an 8 bit Traffic Class field. The 6 most significant bits are treated the same as IPv4 DSCP. The least 2 significant bits are not modified by the appliance.

| DSCP Class Name                         | Binary Value | Decimal Value |
|-----------------------------------------|--------------|---------------|
| BE (best effort, default)               | 000000       | 0             |
| AF11 (assured forwarding, see RFC 2597) | 001010       | 10            |
| AF12                                    | 001100       | 12            |
| AF13                                    | 001110       | 14            |
| AF21                                    | 010010       | 18            |
| AF22                                    | 010100       | 20            |
| AF23                                    | 010110       | 22            |
| AF31                                    | 011010       | 26            |
| AF32                                    | 011100       | 28            |
| AF33                                    | 011110       | 30            |
| AF41                                    | 100010       | 34            |

| DSCP Class Name                         | Binary Value | Decimal Value |
|-----------------------------------------|--------------|---------------|
| AF42                                    | 100100       | 36            |
| AF43                                    | 100110       | 38            |
| CS1 (class selector)                    | 001000       | 8             |
| CS2                                     | 010000       | 16            |
| CS3                                     | 011000       | 24            |
| CS4                                     | 100000       | 32            |
| CS5                                     | 101000       | 40            |
| CS6                                     | 110000       | 48            |
| CS7                                     | 111000       | 56            |
| EF (expedited forwarding, see RFC 2598) | 101110       | 46            |

#### Related topic

[How Exinda Uses the ToS/DiffServ Field](#)

[How Exinda Uses the ToS/DiffServ field](#)

All Exinda products can read and write the ToS/DiffServ field, allowing users to:

- » Match packets with a ToS/DSCP value and apply optimizer policies to this traffic.
- » Mark the packets with a ToS/DSCP value based on source/destination host/subnet, source/destination port, layer 7 application, time of day, vlan id, etc.

#### Related topics

- » [Match Packets to ToS/DSCP Values](#)
- » [Mark Packets with ToS/DSCP Values](#)

#### Matching packets to ToS/DSCP values

When defining Optimizer Policies on the Exinda appliance, there is a ToS/DSCP drop down that allows users to match only those packets with the specified ToS/DSCP value.

Screenshot 168: Optimizer Policy configuration page.

Users can select the appropriate DSCP/ToS value from this drop down field and any packets that match this ToS/DSCP value will be applied to this policy.

### EXAMPLE

VoIP equipment in a user’s network may be configured to mark all outgoing packets as DSCP EF (decimal 46). VoIP is a real-time application and the user wishes to prioritize this with a high priority policy that guarantees VoIP a certain amount of WAN bandwidth. To achieve this, the user selects ‘DSCP 46’ from the ToS/DSCP drop down and configures the appropriate bandwidth allocation in this policy.

### Marking packets with ToS/DSCP values

Users may want to mark certain packets with a ToS/DSCP value so that external routers can treat the traffic appropriately. The same policy configuration screen above (see Figure 1) allows users to configure such an action.

When the policy action is set to ‘Optimize’, several options are available on the right-hand side, one of which is the ‘ToS/DSCP Mark’ checkbox. Users will need to enable this feature by checking the box and selecting the appropriate ToS/DSCP mark from the drop down.

Any traffic that matches the corresponding filter rules is then marked with the specified value and should be treated appropriately by routing equipment down the line.

### EXAMPLE

Service Providers may provide users with a table similar to the one below (example only). Each class has different guaranteed service and pricing levels. This information should be used in conjunction with optimizer policies to implement and ensure quality of service. See Table 1 to convert the DSCP Settings to a decimal value that can be used in the Optimizer Policies.

| Traffic Priority Class         | IETF DiffServ Traffic Priority Class | DSCP Setting |
|--------------------------------|--------------------------------------|--------------|
| Real Time (Gold)               | Expedited Forwarding                 | EF           |
| Mission Critical (Silver High) | Assured Forwarding                   | AF31         |
| Business Critical (Silver Low) | Assured Forwarding                   | AF32/33      |
| General Business (Bronze)      | Best Effort                          | BE           |

## Creating policy sets in the EMC

Policy sets are an ordered list of policies that can be applied to one or more virtual circuits in one or more appliance groups.

Policy sets can be created in the Policy Set Library directly and then later assigned to an Optimizer Policy Tree, or they can be created in an Optimizer Policy Tree. When saved in an Optimizer Policy Tree, they are also saved to the Configuration Library for use elsewhere. The policy sets within a tenant must have unique names. When a policy set is modified, all uses of it are modified.

There are a few default policy sets that you can use or modify. They correspond to the defined configuration when the wizard is run on the appliance. The different policy sets are due to varying answers to the wizard questions.

To learn how Circuits, Virtual Circuits, Policy Sets, and Policy Rules work together, see [Policy Tree](#).

### Where do I find policy sets?

Policy set library items can be found in **Library > Policy Sets**. Also, the policy sets linked to appliances can also be found in the policy trees for each appliance group. Go to the desired appliance group in the **Optimizer Policy Tree**.

### To create a policy set in the Library

1. Go to **Library > Policy Sets**.
2. Click **Create new policy set**
3. In the **Name** section, key-in a name for the policy. The name must be unique within the tenant.
4. In the **Policies** section, add policies to the list.
  - You can select a policy from the library by clicking **Add policy from library**.
  - You can create a new policy by clicking **Create new policy**.

Policies created in the policy set are automatically added to the policy library. For more information, refer to [Policies](#) (page 430).

4. To reorder the policies, drag and drop the policy rule to the desired location.
5. Click **Create**. The policy set is added to the library list.

#### NOTE

While creating the policy set, if you create a new policy but cancel creating the policy set, the new policy is retained in the library.

### To create a new policy set directly in the Optimizer Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Create new policy set...**
3. Click **Create new policy set**
4. In the **Name** section, key-in a name for the policy. The name must be unique within the tenant.
5. In the **Policies** section, add policies to the list.

- You can select a policy from the library by clicking **Add policy from library**.
- You can create a new policy by clicking **Create new policy**.

Policies created in the policy set are automatically added to the policy library. For more information, refer to [Policies](#) (page 430).

4. To reorder the policies, drag and drop the policy rule to the desired location.

5. Click **Save**. The policy set is added to the Optimizer Policy tree for the current appliance group and is also saved to the Configuration Library.

### To add a policy set from the library to a Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add policy set from library**.
3. Select the desired policy set from the drop-down list.

## 0.13.5 The optimizer wizard

The Optimizer Wizard is a convenient way to populate the Optimizer with default policies.

### CAUTION

Running the Optimizer Wizard will delete any existing Optimizer Policies and Optimizer Configuration.

To start the optimizer wizard:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Configuration > Optimizer > Wizard**.
5. Select one of the options for each step.
6. Click **Apply**.

### NOTE

When asked if you have asymmetric traffic, answer "YES" if the inbound and outbound direction of traffic are flowing through different links. Otherwise, select "NO".

The first four questions are always the same:

- » **Step 1: Do you want to start Optimization when this wizard is completed?** - Selecting YES will start the Optimizer service automatically when you complete all the steps in the wizard.
- » **Step 2: Do you want to configure optimization policies?** - Selecting YES will cause Steps 3 and 4 to appear and will enable tcp-dual-bridge-bypass and monitoring-dual-bridge-bypass. Selecting NO will disable those settings. Selecting UNKNOWN will leave the settings as they are.
- » **Step 3: Do you want to accelerate?** - Selecting YES will create policies that accelerate WAN applications. You must have another Exinda appliance on the WAN for this to work.
- » **Step 4: Do you want to apply QoS?** - Selecting YES will apply traffic shaping. You must have another Exinda appliance on the WAN for this to work.

Depending on your answers to Steps 3 and 4, the following scenarios are possible:

- » Scenario 1: Acceleration & QoS (Quality of Service, i.e. traffic shaping)
- » Scenario 2: QoS Only
- » Scenario 3: Acceleration Only

### Scenario 1: Acceleration & QoS (Quality of Service, i.e. traffic shaping)

- » **Step 3: Do you have asymmetric traffic?** - YES
- » **Step 4: Do you want to enable Optimization?** - YES

| Optimizer Wizard                                                                                                                                                             |                                                                                                                                                                                                                                |                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Step 1:                                                                                                                                                                      | Do you want to start Optimization when this wizard is completed?                                                                                                                                                               | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 2:                                                                                                                                                                      | Do you want to configure new Optimization Policies?<br><i>Selecting YES will overwrite any existing policies you have configured.</i>                                                                                          | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 3:                                                                                                                                                                      | Do you have asymmetric traffic?<br><i>Selecting YES will enable tcp-dual-bridge-bypass and monitoring-dual-bridge-bypass. Selecting NO will disable those settings. Selecting UNKNOWN will leave the settings as they are.</i> | <input checked="" type="radio"/> Yes <input type="radio"/> No <input type="radio"/> Unknown |
| Step 4:                                                                                                                                                                      | Do you want to enable Optimization?<br><i>Selecting YES will create policies that optimize and accelerate WAN applications. Note: You must have another Exinda appliance on the WAN for this to work.</i>                      | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 5:                                                                                                                                                                      | Do you want to enable Control?<br><i>Selecting YES will apply traffic shaping.</i>                                                                                                                                             | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 6:                                                                                                                                                                      | Select the topology type WAN or WAN + Internet?                                                                                                                                                                                | <input checked="" type="radio"/> WAN <input type="radio"/> WAN + Internet                   |
|  <p>Internet traffic for this site is routed over the WAN, usually via another site.</p> |                                                                                                                                                                                                                                |                                                                                             |
| Step 7:                                                                                                                                                                      | Enter inbound bandwidth (kbps)?<br><i>(MAX = 1024000)</i>                                                                                                                                                                      | <input type="text" value="1024000"/> kbps                                                   |
| Step 8:                                                                                                                                                                      | Enter outbound bandwidth (kbps)?<br><i>(MAX = 1024000)</i>                                                                                                                                                                     | <input type="text" value="1024000"/> kbps                                                   |

Screenshot 169: The Optimizer Wizard configuration.

This will enable both QoS (traffic shaping) and Application Acceleration. You will need to select the WAN topology that best represents your deployment and also type the inbound and outbound bandwidths for this Exinda appliance.

### Scenario 2: QoS Only

- » **Step 3: Do you have asymmetric traffic?** - NO
- » **Step 4: Do you want to enable Optimization?** - YES

| Optimizer Wizard                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                        |                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Step 1:                                                                                                                                                                            | Do you want to start Optimization when this wizard is completed?                                                                                                                                                                                                                                                       | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 2:                                                                                                                                                                            | Do you want to configure new Optimization Policies?<br><i>Selecting YES will overwrite any existing policies you have configured.</i>                                                                                                                                                                                  | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 3:                                                                                                                                                                            | Do you have asymmetric traffic?<br><i>Selecting YES will enable tcp-dual-bridge-bypass and monitoring-dual-bridge-bypass. Selecting NO will disable those settings. Selecting UNKNOWN will leave the settings as they are.</i>                                                                                         | <input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Unknown |
| Step 4:                                                                                                                                                                            | Do you want to enable Optimization?<br><i>Selecting YES will create policies that optimize and accelerate WAN applications. Note: You must have another Exinda appliance on the WAN for this to work.</i>                                                                                                              | <input type="radio"/> Yes <input checked="" type="radio"/> No                               |
| Step 5:                                                                                                                                                                            | Do you want to enable Control?<br><i>Selecting YES will apply traffic shaping.</i>                                                                                                                                                                                                                                     | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 6:                                                                                                                                                                            | Do you want to apply an Enterprise or Service Provider policy template?<br><i>Enterprise policies are strict in capping usage for P2P &amp; recreational applications. Service Provider policies are more generous with P2P &amp; recreational traffic usage but these traffic groups are still bandwidth limited.</i> | <input checked="" type="radio"/> Enterprise <input type="radio"/> Service Provider          |
| Step 7:                                                                                                                                                                            | Select the topology type WAN or WAN + Internet?                                                                                                                                                                                                                                                                        | <input checked="" type="radio"/> WAN <input type="radio"/> WAN + Internet                   |
|  <p><i>Internet traffic for this site is routed over the WAN, usually via another site.</i></p> |                                                                                                                                                                                                                                                                                                                        |                                                                                             |
| Step 8:                                                                                                                                                                            | Enter inbound bandwidth (kbps)<br><i>(MAX = 10240000)</i>                                                                                                                                                                                                                                                              | <input type="text" value="10240000"/> kbps                                                  |
| Step 9:                                                                                                                                                                            | Enter outbound bandwidth (kbps)<br><i>(MAX = 10240000)</i>                                                                                                                                                                                                                                                             | <input type="text" value="10240000"/> kbps                                                  |

Screenshot 170: The Optimizer Wizard configuration.

This will enable QoS (traffic shaping) only. You have the choice of the type of default policy template to apply- one is better suited for Enterprise, the other is better suited for Service Providers. You will need to select the WAN topology that best represents your deployment. And you will also need to type the inbound and outbound bandwidths for this Exinda appliance.

### Scenario 3: Acceleration Only

- » **Step 3: Do you asymmetric traffic?** - YES
- » **Step 4: Do you want to enable Optimization?** - NO

| Optimizer Wizard |                                                                                                                                                                                                                                |                                                                                             |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Step 1:          | Do you want to start Optimization when this wizard is completed?                                                                                                                                                               | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 2:          | Do you want to configure new Optimization Policies?<br><i>Selecting YES will overwrite any existing policies you have configured.</i>                                                                                          | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 3:          | Do you have asymmetric traffic?<br><i>Selecting YES will enable tcp-dual-bridge-bypass and monitoring-dual-bridge-bypass. Selecting NO will disable those settings. Selecting UNKNOWN will leave the settings as they are.</i> | <input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Unknown |
| Step 4:          | Do you want to enable Optimization?<br><i>Selecting YES will create policies that optimize and accelerate WAN applications. Note: You must have another Exinda appliance on the WAN for this to work.</i>                      | <input checked="" type="radio"/> Yes <input type="radio"/> No                               |
| Step 5:          | Do you want to enable Control?<br><i>Selecting YES will apply traffic shaping.</i>                                                                                                                                             | <input type="radio"/> Yes <input checked="" type="radio"/> No                               |

Screenshot 171: The Optimizer Wizard configuration.

This will enable Application Acceleration only.

The optimizer will populate the policy tree with circuits, virtual circuits, and policies. To learn how circuits, virtual circuits, and policies work together, see [Policy Tree](#).

To access this functionality, go to **Configuration > Traffic Policies > Optimizer > Wizard**.

The following tables shows the policies that will be available to some of the default application groups after running the Optimizer wizard.

| No. | Name                                 | Min BW% | Max BW% | Priority | Accelerate |
|-----|--------------------------------------|---------|---------|----------|------------|
| 1   | Ignore                               | -       | -       | -        | -          |
| 2   | Accelerate                           | -       | -       | -        | X          |
| 3   | Choke 1%-3%                          | 1       | 3       | 10       | -          |
| 4   | Limit Low 2%-10%                     | 2       | 10      | 10       | -          |
| 5   | Limit Med 3%-50%                     | 3       | 50      | 9        | -          |
| 6   | Limit High 4%-70%                    | 4       | 70      | 8        | -          |
| 7   | Guarantee Low 5%-100%                | 5       | 100     | 7        | -          |
| 8   | Guarantee Med 8%-100%                | 8       | 100     | 5        | -          |
| 9   | Guarantee High 10%-100%              | 10      | 100     | 3        | -          |
| 10  | Guarantee Critical 15%-100%          | 15      | 100     | 1        | -          |
| 11  | Guarantee Low 5%-100% - Accelerate   | 5       | 100     | 6        | X          |
| 12  | Guarantee Med 8%-100% - Accelerate   | 8       | 100     | 4        | X          |
| 13  | Guarantee High 10%-100% - Accelerate | 10      | 100     | 2        | X          |

### 0.13.6 Calculating network performance metrics

A transaction is defined as a client request followed by a server reply, including both TCP and UDP flows. With each read and write transaction between a client and a server, the following values are measured and used to calculate how long the transaction takes to complete:

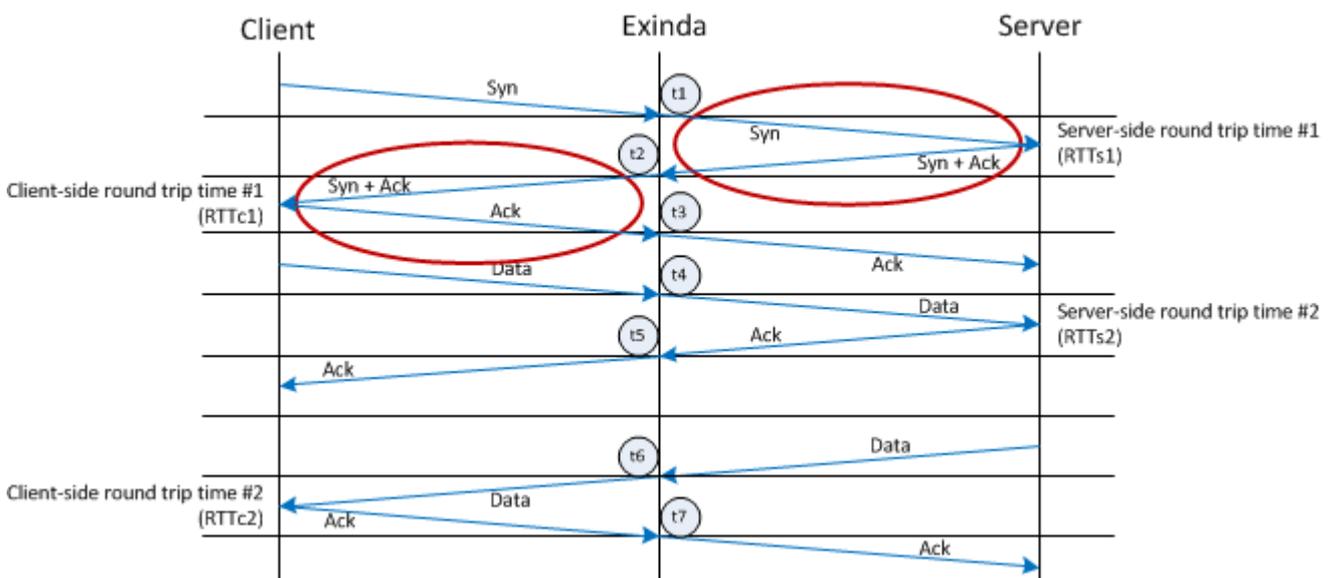
- » **Round Trip Time** — the time taken for a very small packet to travel across the network and return
- » **Network Delay**— the overall time taken for data to cross from a client to a server, or from the server to a client.
- » **Server Delay**— the time taken for a server to respond to a request.
- » **Total Transaction Delay**— The time taken for data to cross the network from a client to the server and back. Unlike round trip time, this could include large packets and could result in multiple packets being sent to the server, or received from the server.
- » **Network Jitter**— measures the variability of the network delay time. This is expressed as a multiple of one standard deviation.
- » **Packet Loss**— measures when one or more packets within a transmission are successfully sent, but fail to arrive at the destination.

#### Calculating round trip time

Round Trip Time (RTT) is the measure of how long it takes for a very small packet to travel across the network and for an acknowledgment of that packet to be returned. Consider the typical topology where an Exinda appliance is positioned between the client and the server.

As each packet is intercepted by the Exinda appliance, it is time stamped with a highly accurate nanosecond resolution clock source. Since the Exinda appliance intercepts the packet after the client sends the packet, the start time is not known and so the RTT is determined by summing the round trip time from the appliance to the server and back (Server RTT), and the round trip time from the appliance to the client and back (Client RTT). As more packets are sent from the client through the Exinda appliance make the round trip, the RTT estimate is updated by averaging new information.

The following diagram illustrates how the round trip time is calculated:



Screenshot 172: This diagram shows how RTT is calculated.

Server RTT:

- »  $RTTs1 = t2 - t1$
- »  $RTTs2 = t5 - t4$

Client RTT:

- »  $RTTc1 = t3 - t2$
- »  $RTTc2 = t7 - t6$

Average RTT:

Average Server RTT =  $(RTTs1 + RTTs2) / 2$

Average Client RTT =  $(RTTc1 + RTTc2) / 2$

Average Total RTT = avRTTs + avRTTc

#### Related topics

- » [Network and Server Delay in a Read Transaction](#)
- » [Network and Server Delay in a Write Transaction](#)
- » [Delay Normalization](#)
- » [Packet loss](#)

### Normalizing transaction delay metrics

To create accurate comparisons of the network delay experienced by a transaction, the appliance must analyze packets of the same size (normalized). All other factors being equal, the transaction delays should increase with the amount of data transferred or the transaction size.

To make the APS score independent of transaction size, the transaction delay metrics are normalized using a constant of 1024 bytes. The normalized network delay is calculated as follows:

Normalized Network Delay = Total Network delay \* 1024 / transaction bytes

#### NOTE

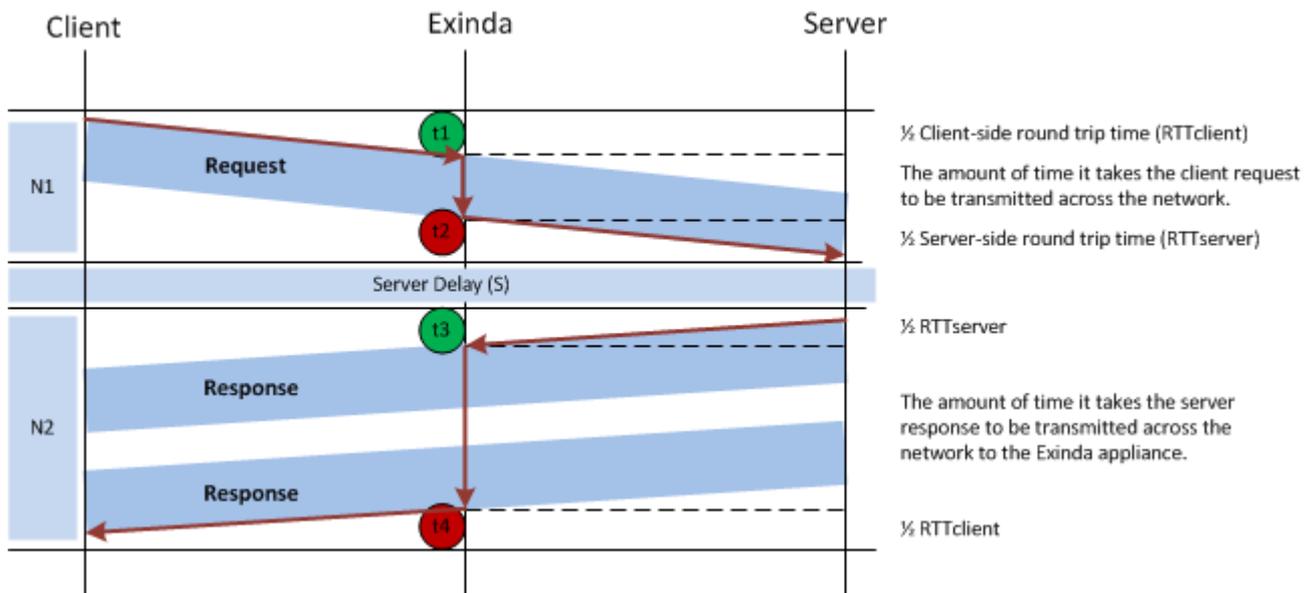
Due to the nature of normalization, protocols that naturally use very small sized packets (like RDP and Citrix) will produce APM values higher than the real ones given such that the results are being multiplied by 1024. The number of bytes used to normalize the calculation of the network delay during a transaction can be configured through the CLI, as well as disabling normalization altogether. For more information refer to [CLI: APM](#).

Related Topics:

- » [Round Trip Time](#)
- » [Network and Server Delay in a Read Transaction](#)
- » [Network and Server Delay in a Write Transaction](#)
- » [Packet loss](#)

### Calculating network and server delay in a read transaction

When a client computer requests information from the server, the request and response are tracked to determine how long it takes for the client to send the request, and the server to send the requested data back to the client. The diagram below shows the flow of information between the client, the Exinda appliance, and the server, and identifies the points in the transaction where time stamps are acquired.



Screenshot 173: A visual presentation of a network read transaction.

### Network Delay for Read Request (N1)

- » The client sends a request to the server.
- » When the request passes through the Exinda, the time stamp is noted as the beginning of the request ( $t_1$ ).
- » When the end of the request passes through the Exinda, the time stamp is noted ( $t_2$ ).  $t_2 - t_1$  = The amount of time it takes the client request to pass through the Exinda appliance.
- » The server receives the complete client request.

### Server Delay for Read Request (S)

- » After the server receives a request from the client, the server takes some time to process the request. This is the Server delay (S).

### Network Delay for Read Response (N2)

- » The server's response to the client request is sent, and may be sent in a number of packets.
- » When the first response passes through the Exinda, the time stamp is noted ( $t_3$ ).
- » When the end of the last response passes through the Exinda, the time stamp is noted ( $t_4$ ).  $t_4 - t_3$  = The amount of time it takes the data requested by the client to pass through the Exinda appliance.
- » The client receives the data requested from the server.

### Total Time for Read Transaction

The total transaction time for a Read transaction is calculated as  $\text{Transaction time} = N1 + S + N2$  where:

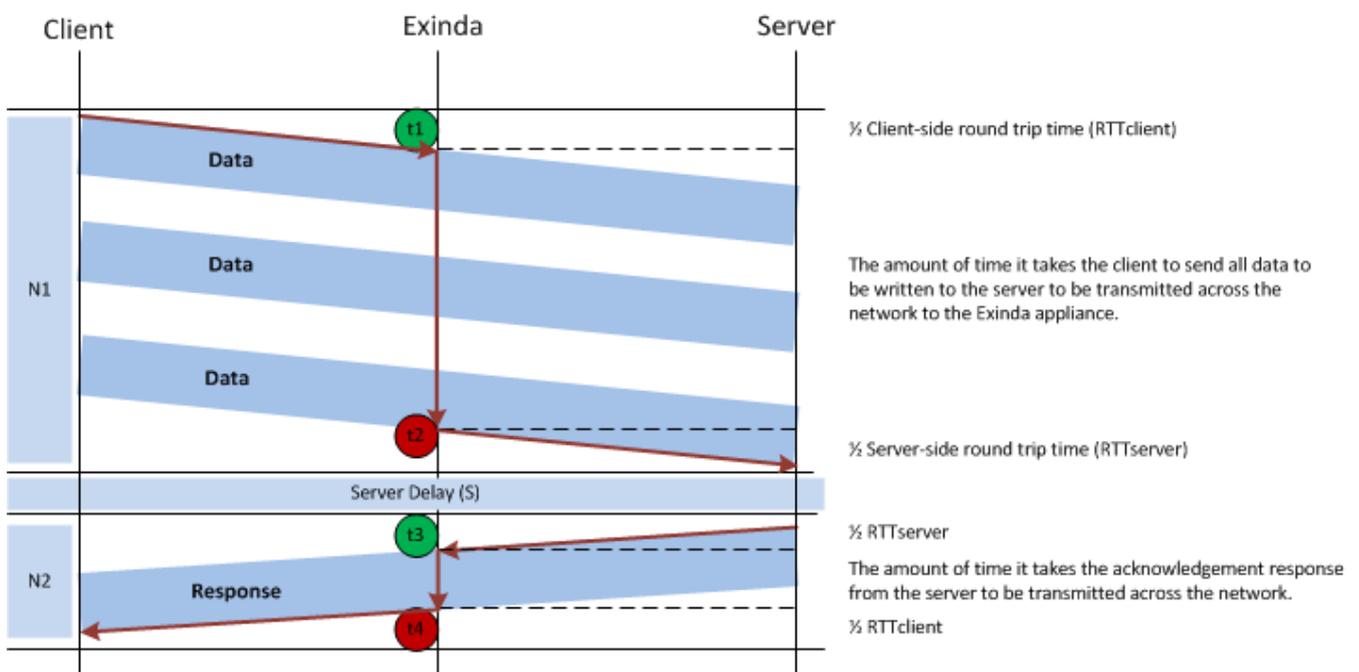
- »  $N1 = \frac{1}{2} \text{RTT}_{\text{client}} + (t_2 - t_1) + \frac{1}{2} \text{RTT}_{\text{server}}$
- »  $S = (t_3 - t_2) - \text{RTT}_{\text{server}}$
- »  $N2 = \frac{1}{2} \text{RTT}_{\text{server}} + (t_4 - t_3) + \frac{1}{2} \text{RTT}_{\text{client}}$

## Related Topics

- » Round Trip Time
- » Network and Server Delay in a Write Transaction
- » Delay Normalization
- » Packet loss

## Calculating network and server delay in a write transaction

When a client computer sends information to be written to the server, the request and response are tracked to determine how long it takes for the client to send the data to the server, and the server to send an acknowledgment of receiving the data back to the client. The diagram below shows the flow of information between the client, the Exinda appliance, and the server, and identifies the points in the transaction where time stamps are acquired.



Screenshot 174: A visual representation of a network write transaction.

### Network Delay for Write Request (N1)

- » The client sends data to be written on the server, and may be sent in a number of packets.
- » When the first data packet starts passing through the Exinda, the time stamp is noted as the beginning of the packet ( $t_1$ ).
- » When the end of the last data packet passes through the Exinda, the time stamp is noted ( $t_2$ ).  $t_2 - t_1$  = The amount of time it takes the client to send data through the Exinda appliance.
- » The server receives all the data from the client.

### Server Delay for Write Request (S)

- » There is a very small delay between receiving the data from the client and the acknowledgement that is sent from the Server back to the client. This is the Server delay (S).

### Network Delay for Write Response (N2)

- » The server's acknowledgment response to the client that the data has been received is sent.
- » When the response passes through the Exinda, the time stamp is noted ( $t_3$ ).
- » When the end of the response passes through the Exinda, the time stamp is noted ( $t_4$ ).  $t_4 - t_3$  = The amount of time it takes the server response to pass through the Exinda appliance.
- » The client receives the response from the server.

### Total Time for Write Transaction

The total transaction time for a Write transaction is calculated as  $\text{Transaction time} = N1 + S + N2$  where:

- »  $N1 = \frac{1}{2} \text{RTT}_{\text{client}} + (t_2 - t_1) + \frac{1}{2} \text{RTT}_{\text{server}}$
- »  $S = (t_3 - t_2) - \text{RTT}_{\text{server}}$
- »  $N2 = \frac{1}{2} \text{RTT}_{\text{server}} + (t_4 - t_3) + \frac{1}{2} \text{RTT}_{\text{client}}$

### Related Topics

- » [Round Trip Time](#)
- » [Network and Server Delay in a Read Transaction](#)
- » [Delay Normalization](#)
- » [Packet loss](#)

### Calculating network efficiency using packet loss and retransmitted data

Packet loss occurs when one or more packets within a transmission are successfully sent, but fail to arrive at the destination. Packet loss can be caused by a variety of factors including network congestion, faulty network components such as hardware or drivers, or corrupted packets within the transmission. If the transmission experiences packet loss, it may cause the following:

- » Jitter in video conferences
- » Gaps in audio during VoIP communications
- » Performance issues when streaming media

To recover from packet loss, data must be retransmitted to the destination to complete requests successfully. The amount of data retransmitted per flow is used to calculate the Network Efficiency metric.

- »  $\text{Efficiency} = 100\% * (\text{transferred} - \text{retransmitted}) / \text{transferred}$
- »  $\text{Network Loss} = 100 - \text{Efficiency}$

#### NOTE

Network loss, not efficiency, is used when calculating APS.

### Related topics

- » [Round Trip Time](#)
- » [Network and Server Delay in a Read Transaction](#)
- » [Network and Server Delay in a Write Transaction](#)
- » [Delay Normalization](#)

## 0.14 Configuring for common use cases and scenarios

The following list of topics demonstrate some common ways you can use the features of your Exinda Appliance to monitor, manage and control network traffic.

---

|                                                                                                 |     |
|-------------------------------------------------------------------------------------------------|-----|
| 0.14.1 Monitoring and controlling traffic in a captive portal system .....                      | 398 |
| 0.14.2 Backhauling Internet traffic .....                                                       | 403 |
| 0.14.3 Setting and enforcing quotas .....                                                       | 405 |
| 0.14.4 Creating Applications from DSCP-marked traffic (like Riverbed accelerated traffic) ..... | 409 |
| 0.14.5 Clustering and high availability .....                                                   | 410 |
| 0.14.6 Controlling anonymous proxy traffic .....                                                | 418 |

---

### 0.14.1 Monitoring and controlling traffic in a captive portal system

#### Integrating your Exinda Appliance with a captive portal

The Exinda Appliance can be a part of a captive portal system by using the HTTP Redirect policy to redirect unauthenticated traffic to the login page of your captive portal system.

When users have authenticated with your captive portal, your portal system sends the login user information and the group to which you want the user to belong, to the Exinda appliance through the appliance Active Directory API. You may choose to have a single authenticated user group or you may choose to have several user groups to indicate a level of service, such as Gold, Silver, and Bronze service levels.

On the Exinda Appliance, you need to create user-group network objects that link with these Active Directory groups, which then identify the traffic as belonging to authenticated users. In the Optimizer policy tree, you need to create policies for the authenticated traffic using the user group network objects and you will also need to create a policy to redirect unauthenticated HTTP traffic to your captive portal, and another policy to block other types of unauthenticated traffic. You should note that you should ensure that DNS traffic for the unauthenticated users is not blocked.

Since the Exinda Appliance matches traffic to the filters in the policies (and virtual circuits) from the top of the Optimizer policy tree, you need to ensure that the most specific filters appear first in the tree. The policies should appear in the following order.

1. Authenticated traffic, to which you may want to apply various policies
2. Remaining (unauthenticated) HTTP traffic, which you will redirect to a login URL using the HTTP Redirect policy
3. Remaining (unauthenticated, non-HTTP) traffic, which you may want to block or throttle.

#### NOTE

When redirecting traffic to a URL, ensure that traffic to that URL does not fall into the redirect policy as this will cause a redirect loop. This will not be an issue if the captive portal is on your local network such that traffic from the unauthenticated users can get to the captive portal directly without going through the Exinda appliance.

If the captive portal is positioned such that the unauthenticated users traffic goes through the Exinda appliance, then add a policy before the redirect policy that will capture the traffic to the captive portal and will allow it through.

If you have a single authenticated users group, you can create separate virtual circuits for the authenticated and unauthenticated users. Multiple policies can then be added to the authenticated user virtual circuit to manage these users.

Note that both virtual circuits can be specified as requiring 100% of the available bandwidth with automatic handling of over-subscription. This allows the bandwidth to be shared. However, since unauthenticated traffic requires very little bandwidth, the authenticated users will get almost all of the bandwidth.

**Virtual Circuit 10 - Authenticated Users (100% to / from 'Users')**

- 10 **Recreational - Limit Low 2%-10%** (Optimize 2% - 10%, Priority 10)
- 20 **Voice - Guarantee Critical 15%-100%** (Optimize 15% - 100%, Priority 1)
- 30 **ALL - Guarantee Low 5%-100%** (Optimize 5% - 100%, Priority 7, Application Acceleration)

Order:  Policy: ALL - Accelerate

[Create New Policy...](#)

**Virtual Circuit 20 - Unauthenticated traffic (100% to / from 'ALL')**

- 10 **Redirect HTTP** (HTTP Redirect - http://mysystem.mycompany.com/login)
- 15 **Allow DNS** (Ignore)
- 20 **Block All** (Discard)

Order:  Policy: ALL - Accelerate

[Create New Policy...](#)

Screenshot 175: Using virtual circuits to filter for authenticated traffic

However, if you have multiple groups of authenticated users where you want one group to have preferential treatment over the other groups, you need to create a series of policies within a single virtual circuit where each policy explicitly filters in favour of the authenticated users.

Using virtual circuits to filter authenticated traffic is easier if you have many policies that you want applied to the authenticated traffic. However, since only policies, not virtual circuits, can ensure preferential treatment of the traffic, you need to use policies to filter the authenticated user groups.

**Virtual Circuit 30 - User Traffic (100% to / from 'ALL')**

- 10 **Gold Authenticated Users** (Optimize 50% - 100%, Priority 1)
- 20 **Silver Authenticated Users** (Optimize 30% - 100%, Priority 3)
- 30 **Bronze Authenticated Users - Choke streaming** (Optimize 1% - 10%, Priority 10)
- 40 **Bronze Authenticated Users - All traffic** (Optimize 19% - 100%, Priority 5)
- 50 **Redirect HTTP** (HTTP Redirect - http://mysystem.mycompany.com/login)
- 60 **Allow DNS** (Ignore)
- 70 **Block All** (Discard)

Screenshot 176: Using the policies themselves to filter for authenticated traffic

## Related Topics

- » [Configure Network User Groups](#)
- » [Creating a Virtual Circuit for Authenticated Users](#)
- » [Creating a Virtual Circuit for Unauthenticated Users](#)
- » [Creating policies to filter for Authenticated Users with no Virtual Circuit available](#)
- » [Creating policies that redirect traffic](#)
- » [Creating policies that block unwanted unauthenticated traffic](#)

## Creating a virtual circuit for authenticated users

Use the following instructions to create a virtual circuit that filters for authenticated users.

1. Go to **Configuration > Optimizer**.

The page to the right refreshes.

2. If not already open, select the **Optimizer** tab.
3. Click **Create New Virtual Circuit**.
4. In the **Virtual Circuit Name** field, type a suitable name.
5. Specify the **Virtual Circuit Bandwidth** to be 100%.
6. Specify **Oversubscription** handling to be **Automatic**.
7. Specify the **Network Object** to be the network object for the Active Directory authenticated users group.
8. Click **Add New Virtual Circuit**.

The page refreshes and the new VC appears on the **Optimizer** tab.

### Related Topics

- » [Creating a Virtual Circuit for Unauthenticated Users](#)
- » [Creating policies to filter for Authenticated Users with no Virtual Circuit available](#)
- » [Creating policies that redirect traffic](#)
- » [Creating policies that block unwanted unauthenticated traffic](#)

## Creating a virtual circuit for unauthenticated users

Use the following instructions to create a virtual circuit to filter for unauthenticated users.

To create the virtual circuit:

1. Go to **Configuration > .**
- The page to the right refreshes.
2. If not already open, select the **Optimizer** tab.
  3. Click **Create New Virtual Circuit**.
  4. In the **Virtual Circuit Name** field, type a suitable name.
  5. Specify the **Virtual Circuit Bandwidth** to be 100%.
  6. Specify the **Oversubscription** handling to be **Automatic**.
  7. Set each of the **Filter Options** to **ALL**, including the **Network Object** field.
  8. Click **Add New Virtual Circuit**.

The page refreshes and the new VC appears on the tab.

### Related Topics

- » [Creating a Virtual Circuit for Authenticated Users](#)
- » [Creating policies to filter for Authenticated Users with no Virtual Circuit available](#)

- » [Creating policies that redirect traffic](#)
- » [Creating policies that block unwanted unauthenticated traffic](#)

### Creating policies to filter for authenticated users with no virtual circuit available

Use the following instructions to create policies that filter for authenticated users when a virtual circuit is not available.

1. Go to **Configuration > Optimizer**.

The right side of the screen refreshes.

2. Select the **Policies** tab.

**Edit Policy**

Policy Name:

Schedule:

Action:

Policy Enabled:

**Filter Rules:**

| VLAN                             | Source                                    | Direction                         | Destination                      | ToS/DSCP                         | Application                            |
|----------------------------------|-------------------------------------------|-----------------------------------|----------------------------------|----------------------------------|----------------------------------------|
| <input type="text" value="ALL"/> | <input type="text" value="Bronze Users"/> | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="Streaming"/> |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |

**Guaranteed Bandwidth:**   %

**Burst (Max) Bandwidth:**  %

**Burst Priority:**

**Acceleration:**

**Packet Marking**

Screenshot 177: Editing a policy.

3. In the **Policy Name** field, create a suitable name for the policy.

4. From the **Action** drop-down list, select **Optimize**.

5. Set the **Guaranteed Bandwidth**, **Acceleration**, and **Packet Marking** settings as needed.

6. Under **Filter Rules**, do the following:

- Set the traffic **Source** to be the network object for the Active Directory authenticated users group.
- Set traffic **Direction** to **Both**
- Set any other options that you need, such as setting a particular application or application group.

#### Related Topics

- » [Creating a Virtual Circuit for Authenticated Users](#)
- » [Creating a Virtual Circuit for Unauthenticated Users](#)
- » [Creating policies that redirect traffic](#)
- » [Creating policies that block unwanted unauthenticated traffic](#)

### Creating policies that redirect traffic

Use the following instructions to create policies that redirect unauthenticated users to an identified URL.

**Edit Policy**

Policy Name:       Redirect URL:

Schedule:

Action:

Policy Enabled:

| VLAN                             | Source                           | Direction                         | Destination                      | ToS/DSCP                         | Application                           |
|----------------------------------|----------------------------------|-----------------------------------|----------------------------------|----------------------------------|---------------------------------------|
| <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="HTTP"/>     |
| <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="HTTP-ALT"/> |
| <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="HTTPS"/>    |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                  |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                  |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                  |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                  |

Screenshot 178: Configuring a HTTP Redirect policy

1. Go to **Configuration > Optimizer**.

The right side of the screen refreshes.

2. Select the **Policies** tab.

3. In the **Policy Name** field, create a suitable name for the policy.

4. From the **Action** drop-down list, select **HTTP Redirect**.

5. In the **Redirect URL** field, type the URL to which that unauthenticated user directed.

Traffic matching this policy is forwarded to the specified URL, which causes the specified URL to be presented to the client.

6. Set up the **Filter Rules**.

**NOTE**

The only allowable applications are HTTP, HTTP-ALT, and HTTPS. The recommendation is to add three filter rules, one for each of these applications. Similar to other policy configurations, in conjunction with the application, if needed you can specify VLAN, Host/Direction/Host, or ToS/DSCP.

7. When complete, click **Create New Policy**.

**Related Topics**

- » [Creating policies to filter for Authenticated Users with no Virtual Circuit available](#)
- » [Creating policies that block unwanted unauthenticated traffic](#)
- » [Creating a Virtual Circuit for Authenticated Users](#)
- » [Creating a Virtual Circuit for Unauthenticated Users](#)

**Creating policies that block unwanted unauthenticated traffic**

Use the following instructions to block any unauthenticated traffic that is not already identified in other policies.

To create the policy:

1. Go to **Configuration > Optimizer**.

The right side of the screen refreshes.

2. Select the **Policies** tab.

**Edit Policy**

Policy Name:

Schedule:

Action:

Policy Enabled:

Guaranteed Bandwidth:  %

Burst (Max) Bandwidth:  %

Burst Priority:

Acceleration:

Packet Marking

| VLAN                             | Source                                    | Direction                         | Destination                      | ToS/DSCP                         | Application                            |
|----------------------------------|-------------------------------------------|-----------------------------------|----------------------------------|----------------------------------|----------------------------------------|
| <input type="text" value="ALL"/> | <input type="text" value="Bronze Users"/> | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="Streaming"/> |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |
| <input type="text"/>             | <input type="text"/>                      | <input type="text" value="Both"/> | <input type="text"/>             | <input type="text"/>             | <input type="text"/>                   |

Screenshot 179: Editing policies.

3. In the **Policy Name** field, type a suitable name for the policy.

4. From the **Action** drop-down list, select **Discard**.

5. If you want to block all unauthenticated traffic, then do not check the **Discard only the first packet of the connection** checkbox.

6. Do not set any filter rules.

#### Related Topics

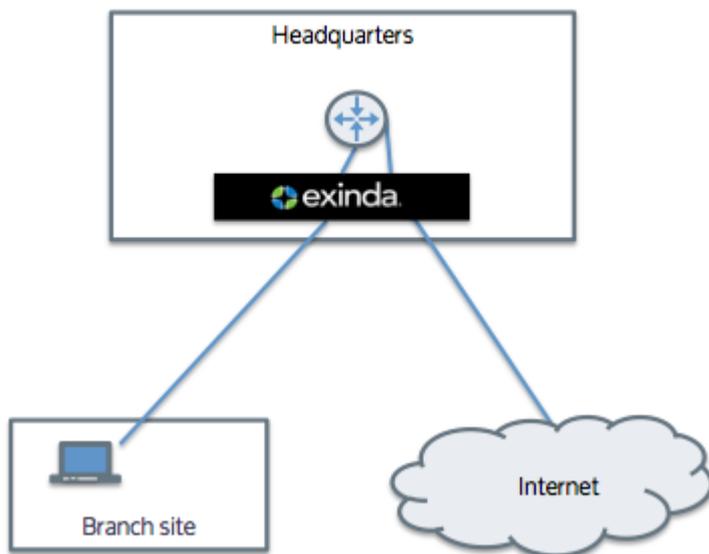
- » [Creating a Virtual Circuit for Authenticated Users](#)
- » [Creating a Virtual Circuit for Unauthenticated Users](#)
- » [Creating policies to filter for Authenticated Users with no Virtual Circuit available](#)
- » [Creating policies that redirect traffic](#)

### 0.14.2 Backhauling Internet traffic

A backhauled topology transports traffic between a remote site and the Internet via a centralized backbone, such as the headquarters of an organization .

Because of the layout, the traffic may go through an Exinda Appliance at the headquarters twice. The traffic flows from the client through the appliance, turns around at a router, and goes back through the appliance to the destination.

Backhauling traffic introduces various issues that need to be considered when configuring your Exinda Appliance.



## Multi-Bridge Circuits

Backhauling is problematic for accelerated traffic because you do not want to re-accelerate the traffic. The dual bridge bypass mode prevents traffic from being re-accelerated. By default, this mode is enabled.

In a backhauled topology, the traffic is transported over multiple bridges and you want each bridge to treat traffic differently so that the traffic is accelerated on the way in from the branch site and bypasses acceleration on the way out to the Internet. The dual bridge bypass mode ensures that only one bridge handles acceleration. The traffic is passed through the second bridge with no additional acceleration handling, that is, the traffic bypasses acceleration handling on the second bridge.

Consider traffic passing through the appliance from the WAN to the Internet as

WAN <-> br10 <-> router <-> br20 <-> internet

- » For an incoming accelerated connection, the acceleration processing happens on the first WAN interface to see the SYN. In this case, where the connection is going to the Internet, the accelerated connection is processed on br10. That same traffic when seen on br20 is then passed through untouched (the traffic bypasses br20).
- » For an outgoing accelerated connection, the acceleration processing will happen on the first bridge with a matching acceleration policy. In this case, where the connection is from the Internet, if there is an acceleration policy on br20, the acceleration policy is enacted on br20 and the traffic is left untouched on br10. Or if the acceleration policy was on br10, then the traffic is passed through br20 untouched and accelerated on br10.

### NOTE

There are separate dual bridge bypass settings for acceleration and for monitoring. To learn more about dual bridge bypass, see [Dual Bridge Bypass](#).

## Sending Exinda-Specific Option Codes to the Internet

Consider backhauled traffic passing from the client to the Internet through two accelerated Exinda appliances - one at the client's branch site and one at headquarters:

Client -> Branch Exinda -> WAN -> Headquarters Exinda -> Internet

Exinda Appliances use TCP option 30 on the SYN packets to detect which appliances will participate in accelerating the TCP connection. Normally, the Exinda appliance at the headquarters would send a SYN with an attached TCP option 30 to the server on the Internet just in case there is another Exinda appliance closer to the server. The **End Acceleration**

feature enabled on the Exinda appliance at headquarters identifies that appliance as the end of the acceleration chain and will prevent the appliance from sending the TCP option 30 on the SYN packet to the Internet.

If, however, you have a hub-and-spoke topology, where the Exinda Appliance at headquarters could either be the end of the acceleration chain when backhauling Internet traffic or it could be an Exinda in the middle for traffic transported between distributed branches via the headquarters, you will not want to enable this feature.

To learn more about this feature, see [Configure TCP Acceleration](#).

### 0.14.3 Setting and enforcing quotas

Quotas are an effective way to enforce fair sharing of the network or to ensure customers receive only the amount of access to the network for which they have paid. Quotas can enforce caps based on a data transfer amount or based on the amount of time on the network. After the quota has been reached, a variety of actions could take place, such as throttling or blocking all data, or throttling only particular types of traffic, or redirecting the user to a particular webpage.

To support quota enforcement scenarios, you need to configure the following:

1. Create an adaptive response limit object to define how the quota is measured and to identify the users that have exceeded their quota by using a named network object. The adaptive response object can specify whether to set a network-traffic data-volume limit or a time limit. The adaptive response object identifies the traffic that is monitored against the specified quota as a network object. The network object can either be based on IP addresses, or based on Active Directory users or user groups. The adaptive response object tracks those that have exceeded their quota by dynamically adding them to a named network object.
2. Add a policy (or policies) to the Optimizer policy tree for those who are over their limit. The policy that addresses those that have exceeded their quota is defined according to your business needs. You can choose to throttle their traffic or block it entirely. When they have HTTP traffic, you can also choose to redirect them to a webpage that you host or respond with a webpage that the Exinda Appliance hosts. If needed you can combine these, such that the first policy filters for HTTP traffic and then shows a webpage, but then other types of traffic are caught by a second policy that blocks the remaining traffic.
3. Add policies to the Optimizer policy tree for those under the limit. The remaining policies define how to deal with the traffic of the users who have not yet exceeded their quota.

#### NOTE

Since the Exinda Appliance attempts to match the traffic to the filters in the policies (and virtual circuits) in the top-down order defined in the Optimizer policy tree, you need to set up the series of policies with the most specific filter criteria appearing first in the policy tree, which means the policies should appear in the following generalized order.

- » Those who have exceeded their quota and have HTTP traffic
- » Those who have exceeded their quota and have other types of traffic
- » Remaining traffic (that is, those who have not exceeded their quota)

To learn more about the individual components needed for quota enforcement, see [Configure Adaptive Response Limit](#), [Configure Network Objects](#), [Configure Network User Groups](#), [Configure HTML Response Object](#), [Policy Tree](#), and [Policies](#).

#### Example: Each user has a 10GB capped data quota

Consider an educational institution that has a group of students who have IP addresses in the subnet 192.168.0.0/16. Each student is allowed 10GB data transfer (uploads and downloads) per month. After the limit is reached, they are allowed no more data.

1. Create a network object to represent the students.

OPTION 1: Create a static network object using the **Configuration > Objects > Network > Network Objects** page.

**Add New Network Object**

Name:

Location:

Subnet Report:

Subnets:

| IP Network Address / Mask Length         |                                   |
|------------------------------------------|-----------------------------------|
| <input type="text" value="192.168.0.0"/> | <input type="text" value="/ 16"/> |
| <input type="text"/>                     | <input type="text"/>              |
| <input type="text"/>                     | <input type="text"/>              |
| <input type="text"/>                     | <input type="text"/>              |

Screenshot 180: Adding a network object.

OPTION 2: Create a network user group object using the **Configuration > Objects > Users & Groups > Network Groups** page.

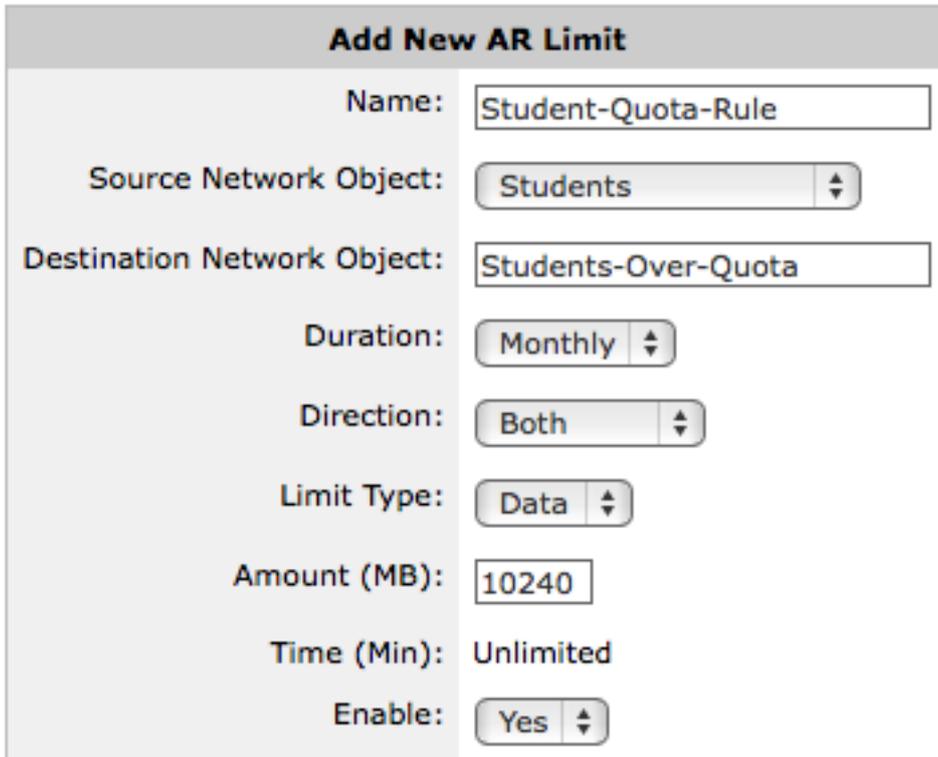
[0-9](#) | [\[A\]](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#) | [Y](#) | [Z](#) | [Others](#) | [ALL](#)

| <input type="checkbox"/> | User (Domain)     | IP | Network Object                      |
|--------------------------|-------------------|----|-------------------------------------|
| <input type="checkbox"/> | Students (EXANET) |    | <input checked="" type="checkbox"/> |
| <input type="checkbox"/> | Tac (EXANET)      |    | <input checked="" type="checkbox"/> |
|                          | :                 |    | :                                   |

Screenshot 181: Choosing the user (domain).

2. Create an adaptive response limit object that defines the 10GB limit as well as the destination dynamic network object that will contain the students who exceeded their quota using the **Configuration > Objects > Adaptive**

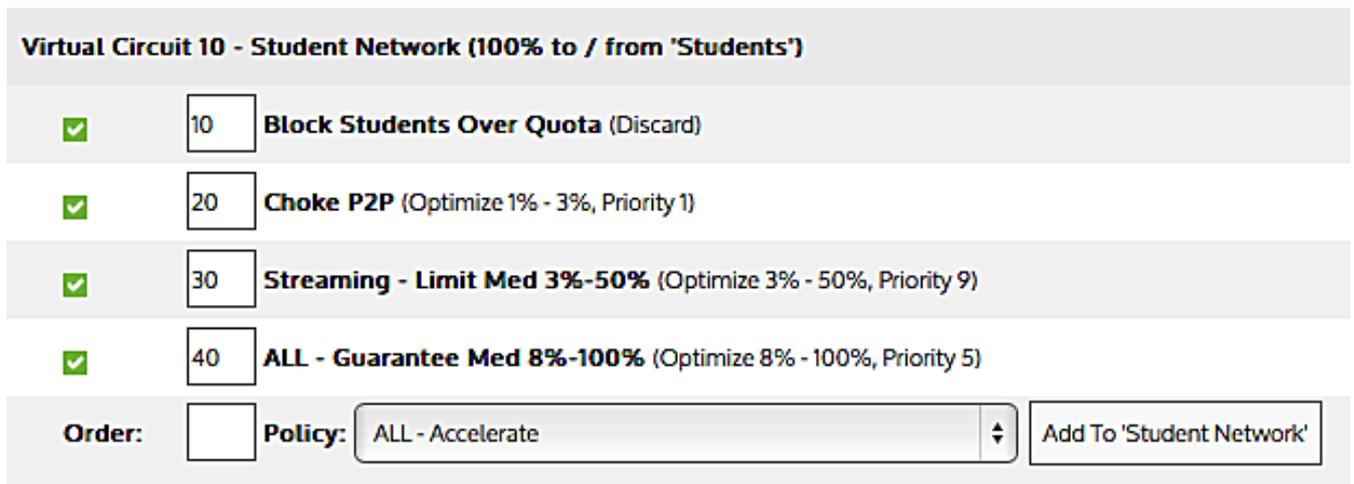
Response page.



|                                    |                                                  |
|------------------------------------|--------------------------------------------------|
| <b>Name:</b>                       | <input type="text" value="Student-Quota-Rule"/>  |
| <b>Source Network Object:</b>      | <input type="text" value="Students"/>            |
| <b>Destination Network Object:</b> | <input type="text" value="Students-Over-Quota"/> |
| <b>Duration:</b>                   | <input type="text" value="Monthly"/>             |
| <b>Direction:</b>                  | <input type="text" value="Both"/>                |
| <b>Limit Type:</b>                 | <input type="text" value="Data"/>                |
| <b>Amount (MB):</b>                | <input type="text" value="10240"/>               |
| <b>Time (Min):</b>                 | <input type="text" value="Unlimited"/>           |
| <b>Enable:</b>                     | <input type="text" value="Yes"/>                 |

Screenshot 182: Adding an AR limit.

3. Configure the policy tree such that the students over quota are blocked from further data. In the virtual circuit that will process the student data, create a policy that will block the students who have exceeded their quota and ensure that it is first in the virtual circuit. The rest of the policies can manage the traffic however you like, perhaps choking P2P and throttling streaming.



| Virtual Circuit 10 - Student Network (100% to / from 'Students') |                                                                        |
|------------------------------------------------------------------|------------------------------------------------------------------------|
| <input checked="" type="checkbox"/>                              | 10 <b>Block Students Over Quota</b> (Discard)                          |
| <input checked="" type="checkbox"/>                              | 20 <b>Choke P2P</b> (Optimize 1% - 3%, Priority 1)                     |
| <input checked="" type="checkbox"/>                              | 30 <b>Streaming - Limit Med 3%-50%</b> (Optimize 3% - 50%, Priority 9) |
| <input checked="" type="checkbox"/>                              | 40 <b>ALL - Guarantee Med 8%-100%</b> (Optimize 8% - 100%, Priority 5) |
| <b>Order:</b>                                                    | <input type="text" value=""/>                                          |
| <b>Policy:</b>                                                   | <input type="text" value="ALL - Accelerate"/>                          |
| <input type="button" value="Add To 'Student Network'"/>          |                                                                        |

Screenshot 183: Setting parameters for a virtual circuit.

4. Create an HTML Response object that defines what the webpage will look like once the shoppers have exceeded 2 hours of usage. See the **Configuration > Objects > HTML Response** page.

### Add New HTML Response

Name:

HTML Response:

```
<html>
<head>
<style>
hr {color:sienna;}
p {
margin-left:20px;
font-family:Verdana;
}
</style>
</head>
<body style="background:#e3e9ef">
<div style="padding:30px;margin-left:auto;margin-right:auto;width:80%;background:white">
<p style="color:#79c157;font-size:20pt;">OOPS! YOU'VE REACHED </br>
<span style="color:#3da58e;font-size:20pt;">YOUR DAILY 2 HOUR LIMIT</span></p>
<hr>
<p style="color:#515151;">Thank you for visiting the Springfield Mall. Please come again.</p>
</div>
</body>
</html>
```

Screenshot 184: Adding an HTML response.

5. Configure the policy tree such that the shoppers over quota are presented with a HTML response web page when accessing web traffic and all other data access for those shoppers is blocked, followed by policy for shoppers who have had access for less than 2 hours.

|                                     |    |                                                                                                                                    |
|-------------------------------------|----|------------------------------------------------------------------------------------------------------------------------------------|
| <input checked="" type="checkbox"/> | 10 | Redirect students over quota to quota-policy webpage (HTTP Redirect - http://mysystem.mycompany.com/residence/network-data-policy) |
| <input checked="" type="checkbox"/> | 20 | Block students over quota (Discard)                                                                                                |
| <input checked="" type="checkbox"/> | 30 | ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)                                                                       |

Screenshot 185: Setting policy tree parameters.

To create the policy that presents the HTML response web page:

### Add New VC Policy

Policy Name:  HTML Response Objects:  Select the HTML response to send

VC Policy Number:

Schedule:

Action:

Policy Enabled:

Filter Rules:

| VLAN                 | Source                                           | Direction                         | Destination          | ToS/DSCP             | Application                           |
|----------------------|--------------------------------------------------|-----------------------------------|----------------------|----------------------|---------------------------------------|
| <input type="text"/> | <input type="text" value="Shoppers-over-quota"/> | <input type="text" value="Both"/> | <input type="text"/> | <input type="text"/> | <input type="text" value="HTTP"/>     |
| <input type="text"/> | <input type="text" value="Shoppers-over-quota"/> | <input type="text" value="Both"/> | <input type="text"/> | <input type="text"/> | <input type="text" value="HTTP-ALT"/> |
| <input type="text"/> | <input type="text" value="Shoppers-over-quota"/> | <input type="text" value="Both"/> | <input type="text"/> | <input type="text"/> | <input type="text" value="HTTPS"/>    |
| <input type="text"/> | <input type="text"/>                             | <input type="text" value="Both"/> | <input type="text"/> | <input type="text"/> | <input type="text"/>                  |

Screenshot 186: Adding a new VC policy.

1. Select **Return HTML Response** as the policy action.
2. Select the HTML Response Object that you created in step 3.

Web traffic matching this policy will be sent back an HTML response with the contents of the HTML Response object, which will cause the a web page to be presented to the client.

3. Type the **Filter Rules**.

The only allowable applications are HTTP, HTTP-ALT, and HTTPS. It is recommended to add three filter rules - one for each of these applications.

For each of the filter rules specify the **Filter traffic Source** to be the destination network object that was created as part of the adaptive response limit object and specify the **Filter traffic Direction** to be **Both**.

To create a policy that blocks remaining traffic for the shoppers who are over quota:

1. Select **Discard** as the policy action.
2. If you want to block all traffic, then do not check the **Discard only the first packet of the connection** checkbox.
3. Type the **Filter Rules** specifying the shoppers over quota network object.

Specify the **Filter traffic Source** to be the destination network object that was created as part of the adaptive response limit object and specify the **Filter traffic Direction** to be **Both**.

#### 0.14.4 Creating Applications from DSCP-marked traffic (like Riverbed accelerated traffic)

Some times you might find that the Exinda Appliance is located between two other types of network appliances that modify the traffic in a way that obscures the visibility of the Exinda appliance into the traffic. An example of this would be having the Exinda Appliance located between two Riverbed appliances doing acceleration. The Riverbed appliances will send all of the accelerated traffic over port 7800. The Exinda Appliance will then see all accelerated traffic as a single application even if there are multiple types of application traffic within that single Riverbed tunnel.

In this situation, DSCP marking of traffic can restore the visibility of different traffic types to the Exinda Appliance. Often these other network appliances will have the ability to apply DSCP marks to traffic based on the type of application traffic they are accelerating. For instance, these other appliances may be configured to apply DSCP mark 2 to all HTTP traffic and DSCP mark 3 to all SMB traffic. Then, on the Exinda appliance, you can create applications based on the DSCP marks.

Consider a Riverbed appliance and Exinda appliance in the same environment. By default the Exinda Appliance will report the accelerated traffic streams from the Riverbed appliance as a single traffic type. If you configure the Riverbed appliances to add DSCP markings based on specific traffic types, then on the Exinda appliance create applications based on the DSCP markings and the Riverbed tunnel port number. Each of these applications appears on the Exinda Appliance like other built in applications. That is, these new Riverbed accelerated applications appear on the applications monitor and on the real-time screen. They can also be added to application groups and can have policy applied to them.

##### VERSION INFO

The ability to create applications based on DSCP-marked traffic is available in 7.0.1 and above (in the 7.0 firmware product line) and in 6.4.5 and above (in the 6.4 firmware product line).

To create applications based on DSCP marked applications in Riverbed-accelerated traffic:

1. On your Riverbed appliance, add DSCP markings for each traffic type that you are interested in.

See your Riverbed documentation to learn how to do this.

2. On your Exinda appliance, add an application for each DSCP-marked application type.

On the **Add new Application** form, type the DSCP mark number in the **DSCP** field and optionally fill in the **Ports/Protocols** field by selecting **TCP Port Range** and type 7800.

For instance, you could create a Riverbed SMB application and a Riverbed HTTP application.

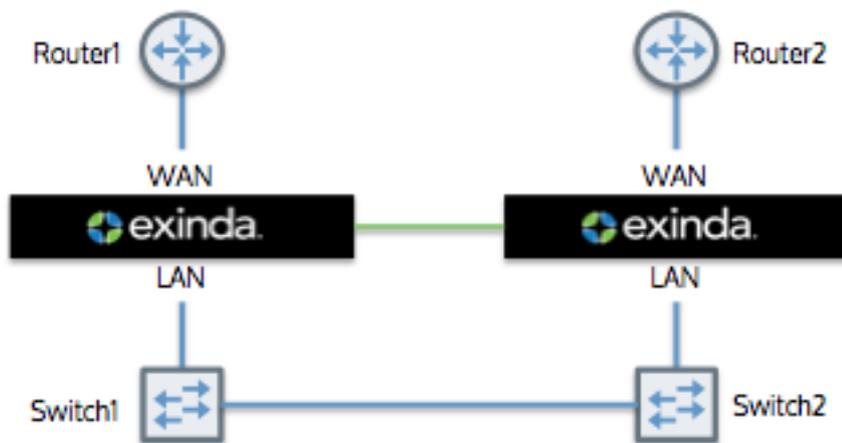
Note that application names must be unique. SMB and HTTP are already named applications on the application, so you will have to name these new applications differently.

For more information, refer to [Predefined Applications and Supported L7 Signatures](#) (page 756).

3. Use these applications are you would the predefined applications.

### 0.14.5 Clustering and high availability

Exinda Appliances can be seamlessly deployed into high availability and load balanced network architectures. The Exinda clustering options allows multiple Exinda Appliances to be deployed into these architectures and to operate as if they were a single appliance. A typical deployment topology is illustrated below.



Screenshot 187: A typical clustering topology.

In this example, there are two physical links: from Router1 to Switch 1 and from Router2 to Switch2. An Exinda appliance is deployed between each switch and router, and a cable is connected between the two appliances for synchronization.

The appliances share configuration, monitoring information, and optimizer policies, as if they were a single appliance. For the purposes of configuration synchronization, one appliance is considered as the master appliance, the other appliance(s) are the slaves. The master appliance shares its configuration with the slaves.

Once the appliances are configured, the appliances will auto-discover each other and one will be elected as the Cluster Master. All configuration must be done on the Cluster Master, so when accessing the cluster, it is best to use the Cluster Master IP address when managing a cluster.

#### CAUTION

When upgrading the firmware of appliances that are part of a cluster, Exinda recommends that you break the cluster before starting the upgrade (Either by disconnecting the cluster link or by clearing the "Cluster" option for the appropriate interface). After all appliances in the cluster have been upgraded to the same firmware, the cluster can be put back together.

#### IMPORTANT

Appliances in a cluster must be the same model and have the same number of bridge pairs. The time on all the appliances must use NTP to ensure that all appliances have exactly the same time.

### Shared Configurations

As part of normal cluster operations, the Cluster Master synchronizes parts of the system configuration to all other nodes in the cluster. Some configuration is specific to an individual appliance (for example IP addressing and licensing), however, most of the system configuration is synchronized throughout the cluster, including:

- » Optimizer Policies (see note below)
- » Network Objects

- » Protocol and VLAN Objects
- » Applications and Application Groups
- » Optimizer Schedules
- » Monitoring and Reporting Settings
- » SDP and Remote SQL Settings
- » Time-zone and NTP Settings
- » Logging Settings
- » Email and SNMP Notification Settings

#### **NOTE**

Optimizer policies are also implemented globally across all cluster nodes. For example, if there were a single policy to restrict all traffic to 1Mbps, this would be applied across all cluster nodes. So, the sum of all traffic through all cluster nodes would not exceed 1Mbps.

### Shared Monitoring Information

All appliances in the cluster can be monitored. Most monitoring information is shared across the cluster. Some reports don't make sense to share (e.g. Interface reports); however, most reports are synchronized, including:

- » Realtime
- » Network
- » AQS
- » Applications and URLs
- » Hosts
- » Conversations
- » Subnets

#### **NOTE**

When monitoring information is shared across the cluster, the timestamp on this information is not shared across the cluster. New timestamps are added to the data when it enters the other appliances in the cluster. If there is a delay in sharing this information, which could be due to the appliances in the cluster being separated physically by a great distance or by not providing enough bandwidth between the clustered appliances, the reports may not appear similar on the different appliances in the cluster.

### When a cluster node fails

In the event that a node in the cluster fails, is rebooted, or powered off, bypass mode is enacted and traffic passes though unaffected. When the appliance comes back online, the node is updated with the latest configuration settings from the Cluster Master and normal operations resume. Monitoring and reporting information during the downtime is not synchronized after the fact.

In the event that the Cluster Master fails, is rebooted, or powered off, a new Cluster Master is automatically elected. The offline node (previously the Cluster Master) is treated as a regular offline node. When it comes back online, it does not necessarily become the Cluster Master.

## Cluster Terminology

| Term                | Meaning                                                                                                                                                                                                    |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster             | A group of Exinda appliances (cluster nodes) configured to operate as a single Exinda appliance.                                                                                                           |
| Cluster External IP | An IP address assigned to the management port of the cluster master. Whichever node is the cluster master has this IP address assigned to its management port.                                             |
| Cluster Node        | An Exinda appliance that is a member of a cluster.                                                                                                                                                         |
| Cluster Interface   | The physical interface that a node in the cluster uses to connect to other cluster nodes (also referred to as the HA or AUX interface).                                                                    |
| Cluster Internal IP | A private IP address assigned to each cluster node that enables communication with other nodes in the cluster.                                                                                             |
| Cluster Master      | The node responsible for synchronizing configuration changes with all other cluster nodes. The cluster master is automatically elected. Configuration changes should only be made from the cluster master. |
| ID                  | The cluster-assigned unique identifier for each node.                                                                                                                                                      |
| Management IP       | The cluster management (alias) address. The cluster is always reachable at this address as long as at least one node is online.                                                                            |
| Role                | The current role of a given node within the cluster (master or standby).                                                                                                                                   |
| State               | The state (online or offline) of a given node.                                                                                                                                                             |

## Create a cluster of Exinda Appliances

Configuring the appliances in the network to behave as a cluster, allowing for high availability and failover, involves two steps:

1. Adding Exinda Appliances to the cluster.
2. Specifying what data is synchronized between the cluster members.

Once the appliances are configured, the appliances will auto-discover each other and one will be elected as the Cluster Master. All configuration must be done on the Cluster Master, so when accessing the cluster, it is best to use the Cluster Master IP address when managing a cluster.

### CAUTION

When upgrading the firmware of appliances that are part of a cluster, Exinda recommends that you break the cluster before starting the upgrade (Either by disconnecting the cluster link or by clearing the "Cluster" option for the appropriate interface). After all appliances in the cluster have been upgraded to the same firmware, the cluster can be put back together.

To add Exinda Appliances to the cluster

Configure the appliances with an IP address used within the cluster, as well as the IP address of the cluster master.

1. Click **Configuration > System > Network > IP Address**.
2. In eth1, type the management port IP address of the appliance in the **Static Addresses** field.
3. In eth2, select **Cluster**, and type the internal IP address for this node of the cluster in the **Static Addresses** field.

**NOTE**

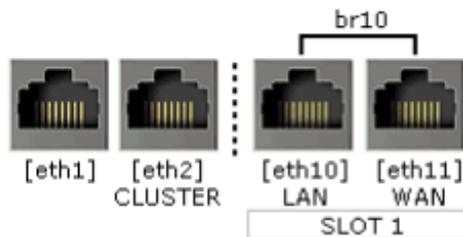
The Cluster Internal IP for each appliance in the cluster must be in the same subnet and should be an isolated and unused subnet within the network. The cluster subnet is used exclusively for communications between cluster nodes so should be private and not publicly routable.

4. In the **Cluster Master Settings** area, select eth1 and type the external address used to access the appliances.

5. Repeat these steps all each Exinda Appliance joining the cluster.

Once these settings are saved, the appliances will auto-discover each other and one will be elected as the Cluster Master. All configuration must be done on the Cluster Master, so when accessing the cluster, it is best to use the Cluster Master IP address when managing a cluster.

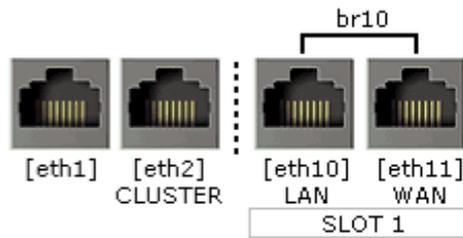
Example for a two appliance cluster:



| Interface Settings                                         |                                                                                                                                                                                              |
|------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| eth1                                                       | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP                                                                                         |
|                                                            | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                                                                           |
|                                                            | Dynamic Addresses: fe80::224:e8ff:fe69:ef23/64                                                                                                                                               |
|                                                            | Static Addresses: <input type="text" value="192.168.0.161"/> / <input type="text" value="24"/><br><input type="text"/> / <input type="text"/><br><input type="text"/> / <input type="text"/> |
| Comment: <input type="text" value="Management Interface"/> |                                                                                                                                                                                              |
| eth2                                                       | Role: <input checked="" type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP                                                                              |
|                                                            | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                                                                           |
|                                                            | Dynamic Addresses: fe80::224:e8ff:fe69:ef24/64                                                                                                                                               |
|                                                            | Static Addresses: <input type="text" value="192.168.1.1"/> / <input type="text" value="24"/><br><input type="text"/> / <input type="text"/>                                                  |
| Comment: <input type="text" value="Cluster Internal"/>     |                                                                                                                                                                                              |
| br10 <input checked="" type="checkbox"/>                   | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC                                                                                                           |
|                                                            | Dynamic Addresses: fe80::2e0:edff:fe14:5e32/64                                                                                                                                               |
|                                                            | Static Addresses: <input type="text"/> / <input type="text"/>                                                                                                                                |
|                                                            | Comment: <input type="text" value="Bridge interface"/>                                                                                                                                       |
| Gateway Settings                                           |                                                                                                                                                                                              |
| IPv4: <input type="text" value="192.168.0.1"/>             |                                                                                                                                                                                              |
| IPv6: <input type="text"/>                                 |                                                                                                                                                                                              |

**Apply Changes**

Screenshot 188: IP address configuration page on Exinda 1.



### Interface Settings

|                                     |                    |                                                   |                                      |                                 |
|-------------------------------------|--------------------|---------------------------------------------------|--------------------------------------|---------------------------------|
| <b>eth1</b>                         | Role:              | <input type="checkbox"/> Cluster                  | <input type="checkbox"/> Mirror      | <input type="checkbox"/> WCCP   |
|                                     | Autoconf:          | IPv4: <input type="checkbox"/> DHCP               | IPv6: <input type="checkbox"/> SLAAC |                                 |
|                                     | Dynamic Addresses: | fe80::224:e8ff:fe69:ef23/64                       |                                      |                                 |
|                                     | Static Addresses:  | <input type="text" value="192.168.0.162"/>        | /                                    | <input type="text" value="24"/> |
|                                     |                    | <input type="text"/>                              | /                                    | <input type="text"/>            |
|                                     | Comment:           | <input type="text" value="Management Interface"/> |                                      |                                 |
| <b>eth2</b>                         | Role:              | <input checked="" type="checkbox"/> Cluster       | <input type="checkbox"/> Mirror      | <input type="checkbox"/> WCCP   |
|                                     | Autoconf:          | IPv4: <input type="checkbox"/> DHCP               | IPv6: <input type="checkbox"/> SLAAC |                                 |
|                                     | Dynamic Addresses: | fe80::224:e8ff:fe69:ef24/64                       |                                      |                                 |
|                                     | Static Addresses:  | <input type="text" value="192.168.1.2"/>          | /                                    | <input type="text" value="24"/> |
|                                     |                    | <input type="text"/>                              | /                                    | <input type="text"/>            |
|                                     | Comment:           | <input type="text" value="Cluster Internal"/>     |                                      |                                 |
| <b>br10</b>                         | Autoconf:          | IPv4: <input type="checkbox"/> DHCP               | IPv6: <input type="checkbox"/> SLAAC |                                 |
| <input checked="" type="checkbox"/> | Dynamic Addresses: | fe80::2e0:edff:fe14:5e32/64                       |                                      |                                 |
|                                     | Static Addresses:  | <input type="text"/>                              | /                                    | <input type="text"/>            |
|                                     | Comment:           | <input type="text" value="Bridge interface"/>     |                                      |                                 |

### Gateway Settings

|       |                                          |
|-------|------------------------------------------|
| IPv4: | <input type="text" value="192.168.0.1"/> |
| IPv6: | <input type="text"/>                     |

Screenshot 189: IP address configuration page on Exinda 2.

### Cluster Master Settings

|                 |                                                                              |
|-----------------|------------------------------------------------------------------------------|
| Interface:      | <input type="text" value="eth1"/> <input type="button" value="v"/>           |
| Master Address: | <input type="text" value="192.168.0.160"/> / <input type="text" value="24"/> |

Screenshot 190: Cluster Master (External) configuration on both Exinda Appliances.

In the example above, Exinda 1 has a Management IP of 192.168.0.161 and Exinda 2 has a Management IP of 192.168.0.162. The Cluster External IP is configured as 192.168.0.160 on both appliances – regardless of which of these two appliances becomes the Cluster Master, it will be reachable on the 192.168.0.160 IP address. The Cluster Internal IP on Exinda 1 is configured as 192.168.1.1 and on Exinda 2 as 192.168.1.2.

Once these settings are saved, the appliances will auto-discover each other and one will be elected as the Cluster Master. All configuration must be done on the Cluster Master, so in this example, access the cluster via 192.168.0.160. Typically, the first appliance to come online is elected the master.

## Link State Mirroring

Even though High Availability protocols are usually handled by the Routers or Multi-Layer switches that act as the “next-hops”, the core switches connected to them also need to detect a link failure in order for the entire network to converge properly. When an Exinda appliance is located in the middle of the connection between a Switch and an HA router, and the link between the HA router and the Exinda goes down, the remaining link (Exinda to Switch) will remain up, avoiding the failure detection by the switch and creating convergence problems. In order to overcome this scenario, you should use the Link State Mirroring feature.

When this option is enabled, the Exinda appliance brings down the second port of a bridge if the first port goes down. This feature allows the Exinda appliance to sit between a WAN router and a switch without blocking detection of switch outages by the router. This is a global setting that is applied to all enabled bridges. Exinda recommends to always enable this setting when the appliance is configured in cluster mode, this option is located under **Configuration > System > Network > NICs** and is disabled by default:

**Link State Mirroring**

Link State Mirroring  Enable

[Apply Changes](#)

### NOTE

An Exinda Appliance goes into “bypass” mode by default if it is shut down. In an HA/Cluster environment, in order to maintain control and visibility on the network, the administrator might want the entire traffic to failover to a backup link if the Exinda appliance on the active one is offline. In order to do this, you need the appliance that went offline to purposefully break the traffic so the HA protocol can hand it over to the backup link. To achieve this functionality, modify the “On Failover” mode for the specific bridge to “NO LINK” mode (This option is located under Configuration > System > Network > NICs):

Use the form below to configure bypass state and failover settings.

**Note: Take care when making changes here as this appliance may become unreachable.**

| Bridge | Status | Running Mode | Enable Failover                     | On Failover                                                                                                                                                                                   |
|--------|--------|--------------|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| br2    | Active | Active ▼     | <input checked="" type="checkbox"/> | <div style="border: 1px solid #ccc; padding: 2px;">                     Bypass ▼<br/>                     Bypass<br/> <span style="border: 2px solid red; padding: 2px;">NoLink</span> </div> |

[Apply Changes](#)

Exinda recommends that you leave at least one of the cluster nodes in “Bypass” failover mode in case there is a power failure that affects the entire cluster.

The “NO LINK” option is not available on the 3062 model. For the models of the 4062 series, it is necessary to purchase an extra Expansion Card given that the on-board bridges do not come with this feature.

To specify what data is synchronized between cluster members

```
[no] cluster sync {all|acceleration|monitor|optimizer|compression}
```

- » all - Acceleration, monitor, optimizer data, and compression are synchronized. This is disabled by default.
- » acceleration - Synchronize acceleration data only

- » `monitor` - Synchronize monitor data only
- » `optimizer` - Synchronize optimizer data only
- » `compression` - Configure cluster compression settings including compression type and compression level.

To view status of all members in a cluster

1. Click **Configuration > System > Maintenance > Clustering**.
2. All appliances in the cluster are displayed.

| Clustering State |                               |                       |        |         |           |             |          |           |        |
|------------------|-------------------------------|-----------------------|--------|---------|-----------|-------------|----------|-----------|--------|
| Host ID          | External IPv4 Address         | Internal IPv4 Address | Status | Role    | Uptime    | Version     | Memory   | Operation |        |
| 0024e83dcaed     | <a href="#">192.168.0.161</a> | 192.168.1.1           | ✔      | Master  | 1h 1m 23s | 6.1.0.16836 | 2050.5MB | Shutdown  | Reboot |
| bc305bd453a8     | <a href="#">192.168.0.162</a> | 192.168.1.2           | ✔      | Standby | 1h 1m 25s | 6.1.0.16836 | 2050.5MB | Shutdown  | Reboot |

It is also possible to reboot and shutdown other nodes in the cluster from this page.

3. To identify the cluster master, the role is displayed in the list of all appliances.

#### TIP

When logged into the Web UI of a cluster node, the role of the node is also shown in the header of the user interface.

To change which appliance is the master

1. Turn clustering off on the appliance that is currently the master.
2. Wait for the standby appliance to become the master.
3. On the original master appliance, turn clustering back on. This appliance will now be the standby appliance.

## Configuring Exinda Appliances for clustering

Before configuring clustering, the Exinda Appliances must be correctly cabled. It is recommended that each appliance in the cluster be connected and configured with a dedicated management port.

In addition, clustering requires a dedicated interface for traffic that is internal to the cluster. Any interface that is not bridged or in use for another role (e.g., Mirror or WCCP) may be used.

The table below lists the suggested cluster interface for each hardware series.

| Hardware Series | Cluster Interface             |
|-----------------|-------------------------------|
| 2000/4000 (EOL) | eth1 (with Bridge 0 disabled) |
| 4060/4061       | eth2                          |
| 5000            | eth1                          |
| 6000            | eth5 (with Bridge 2 disabled) |
| 6010            | eth1                          |
| 6060            | eth2                          |
| 7000            | eth1                          |
| 8060            | eth2                          |
| 10060           | eth2                          |

| Hardware Series | Cluster Interface |
|-----------------|-------------------|
| 10063           | eth2              |
| 10064           | eth2              |
| 12063           | eth2              |
| 12064           | eth2              |

» Where there are two appliances in a cluster, the cluster interfaces may be connected directly to each other with a CAT 5 cross-over cable.

» Where there are more than two appliances in a cluster, the interface for each appliance in a cluster must be connected to a single, dedicated switch, such that each appliance can communicate with every other appliance without requiring a route (they must be on the same Layer 2 LAN segment).

### Example: Configure cluster through the CLI

Configuration using the CLI is very similar to that of the Web UI.

1. Configure a Cluster Internal address. Any interface not bound to a bridge or used in another role (e.g. Mirror or WCCP) may be used. This command will need to be run on each node in the cluster, and each with a unique Cluster Internal address.

```
(config) # cluster interface eth2
(config) # interface eth2 ip address 192.168.1.1 /24
```

2. Configure the Cluster External IP. This command should be executed on all cluster nodes using the same Cluster External IP

```
(config) # cluster master interface eth1
(config) # cluster master address vip 192.168.0.160 /24
```

3. Enable the cluster.

```
(config) # cluster enable
```

#### NOTE

Configuration changes should only be made on the Cluster Master node. The role of the node currently logged into will be displayed in the CLI prompt as shown below.

```
exinda-091cf4 [exinda-cluster: master] (config) #
```

4. It is possible to view the status of all the members of a cluster from the CLI by issuing the following command.

```
(config)# show cluster global brief
Global cluster state summary
=====
Cluster ID: exinda-default-cluster-id
Cluster name: exinda-cluster
Management IP: 192.168.0.160/24
Cluster master IF: eth1
Cluster node count: 2
ID Role State Host External Addr Internal Addr
-----
```

```

1* master online exinda-A 192.168.0.161 192.168.1.1
2 standby online exinda-B 192.168.0.162 192.168.1.2

```

## 0.14.6 Controlling anonymous proxy traffic

Users who want to bypass security policies that prevent tools from monitoring what they are doing, typically find an open and anonymous HTTPS transparent proxy. They set their browsers to proxy all web requests through this anonymous proxy. Those requests will be encrypted with https. The security tools cannot distinguish these transactions from, say, a legitimate access to a financial website.

The Exinda Appliance can identify traffic that is served through an anonymous proxy and classify it as the Anonymous Proxy application. As such, any anonymous proxy traffic appears on the real-time monitoring screen and on other monitoring reports as "Anonymous Proxy", and can then be controlled by the Optimizer.

| Inbound Applications   |                      |                   |            |                                 | Outbound Applications  |                      |                   |            |                                 |
|------------------------|----------------------|-------------------|------------|---------------------------------|------------------------|----------------------|-------------------|------------|---------------------------------|
| Application Name       | Transfer Rate (kbps) | Packet Rate (pps) | Flows      | Distribution (%)                | Application Name       | Transfer Rate (kbps) | Packet Rate (pps) | Flows      | Distribution (%)                |
| <b>Total</b>           | <b>177.348</b>       | <b>52</b>         | <b>129</b> |                                 | <b>Total</b>           | <b>78.713</b>        | <b>57</b>         | <b>131</b> |                                 |
| HTTP                   | 119.675              | 22                | 33         | <div style="width: 33%;"></div> | HTTP                   | 41.652               | 26                | 33         | <div style="width: 33%;"></div> |
| IKE                    | 26.096               | 7                 | 17         | <div style="width: 17%;"></div> | IKE                    | 9.341                | 8                 | 17         | <div style="width: 17%;"></div> |
| HTTPS                  | 15.154               | 4                 | 10         | <div style="width: 10%;"></div> | HTTPS                  | 5.634                | 5                 | 10         | <div style="width: 10%;"></div> |
| IMAP-SSL               | 5.043                | 2                 | 1          | <div style="width: 1%;"></div>  | ssdp                   | 5.268                | 1                 | 2          | <div style="width: 2%;"></div>  |
| Skype                  | 3.448                | 4                 | 36         | <div style="width: 36%;"></div> | SSH                    | 5.072                | 4                 | 2          | <div style="width: 2%;"></div>  |
| SSH                    | 2.672                | 5                 | 2          | <div style="width: 2%;"></div>  | SMTP                   | 3.791                | 2                 | 2          | <div style="width: 2%;"></div>  |
| <b>Anonymous Proxy</b> | <b>2.184</b>         | <b>2</b>          | <b>7</b>   | <div style="width: 7%;"></div>  | <b>Anonymous Proxy</b> | <b>2.745</b>         | <b>2</b>          | <b>7</b>   | <div style="width: 7%;"></div>  |
| SMTP                   | 1.806                | 4                 | 2          | <div style="width: 2%;"></div>  | Skype                  | 2.594                | 4                 | 36         | <div style="width: 36%;"></div> |
| ICMP                   | 0.530                | 1                 | 4          | <div style="width: 4%;"></div>  | IMAP-SSL               | 1.166                | 2                 | 1          | <div style="width: 1%;"></div>  |
| BitTorrent             | 0.506                | 1                 | 3          | <div style="width: 3%;"></div>  | ExindaCom              | 0.620                | 1                 | 12         | <div style="width: 12%;"></div> |
| DNS                    | 0.130                | 0                 | 1          | <div style="width: 1%;"></div>  | ICMP                   | 0.376                | 0                 | 4          | <div style="width: 4%;"></div>  |
| ExindaCom              | 0.104                | 0                 | 12         | <div style="width: 12%;"></div> | BitTorrent             | 0.328                | 1                 | 3          | <div style="width: 3%;"></div>  |
|                        |                      |                   |            |                                 | DNS                    | 0.105                | 0                 | 1          | <div style="width: 1%;"></div>  |

Screenshot 191: The Anonymous Proxy application is shown on the real-time monitoring screen.

The Optimizer policy configuration form below shows how to create an Optimizer policy that will block anonymous proxy traffic.

### NOTE

By default, the Anonymous Proxy application is part of the Recreational application group. This means that any policy that references the **Recreational** application group will also be referencing the Anonymous Proxy application. If you want to block anonymous proxy traffic, the anonymous proxy discard policy must be above (earlier than) any policy that references the Recreational application group in the Optimizer policy tree.

**Add New VC Policy**

Policy Name:       Block Options:  Discard only the first packet of a connection

VC Policy Number:

Schedule:

Action:

Policy Enabled:

Filter Rules:

| VLAN                 | Host                 | Direction                                | Host                 | ToS/DSCP             | Application     |
|----------------------|----------------------|------------------------------------------|----------------------|----------------------|-----------------|
| <input type="text"/> | <input type="text"/> | <input type="text" value="&lt; - &gt;"/> | <input type="text"/> | <input type="text"/> | Anonymous Proxy |
| <input type="text"/> | <input type="text"/> | <input type="text" value="&lt; - &gt;"/> | <input type="text"/> | <input type="text"/> |                 |
| <input type="text"/> | <input type="text"/> | <input type="text" value="&lt; - &gt;"/> | <input type="text"/> | <input type="text"/> |                 |
| <input type="text"/> | <input type="text"/> | <input type="text" value="&lt; - &gt;"/> | <input type="text"/> | <input type="text"/> |                 |

Screenshot 192: Blocking Anonymous Proxies using the Optimizer.

## 0.15 Managing Exinda Appliances with EMC

The appliances list shows the appliances that are being managed in the Exinda Management Center. The list shows inventory details, online/offline communication status, and configuration status. You can move the appliances to configuration groups or to the not deployed (unallocated) group. You can also import the network object configuration and the policy configuration from an appliance into the library. You can launch the Web UI of the appliance by clicking on the IP address. Note that it is a simple launch of the UI. If the appliance is subject to NAT, then the Web UI will not be available.

The Status Column indicates the online or offline status and the Last Communication column shows the date and time of last communication. Note that the system shows an appliance as offline if it has not called in within 15 minutes, which covers three scheduled call-in periods.

The **Config Status** column shows the state of the configuration and the date and time of the last configuration status change:

- » **Needs Sending** – configuration applicable to the appliance has changed, however, you have not clicked the Send Configuration icon
- » **Pending** – configuration applicable to the appliance has changed and you have clicked the Send Configuration icon, however, the appliance has not yet called in to receive the configuration
- » **Delivered** – configuration has been delivered to the appliance.

### Appliances

Move appliances into groups to manage and monitor appliances similarly.

| <input type="checkbox"/> | Status   | Host ID ▲    | Hostname              | IP Address  | Firmware   | Model   | Group       | Last Communication  | Config Status   | Executed            |
|--------------------------|----------|--------------|-----------------------|-------------|------------|---------|-------------|---------------------|-----------------|---------------------|
| <input type="checkbox"/> | ● Online | 392cc3ab6e82 | exinda-ab6e82         | 10.1.2.139  | 7.0.2.3073 | Virtual | Branches    | Sun Dec 14 17:57:17 | ⬇ Needs Sending | Sun Dec 14 17:57:32 |
| <input type="checkbox"/> | ● Online | 3f65a43ca2e8 | exinda-paris-bastille | 172.24.32.3 | 7.0.1.2456 | 4061    | Data Center | Sun Dec 14 17:59:57 | ● Pending       | Sun Dec 14 18:01:12 |
| <input type="checkbox"/> | ● Online | 533b847093b1 | exinda-nice           | 172.49.4.2  | 7.0.1.2456 | 4061    | Branches    | Sun Dec 14 18:00:03 | ● Delivered     | Sun Dec 14 18:00:03 |
| <input type="checkbox"/> | ● Online | 5aabb5ab24b2 | exinda-le-havre       | 10.25.32.9  | 6.4.5.3175 | 2061    | Branches    | Sun Dec 14 18:00:00 | ● Delivered     | Sun Dec 14 18:00:00 |
| <input type="checkbox"/> | ● Online | 681a7582cefb | exinda-dijon          | 172.69.4.5  | 6.4.5.3175 | 6060    | Branches    | Sun Dec 14 18:00:04 | ● Delivered     | Sun Dec 14 18:00:03 |
| <input type="checkbox"/> | ● Online | 821e21b7e1d5 | exinda-marseille      | 172.75.7.5  | 6.4.5.3175 | 4061    | Branches    | Sun Dec 14 18:00:03 | ● Delivered     | Sun Dec 14 18:00:03 |
| <input type="checkbox"/> | ● Online | 8c2c7e0ade41 | exinda-toulouse       | 10.32.3.1   | 7.0.1.2456 | 4061    | Branches    | Sun Dec 14 18:00:03 | ● Delivered     | Sun Dec 14 18:00:02 |

### 0.15.1 Viewing appliances in the tenancy

There are several locations in the Exinda Management Center where you can see the list of appliances in your tenancy, each providing a different scope of which appliances are shown.

- » Selecting **Not Deployed** shows only the appliances that have not been configured. Configuration cannot be applied to the appliances here. However, configuration can be imported into the library from appliances in the Unallocated Appliances list.
- » Selecting **Configured Appliances > Appliances** shows the appliances that can be configured by the Exinda Management Center. In order to apply configuration to an appliance, it must first be moved to the tenant's Configured Appliances group. Appliance groups can be added to the Configured Appliances group so that the appliances can be organized in a way that makes sense to you. Perhaps you want to organize by geography or by function (Data Center vs. Branches, or by circuit size). Groups can be nested. Configuration is applied by appliance groups, so all appliances in a group receive the same configuration.
- » Selecting a nested group within **Configured Appliances > Appliances**, will show the subset of configured appliances that have been moved to that group (and any nested groups) in the configured appliance groups hierarchy.

» Selecting your **Overview > Appliances** will show all the appliances in your tenancy including the unallocated appliances.

## 0.15.2 Moving appliances within the tenancy

### If there is more than one tenant in the system

When appliances first appear in the system within an on-premises EMC deployment with more than one tenant, they are added to the Appliance Pool group. This location is intended for the host of a multi-tenant deployment of the Exinda Management Center. Since the appliance does not know to which tenant it belongs, it appears in the Appliance Pool. The host admin user then moves the appliance to the appropriate tenant.

#### BEST PRACTICE

It is recommended that you initially move appliances to the Not Deployed (Unallocated) Appliance group to let the admin user manage the appliances in the tenant.

### If there is only one tenant in the system

When appliances first appear in the system, they appear in the **Not Deployed** appliances group. From there they can be moved to the **Configured Appliances** group.

#### IMPORTANT

When appliances are moved out of the **Configured Appliances** group to the **Unallocated Appliances** group or the **Appliance Pool**, the configuration that was applied via the edit forms will be automatically removed from the appliances upon the next call into the Exinda Management Center.

### To move an appliance

1. Select the node in the Tenant tree that contains the appliance you need to move. For instance, it could be in the:
  - **Appliance Pool** groups (on-premises EMC deployments only)
  - **Not Deployed Appliances** in the tenant
  - **Configured Appliances** group (or one of its nested groups) in the tenant
2. The system shows the list of appliances.

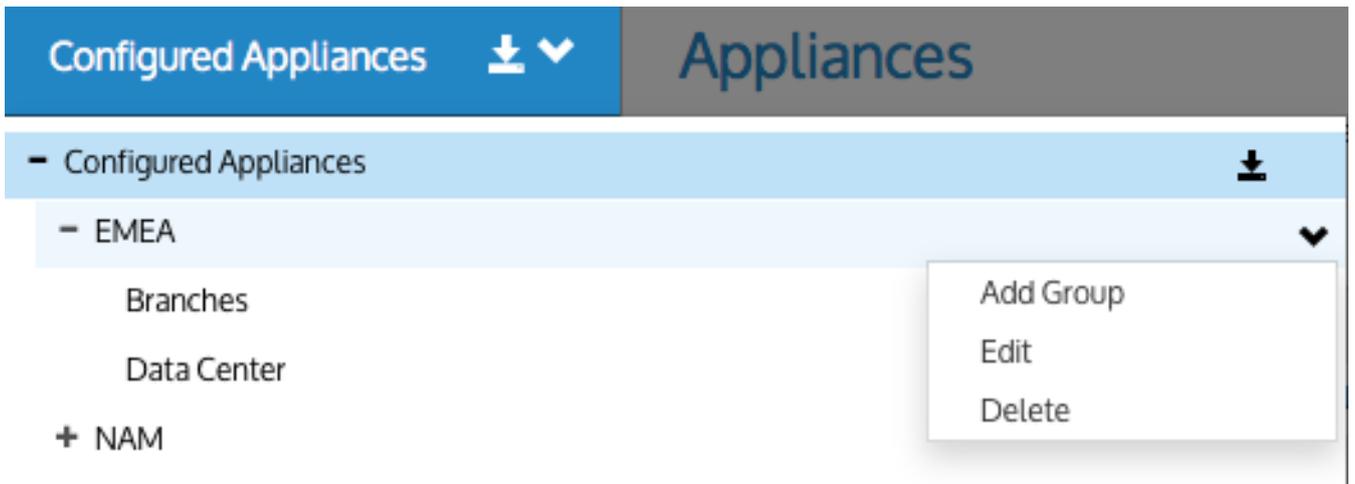
#### TIP

To look at the appliances in the **Configured Appliances** group (or one of its nested groups), you need to select the **Appliances** menu item.

3. Select the checkbox next to the desired appliance(s).
4. Click **Move Appliances**.
5. Select the destination for the appliance(s).
6. Click the **Move Appliances** button.

### To create a new appliance group

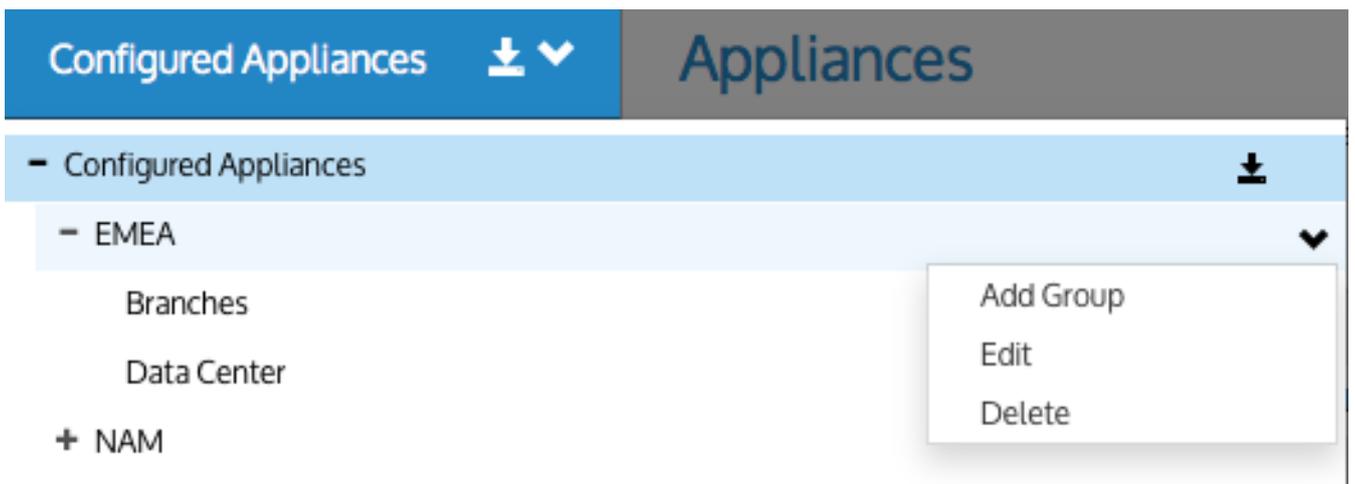
1. Click the appliance group header menu and click the menu icon of the group to which you want to add a group.



2. Click the menu icon of the group where you want to add a group.
3. Select the **Add Group** menu item. A new group is added, with the Edit mode enabled.
4. Type the name of the group and press **Enter** to commit the name.

#### To edit an appliance group name

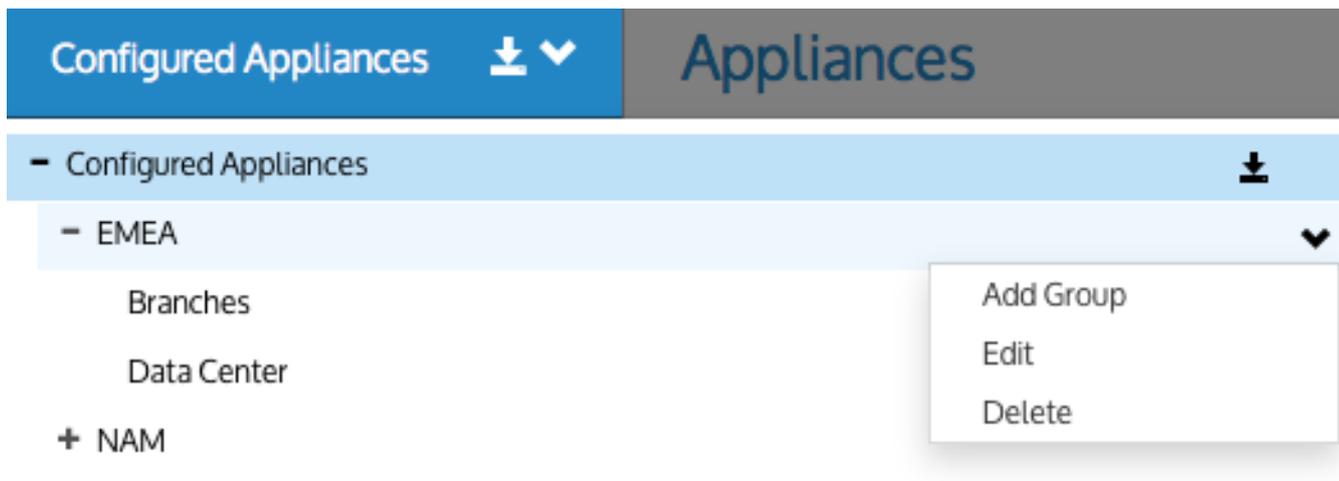
1. Click the appliance group header menu and click the menu icon of the group that you want to rename.



2. Select the **Edit** menu item.
3. Type the name of the group to create the group.

#### To delete an appliance group

1. Click the appliance group header menu and click the menu icon of the group that you want to delete.



2. Select the **Delete** menu item.

#### NOTE

- » If there are appliances in the group or its sub-groups, then the group cannot be deleted.
- » If there is configuration on the group, but no appliances, it can be deleted.

3. Confirm that you want delete the group.

### 0.15.3 Configuration Library

The Configuration Library allows to you create items and save them for reuse in different areas, much like a template. For instance, create a policy set once and then use it in multiple virtual circuits or in multiple Optimizer Policy Trees applied to different appliance groups.

The library items will appear in drop-down lists when configuring other items. For example, when creating a policy, you can select an application from a list; when creating a Optimizer Policy Tree, you can select a circuit from a list, or a virtual circuit from a list, or a policy set from a list, and so on.

Library items are categorized as follows:

- » **Circuits** – Identifies the physical connections to the WAN or Internet by defining the inbound and outbound bandwidth and the named circuit type. The circuits within a tenant must have unique names.
- » **Circuit Types** – An abstract concept used to identify the purpose of the circuit and appliance bridges and to create a logical binding between the circuits and the appliance bridges.
- » **Virtual Circuits** – Logically divide or partition a circuit to define what traffic will be processed in this partition (and when), and how much bandwidth it is allowed.
- » **Dynamic Virtual Circuits** – Dynamic virtual circuits provide a means to configure fair sharing among the hosts, or to configure a limit to the number of hosts so that those hosts get preferential treatment.
- » **Policy Sets** – Ordered list of policies that can be applied to one or more virtual circuits in one or more appliance groups.
- » **Policies** – Define the actions to perform on specific targeted traffic.
- » **Network Objects** – Represent hosts on a network and can include subnets, single hosts, or groups of both. Once on the appliance, network objects are used to determine if host and user traffic data are internal or external to the LAN behind your appliance.

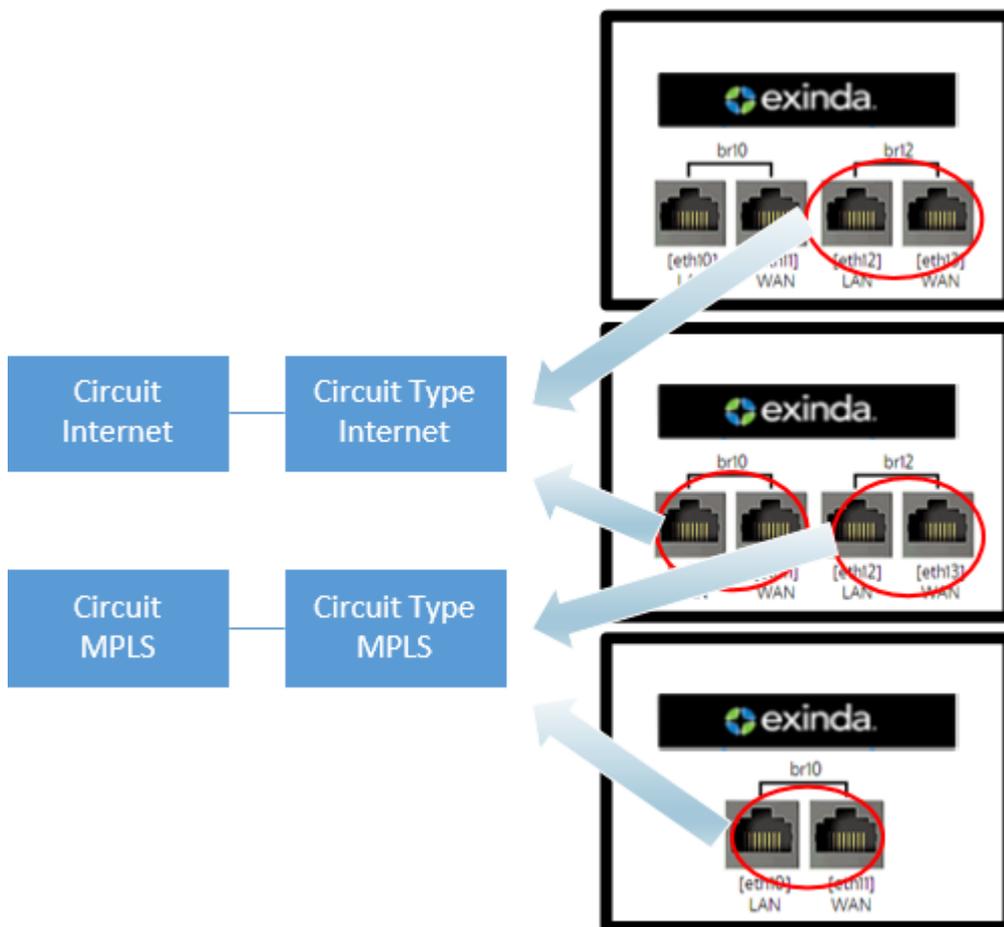
- » **Applications** – Classify traffic by layer 7 signatures OR by a combination of network objects, ports, protocols, and DSCP markings. You can then filter traffic generated by the applications to determine which policy to apply.
- » **Application Group** – Preset and custom groups of applications to monitor and subsequently classify traffic and determine which policy to apply to traffic.
- » **Schedules** – Define a specific timeframe of the week. When used in policies or virtual circuits, the schedule will affect traffic only within the identified timeframe.
- » **Application Performance Scores** – Identify the applications operating on your network whose performance you need to monitor.
- » **Service Level Agreements** – The Service Level Agreement (SLA) objects are used to monitor the availability of a particular IP site.
- » **VLANs** – Virtual LAN (VLAN) Objects are used to logically separate hosts (or groups of hosts) on a functional basis rather than on a physical basis.

## Circuit types in the EMC

Circuit Types are used to identify the purpose of the circuit and the appliance bridges. When the purpose of the circuit and an appliance bridge align, then the circuit is bound to that bridge in the configuration that is sent to the appliances. Note that Circuit Types do not exist on appliances. Circuit Types abstract the binding between the circuits and the appliance bridges. Therefore, the appliances can be treated similarly even when they do not have the same number of bridges and when they are not connected in the same way.

Circuit Types are defined in the Configuration Library and are used by circuits and appliance bridge-to-circuit type mappings.

The following example will further clarify the concept of circuit types:



Screenshot 193: Circuit types example

Consider the following in the example above:

- » Bridge **br12** from first appliance and bridge **br10** on second appliance are both bound to circuit type "Internet". Since, this circuit type "Internet" is tied to the circuit "Internet", all the policies within the "Internet" circuit are applied to bridge br12 on the first appliance and br10 on the second appliance.
- » Bridge **br12** on the second appliance and bridge **br10** on third appliance are bound to circuit type "MPLS". Since, this circuit type "MPLS" is tied to the circuit "MPLS", all the policies within the "MPLS" circuit are applied to bridge br12 on the second appliance and br10 on the third appliance.

The **Circuit Type** library items can be found in **Library > Circuit Types**.

To assign the Circuit Type to a circuit, see [Circuits](#).

To assign Circuit Types to appliance bridges, see [Bridge/Circuit Type Mapping](#).

## Circuits

Circuits define physical connections to the WAN or the Internet. A circuit defines the inbound and outbound bandwidth and the named circuit type.

On an Exinda Appliance, a circuit specifies the named bridge (or bridges) to which it is bound. In the Exinda Management Center, the binding to bridges is through a designated Circuit Type. Circuit Types represent the intended use of a circuit. So you can configure a circuit for multiple appliances without requiring the bridges on the appliances to have the same name. This is favorable where the number of bridges or names of bridges or the cabling of the bridges is not consistent across the appliances.

If multiple bridges on an appliance are mapped to the same Circuit Type, then all those bridges will be bound to the single circuit in the Policy Tree that is configured with that Circuit Type. To learn how Circuit Types are used to determine which circuits are sent to the appliances, read [Circuit Types](#).

Circuits can be created in the Configuration Library directly and then later assigned to an Optimizer Policy Tree, Circuits can also be created in an Optimizer Policy Tree, from where the configuration is saved to the Configuration Library so that it can be used elsewhere. The circuits within a tenant must have unique names. When a circuit is modified, all uses of it are modified.

Circuits are part of the Optimizer Policy Tree. To learn how circuits, Virtual Circuits, policy sets, and policy rules work together, see [Policy Tree](#).

## Where do I find circuits?

Circuit library items can be found in **Library > Circuits**. Also, the circuits that are applied to appliances can also be found in the policy trees for each appliance group. Go to the **(desired appliance group) > Optimizer Policy Tree**.

## To create a circuit in the Configuration Library

1. Go to **Library > Circuits**.
2. Click **Create new circuit in the library ...**
3. In the **Name** section, key-in the name of the circuit. The name must be unique in the tenant.
4. In the **Bandwidth** section, key-in the **Inbound Bandwidth** and the **Outbound Bandwidth**. The bandwidths can be specified in kbps, Mbps, or Gbps.
5. In the **Bind to Circuit Type** section, select the **Circuit Type**.

### NOTE

The circuit type represents the purpose of the circuit; it is a user created object. If the desired circuit type does not exist, you can click **Create new circuit type in the library** to create it. For more information, refer to [Circuit types in the EMC](#) (page 423).

6. Click **Create**. The circuit appears in the library list.

## To create a new circuit directly in the Policy Tree

1. Go to the **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Create new circuit**.
3. In the **Name** section, key-in the name of the circuit. The name must be unique in the tenant.
4. In the **Bandwidth** section, key-in the **Inbound Bandwidth** and the **Outbound Bandwidth**. The bandwidths can be specified in kbps, Mbps, or Gbps.
5. In the **Bind to Circuit Type** section, select the **Circuit Type**.

### NOTE

The circuit type represents the purpose of the circuit; it is a user created object. If the desired circuit type does not exist, you can click **Create new circuit type in the library** to create it. For more information, refer to [Circuit types in the EMC](#) (page 423).

6. Click **Save**. The circuit is added to the Policy Tree and is also saved to the configuration library.

## To add a circuit to a Policy Tree

1. Go to the **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add circuit from library**.
3. Select the desired circuit from the drop-down list.

## Virtual Circuits

Virtual Circuits are created within Circuits and are used to logically divide or partition the circuit. The virtual circuit defines what traffic is processed in this partition, how much bandwidth it is allowed, and whether to enforce fair sharing among the network hosts. Traffic is evaluated against the definition of the virtual circuit. Traffic that does not fall within the virtual circuit is evaluated by the next virtual circuit and so on.

You have the option of creating a virtual circuit within the Configuration Library first, and then later assigning it to a circuit in the Optimizer Policy Tree, or you can create the virtual circuit directly within the Optimizer Policy Tree, which also saves it as a Library item. The virtual circuits within a tenant must have unique names. When a virtual circuit is modified, all instances of its use are modified.

To learn how circuits, virtual circuits, policy sets, and policy rules work together, see [Policy Tree](#).

## Where do I find Virtual Circuits?

Virtual Circuit library items can be found in **Library > Virtual Circuits**. Also, the virtual circuits that will be sent to appliances are found in the policy trees for each appliance group. Go to the **(desired appliance group) > Optimizer Policy Tree**.

## To create a Virtual Circuit in the Configuration Library

1. Go to **Library > Virtual Circuits**.
2. Click **Create new virtual circuit**.
3. In the **Name** section, key-in a name for the virtual circuit. The name must be unique within the tenant.

### NOTE

If you want, you can leave the EMC to define a name for you, It does this based on the configuration of the virtual circuit.

4. In the **Filter** section, select the combination of filters to apply to the virtual circuit. Optionally, type a value to limit the number of connections at one time on this virtual circuit.

### NOTE

The virtual circuit can partition the circuit by filtering the traffic based on these filters. You can apply any combination of these filters. Defined network object library items appear in the [Network Object](#) list, and you can also choose filters from the pre-defined application groups.

5. In the **Bandwidth** section, type the desired bandwidth for this virtual circuit.
6. Also specify how to share bandwidth with other virtual circuits when there is insufficient bandwidth due to over subscription..

**NOTE**

The bandwidth can be specified in %, Kbps, Mbps or Gbps.

7. In the **Dynamic Virtual Circuit** section, set the options that provide the control you need.

**NOTE**

See For more information, refer to [Dynamic Virtual Circuits](#) (page 428). for more information about configuring dynamic virtual circuits.

8. In the **Schedule** section, set the time values for when the virtual circuit will be enforced.

**NOTE**

Options in the list are determined by the library [Schedules](#) category

9. Click the **Create** button. The virtual circuit is added to the Virtual Circuits Library category.

### To create a new Virtual Circuit directly in the Optimizer Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Create new virtual circuit**.
3. In the **Name** section, key-in a name for the virtual circuit. The name must be unique within the tenant.

**NOTE**

If you want, you can leave the EMC to define a name for you, It does this based on the configuration of the virtual circuit.

4. In the **Filter** section, select the combination of filters to apply to the virtual circuit. Optionally, type a value to limit the number of connections at one time on this virtual circuit.

**NOTE**

The virtual circuit can partition the circuit by filtering the traffic based on these filters. You can apply any combination of these filters. Defined network object library items appear in the [Network Object](#) list, and you can also choose filters from the pre-defined application groups.

5. In the **Bandwidth** section, type the desired bandwidth for this virtual circuit.

6. Also specify how to share bandwidth with other virtual circuits when there is insufficient bandwidth due to over subscription..

**NOTE**

The bandwidth can be specified in %, Kbps, Mbps or Gbps.

7. In the **Dynamic Virtual Circuit** section, set the options that provide the control you need.

#### NOTE

See For more information, refer to [Dynamic Virtual Circuits](#) (page 428). for more information about configuring dynamic virtual circuits.

8. In the **Schedule** section, set the time values for when the virtual circuit will be enforced.

#### NOTE

Options in the list are determined by the library [Schedules](#) category

9. Click the **Create** button. The virtual circuit is added to the Circuit within the Optimizer Policy Tree and is also saved to the Configuration Library.

### To add a Virtual Circuit from the Library to a Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add Virtual Circuit from library**.
3. Select the desired virtual circuit from the drop-down list.

### Dynamic Virtual Circuits

You can use Dynamic Virtual Circuits to enforce fair sharing of bandwidth among the hosts, or to limit the number of hosts on the circuit to ensure that those hosts get preferential treatment.

- » For fair sharing, you must specify how you would like the bandwidth in the virtual circuit to be shared among the hosts. You can fix the per host bandwidth and have the system calculate the number of allowed hosts. Note that if there are less than the allowed hosts, each active host can burst to gain more bandwidth (if you have configured the virtual circuit to allow bursting).
- » For limiting the number of hosts, you can have the system calculate the amount of bandwidth that is then allowed to each host. You can specify an automatic calculation of the per host bandwidth and the number of allowed hosts. The system then divides the virtual circuit bandwidth by the number of active hosts.

The options available in the Exinda Management Center appear in the following screenshot:

## Dynamic Virtual Circuit:

Allocate minimum and maximum bandwidths to each unique host that matches the virtual circuit.

Enable dynamic virtual circuit

Allocate a bandwidth for each host on the network

Manually define bandwidth usage to be allocated for each host

 % 

Automatically adjust and share bandwidth for each host

Define maximum bandwidth usage for each host

Manually define maximum bandwidth burst for each host

 % 

Do not allow maximum bandwidth burst for each host

Location of hosts to allocate bandwidth

Internal - LAN side

External - WAN side

Define number of maximum hosts to fall under this dynamic virtual circuit

Manually define number of hosts to be affected by this dynamic virtual circuit

Automatically determine number of hosts to be affected by this dynamic virtual circuit

Screenshot 194: Dynamic virtual circuit options

Consider the following:

- » While allocating bandwidth usage to each host on the network.
  - If you manually define the bandwidth usage for each host, you limit the number of hosts that can be accommodated on the dynamic virtual circuit. You can do this by either defining an actual bandwidth or by defining a percentage of the available bandwidth. By default, though, the EMC can set no less than 10 kbps for any one host, so a hard limit to the maximum number of hosts is the total available bandwidth divided by 10 kbps. This setting allows you set higher bandwidth quotas for a limited number of hosts.

- By allowing the EMC to automatically adjust and share bandwidth, you are letting as many hosts as can be to accommodate within the available bandwidth on the dynamic virtual circuit. The minimum bandwidth that the EMC can provide is 10 kbps, the hard limit to the number of hosts is the total bandwidth divided by 10 kbps.
- » While defining a maximum bandwidth usage for each host, when spare capacity exists on the dynamic virtual circuit because few hosts are active, you can allow the active hosts have greater bandwidth (to burst). You can set the burst rate limit as an actual bandwidth measure (kbps, Mbps, etc.) or specify a percentage of the available bandwidth. You can also disallow bursting.
- » While specifying the location of the hosts as internal or external, remember that this setting allows you specify whether the hosts on the dynamic virtual circuit located within the LAN or outside the LAN.
- » While defining the maximum number of hosts, you can either manually define the maximum number of hosts that can be accommodated on the dynamic virtual circuit, or you can let the EMC control how many hosts are allowed on the circuit.

## Policies

Policies define what actions to perform on specific traffic. The policies can specify whether to optimize the traffic (by bandwidth shaping, acceleration, or packet marking), block the traffic (by discarding the packets), or monitor the traffic (by ignoring the packets). The traffic that the policy affects can be filtered by:

- » Application or application group
- » Hosts or subnets
- » Hosts or subnets that are communicating with other specific hosts or subnets
- » VLAN
- » ToS/DSCP markings
- » Time of day

Any combination of these filters can be applied. For example, the policy could target SAP traffic between a particular branch and headquarters that has particular ToS markings on a particular VLAN during work hours. Furthermore, you can add more than one filter. That is, the policy could target a particular branch site for Netflix and the same branch site for Silverlight.

### VERSION INFO

Exinda Management Center 1.5.0 does not support policies for HTTP Redirect or HTTP Response.

When you create policies, they are added to the Policies Library. If you amend a policy definition, any changes made to it affect all Virtual Circuits that use that policy. To learn how circuits, virtual circuits, policy sets, and policy rules work together, see [Policy Tree](#).

## Where do I find policy rules?

Policy library items can be found in **Library > Policies**.

Also, the policies linked to appliances can also be found in the policy trees for each appliance group. Go to the desired appliance group's **Optimizer Policy Tree**.

## To create a policy in the Configuration Library

1. Got to **Library > Policies**
2. Click **Create new policy in the library**.

3. In the **Name** section, type a name for the policy. The name must be unique within the tenant.

4. In the **Action** section, specify what type of action the rule should take. Select *one* of the following:

**Optimize**– Selecting optimize causes a new action to appear in the UI where you can specify whether you want to apply bandwidth shaping, prioritization, acceleration, or packet marking.

- **Discard**– Select discard to specify that you want to block a particular type of application by discarding the packets.
- **Ignore**– Select ignore to specify that you want to allow packets to pass through without manipulation, that is, traffic monitoring only .

4. In the **Filter** section, specify the type of traffic to which you want to apply the policy. Set any of the following traffic attributes.

- **Application** – Select traffic based on a predefined application or application group from the list. Custom applications that you have created in the library will appear in this drop-down list.
- **Source/Direction/Destination** – Select traffic based on one end of a conversation belonging to a predefined network object or select traffic based on one way or two way conversations between two predefined network objects. For the source, select a network object that filters for the initiation of a conversation. For the destination, select a network object that filters for the destination of the conversation. If hosts are not specified, ALL network objects are assumed. Traffic direction is relative to the Exinda appliance.
- **ToS/DSCP** – Select traffic based on particular ToS/DSCP markings in the IP header. When multiple settings are specified, the intersection of those attributes determines whether the traffic is matched.

5. Click the **Create** button. The policy set will be added to the library list.

#### NOTES

- » While creating the policy set, if you create a new policy but cancel creating the policy set, the new policy is retained in the library.
- » Some applications may not be compatible with all appliance versions. When an application is not compatible, it does not appear in the drop-down list.

### To create a new policy set directly in the Optimizer Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.

2. Click **Create new policy set**.

3. In the **Name** section, type a name for the policy. The name must be unique within the tenant.

4. In the **Action** section, specify what type of action the rule should take. Select *one* of the following:

**Optimize**– Selecting optimize causes a new action to appear in the UI where you can specify whether you want to apply bandwidth shaping, prioritization, acceleration, or packet marking.

- **Discard**– Select discard to specify that you want to block a particular type of application by discarding the packets.
- **Ignore**– Select ignore to specify that you want to allow packets to pass through without manipulation, that is, traffic monitoring only .

4. In the **Filter** section, specify the type of traffic to which you want to apply the policy. Set any of the following traffic attributes.

- **Application** – Select traffic based on a predefined application or application group from the list. Custom applications that you have created in the library will appear in this drop-down list.
- **Source/Direction/Destination** – Select traffic based on one end of a conversation belonging to a predefined network object or select traffic based on one way or two way conversations between two predefined network objects. For the source, select a network object that filters for the initiation of a conversation. For the destination, select a network object that filters for the destination of the conversation. If hosts are not specified, ALL network objects are assumed. Traffic direction is relative to the Exinda appliance.
- **ToS/DSCP** – Select traffic based on particular ToS/DSCP markings in the IP header. When multiple settings are specified, the intersection of those attributes determines whether the traffic is matched.

5. Click the **Save** button. The policy set is added to the Policy Tree and is also saved to the configuration library.

### To add a policy set from the library to a Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add policy set from library**.
3. Select the desired policy set from the drop-down list.

### Creating policy sets in the EMC

Policy sets are an ordered list of policies that can be applied to one or more virtual circuits in one or more appliance groups.

Policy sets can be created in the Policy Set Library directly and then later assigned to an Optimizer Policy Tree, or they can be created in an Optimizer Policy Tree. When saved in an Optimizer Policy Tree, they are also saved to the Configuration Library for use elsewhere. The policy sets within a tenant must have unique names. When a policy set is modified, all uses of it are modified.

There are a few default policy sets that you can use or modify. They correspond to the defined configuration when the wizard is run on the appliance. The different policy sets are due to varying answers to the wizard questions.

To learn how Circuits, Virtual Circuits, Policy Sets, and Policy Rules work together, see [Policy Tree](#).

### Where do I find policy sets?

Policy set library items can be found in **Library > Policy Sets**. Also, the policy sets linked to appliances can also be found in the policy trees for each appliance group. Go to the desired appliance group in the **Optimizer Policy Tree**.

### To create a policy set in the Library

1. Go to **Library > Policy Sets**.
2. Click **Create new policy set**
3. In the **Name** section, key-in a name for the policy. The name must be unique within the tenant.
4. In the **Policies** section, add policies to the list.
  - You can select a policy from the library by clicking **Add policy from library**.
  - You can create a new policy by clicking **Create new policy**.

Policies created in the policy set are automatically added to the policy library. For more information, refer to [Policies](#) (page 430).

4. To reorder the policies, drag and drop the policy rule to the desired location.

5. Click **Create**. The policy set is added to the library list.

#### NOTE

While creating the policy set, if you create a new policy but cancel creating the policy set, the new policy is retained in the library.

### To create a new policy set directly in the Optimizer Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Create new policy set...**
3. Click **Create new policy set**
4. In the **Name** section, key-in a name for the policy. The name must be unique within the tenant.
5. In the **Policies** section, add policies to the list.
  - You can select a policy from the library by clicking **Add policy from library**.
  - You can create a new policy by clicking **Create new policy**.

Policies created in the policy set are automatically added to the policy library. For more information, refer to [Policies](#) (page 430).

4. To reorder the policies, drag and drop the policy rule to the desired location.
5. Click **Save**. The policy set is added to the Optimizer Policy tree for the current appliance group and is also saved to the Configuration Library.

### To add a policy set from the library to a Policy Tree

1. Go to **(desired appliance group) > Optimizer Policy Tree**.
2. Click **Add policy set from library**.
3. Select the desired policy set from the drop-down list.

## Configuring Network Objects

Learn how to defining and apply the network objects for managing aspects of your network.

### Creating network objects in the EMC

A network object can be created in the Library for later use in other components and appliance groups, or it can be created directly in the Optimizer Policy Tree, which also saves it to the Library. You can also import network objects into the Library from an appliance.

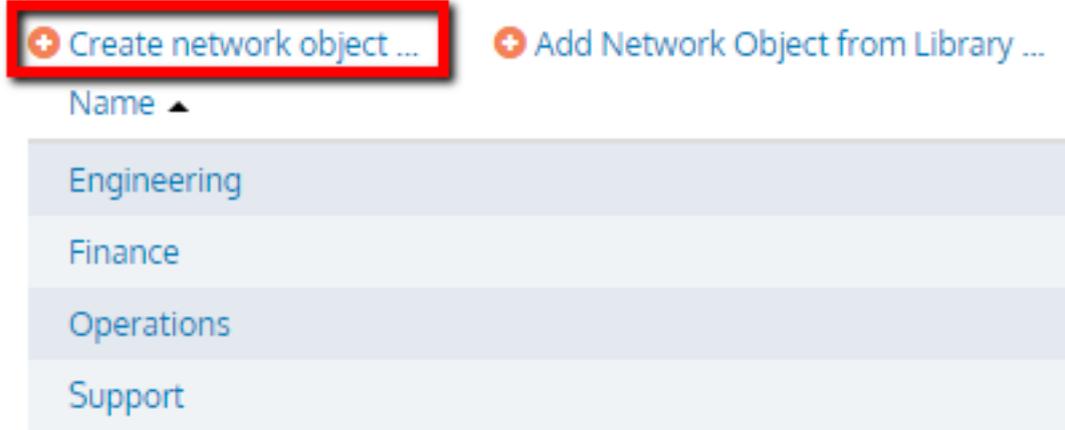
### To create and apply a network object

To create and apply a network object:

1. Go to **Library > Network Objects**.
2. Click **Create new network object in the library....**

# Network Objects

Define network objects to represent subsets of your network, which can



3. Key-in a **Name** for the object.

## Network Object

Define network objects to represent subsets of your network, which can include multiple subnets :

The screenshot shows the 'Network Object' configuration form. It has a 'Name' section with a dropdown arrow and an empty text input field. Below this are two expandable sections: 'Reporting: Include in subnet reporting in the selected appliance group' and 'Subnets'. At the bottom, there are two buttons: 'Create' (green) and 'Cancel' (white).

4. In the **Subnets** section, define subnets to include in this network object by **IP Network Address** and **Mask Length** or by **Fully Qualified Domain Name (FQDN)**.

## Network Object

Define network objects to represent subsets of your network, which can include multiple s

> Name

> Reporting: Include in subnet reporting in the selected appliance group

▼ Subnets

Define which subnets to include in this network object.

IP Network Address / Mask Length

/  ✕

+ Add Another Subnet

Create Cancel

5. Click **Create**. Repeat from step 1, if you want to create another network object.
6. Go to **Configured Appliances > Network Objects**.
7. Click **Add network object from Library...**

## Network Objects

Define network objects to represent subsets of your network, which can include multiple

+ Create network object ...    + Add network object from library ...

Name ▲

|                     |
|---------------------|
| Network Object      |
| Test Network Object |

**Click here**

8. Select one or more network objects to add and click **Add Network Object to 'Configured Appliances'**.
9. To define whether to include monitor information of network objects in subnet reporting, edit the required network objects and check **Include in subnet reporting** option, within the **Reporting** section.

# Network Objects

Define network objects to represent subsets of your network, which can include multiple subnets and multiple hosts. The network objects can then be used to monitor traffic or to configure traffic policy.

[+ Create network object ...](#)

[+ Add Network Object from Library ...](#)

| Name ▲      | IP Network Address | Monitoring               |   |
|-------------|--------------------|--------------------------|---|
| Engineering | 10.10.1.0/24       | ✓                        | 🔒 |
| Finance     | 10.11.0.0/16       | ✓                        | 🔒 |
| Marketing   | 10.50.1.0/24       | <input type="checkbox"/> | ✕ |
| Operations  | 10.30.0.0/16       | ✓                        | 🔒 |
| Support     | 10.10.10.5/32      | ✓                        | 🔒 |

## To use a network object in a policy rule definition

You can apply a network object filter to a policy to filter by subnet (communicating with other specific hosts or subnets). Use the instructions above to create the Network Object in the library, and then refer to [Policy Library](#) for further instructions. The Network Objects saved as Library items appear in the **Filter** section, under the **Source** list.

## To use a network object in a virtual circuit definition

When defining virtual circuits to partition a circuit, you can apply Network Object library items to filter the traffic by subnet. Use the instructions above to create the Network Object in the library, and then refer to [Virtual Circuits](#) for further instructions.

## To use a network object in an Application

When defining applications to classify traffic, you can apply Network Object library items to classify traffic based on a combination of Network Object, TCP Port, UDP Port, DSCP, and Protocols. Use the instructions above to create the Network Object in the library, and then refer to [Applications](#) for further instructions.

## Configuring local network objects in the EMC

Local network objects define which part of the network is considered as the local area network relative to an appliance.

Local network objects take their definition from an IP network address and mask length to identify the range of IP addresses that exist in the LAN behind the appliance. You can define a local network object for each appliance, so that each appliance can differentiate between traffic that is external and internal to the LAN on which it operates. The local network objects cannot be deleted.

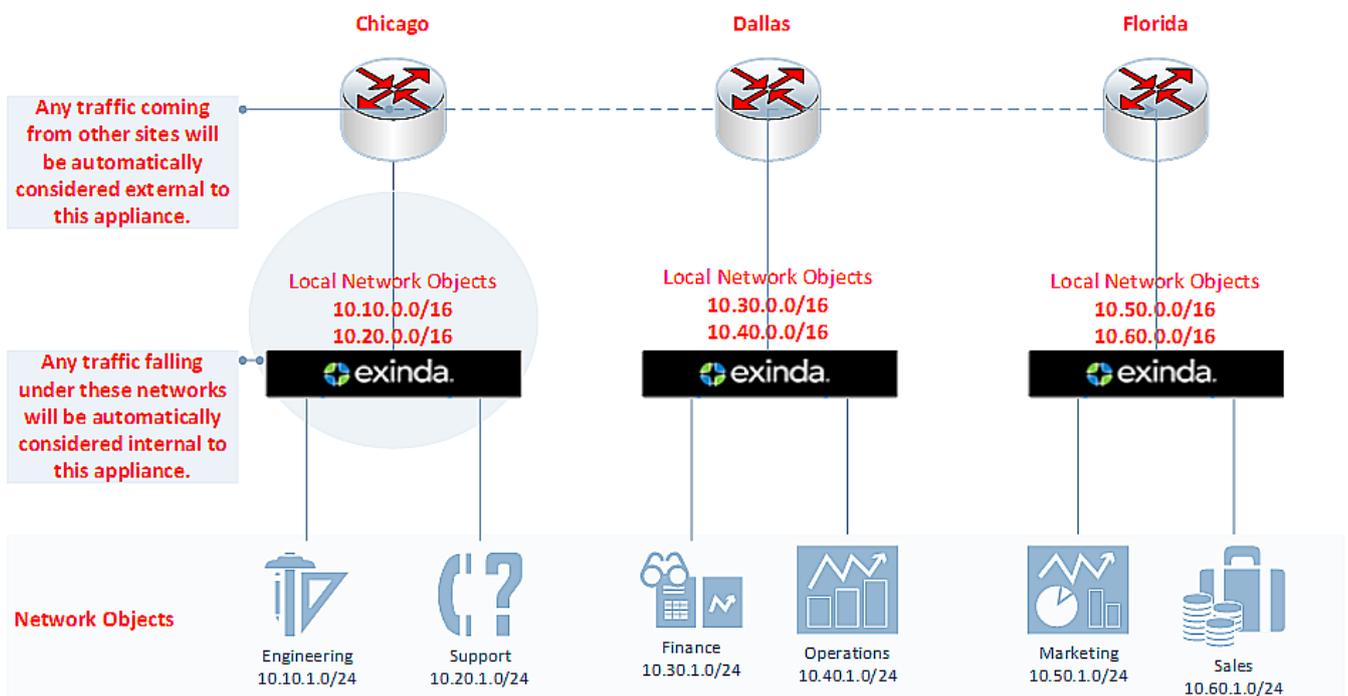
For reporting purposes, hosts and users are defined as internal or external by comparing the IP address with the network objects.

## EXAMPLE

Considering that there are three sites, Chicago, Boston, and Dallas where each site has an appliance, and each appliance has a local network object. Do the following:

- » Create network objects to represent the Chicago site, the Boston site, and the Dallas site.
- » On the Chicago appliance, when compared to the appliance's local network object, the Chicago network object gets set as internal, and the Boston and Dallas network objects are set as external.
- » On the Boston appliance, the Boston network object gets set as internal, and the Chicago and Dallas network objects are set as external.

In the EMC configuration, local network objects are appliance specific, so appliances cannot share these objects.



Screenshot 195: Diagram depicting separated local network objects for each appliance.

To configure the local network object:

1. Go to **Configured Appliances > Local Network Objects**.
2. Click **Local to Appliance (Host ID)** of the appliance to configure the local network object.

## Local Network Objects by Subnet

Local network objects represent the subnet local to each Exinda appliance in the system to-local traffic. The local network objects cannot be deleted.

Local to Appliance (Host ID)

|              |
|--------------|
| 65458c8ea68c |
| 6b06a98e3572 |



3. Key-in the **IP Network Address** and **Mask Length**.

## Local Network Objects by Subnet

Local network objects represent the subnet local to each Exinda appliance in the system to-local traffic. The local network objects cannot be deleted.

Location: Internal To 65458c8ea68c

Reporting: Do NOT include in subnet reporting

Subnets: 172.24.32.0/24

Define which subnets to include in this network object.

| IP Network Address                       | / | Mask Length                     |                                  |
|------------------------------------------|---|---------------------------------|----------------------------------|
| <input type="text" value="172.24.32.0"/> | / | <input type="text" value="24"/> | <input type="button" value="x"/> |



Screenshot 196: Defining local network object by subnet.

4. Click **Save**.

## NOTE

If the IP address and mask length you define are illegal, the EMC returns a warning and provides a suggestion for correcting the issue.

The screenshot shows the 'Local Network Objects by Subnet' configuration page. The left sidebar contains a navigation menu with items like 'Configured Appliances', 'Appliances', 'Optimizer Policy Tree', 'Bridge/Circuit Type Mapping', 'Applications', 'Application Groups', 'Network Objects', 'Local Network Objects', 'Application Performance Scores', 'Service Level Agreements', and 'Configuration via CLI'. The main content area has a breadcrumb trail: 'Tenant > Overview | Configured Appliances | Library | Not Deployed'. Below the breadcrumb, the page title is 'Local Network Objects by Subnet'. A descriptive text states: 'Local network objects represent the subnet local to each Exinda appliance in the system. The local network object to-local traffic. The local network objects cannot be deleted.' The configuration area includes fields for 'Location: Internal To 0024e83dcaed' and 'Reporting: Do NOT include in subnet reporting'. The 'Subnets' field is expanded to show '10.10.0.0/16' with a warning icon. A tooltip message reads: 'The IP network address and mask length combination is illegal. A suggestion has been made. Click save again to use the supplied suggestion.' Below this, there is a section titled 'Define which subnets to include in this network object.' with input fields for 'IP Network Address' (10.10.0.0) and 'Mask Length' (16). There is an 'Add Another Subnet' button and 'Save' and 'Cancel' buttons at the bottom.

Screenshot 197: Defining local network object by subnet.

## Importing network objects

If you have existing appliances that are already configured with network objects, in the Exinda Management Center you can import these network objects into the library (one at a time) from the Not Deployed list. When importing network objects, the importer indicates if the network object already exists in the library (or is included in another network object), or if it conflicts with another network object in the library. By importing your network objects, you can more quickly start building a library of objects to use. The system lets you know if the imported network object is already in the library or if the imported network object definition is a subset of a network object in the library. This allows you to quickly ensure consistency across multiple appliances.

To import network objects:

1. From the **Not Deployed** appliance list in your tenant, select the appliance from which you want to import the network objects.
2. Click **Import Configuration**.

# Import Configuration

Network objects from the configuration file can be imported into the library. The system will warn when a network object is not available for import or already has the same definition.

The system will warn when a network object is not available for import or already has the same definition.

## Step 1 - Network Objects

| <input type="checkbox"/>   | Name              | IP Network Address                     |
|----------------------------|-------------------|----------------------------------------|
| <input type="checkbox"/>   | Exinda Appliances | 1.2.3.5/32<br>1.2.3.6/32<br>1.2.3.4/32 |
| <input type="checkbox"/> ✓ | Toronto           | 10.21.0.0/24                           |
| <input type="checkbox"/>   | Data Center       | 10.10.10.0/24                          |

Screenshot 198: Importing configuration

3. Click **Import Network Objects**. A list of network objects, with their IP addresses, appears. If a network object has the same name as a network object in the library, the following indicate the status:

- ✓ – if the IP addresses are the same, the name has a green check mark before it.
- ⊆ – if all of the object IP addresses are contained within the one in the library, the name has a green subset symbol before it.
- ⚠ – if the imported network object has an IP address that is not in the network object with the same name in the library, the name has a warning icon before it. In this case, you have to resolve this conflict manually by modifying the network object in the Library.

4. Select the network objects that you would like to import.

5. Click **Add Selected Network Objects to the Library**. The network objects are imported into the library.

### NOTE

While importing network objects, neither the location (internal or external) of the network object, nor the reporting flag are imported.

## Configuring applications with the EMC

The Exinda Management Center provides a comprehensive set of built-in Applications for you to use. These cannot be edited, but you can define Custom Applications. You can add the custom application as part of an ignore policy at the bottom of the optimizer tree that is sent to the appliances.

## NOTE

if you use a custom application in the definition of a virtual circuit or policy for a given appliance group, then the custom application is automatically added to the appliance-group configuration.

Applications can be defined in the Library directly, and then later assigned to an Optimizer Policy Tree. Applications can also be created in an Optimizer Policy Tree, in which case they are also saved to the Configuration Library for use elsewhere. All applications within a tenant must have unique names. When an application is modified, all uses of it are also modified.

### Where do I find Applications?

Application library items can be found in **Library > Applications**. You can define custom applications for each appliance group. Go to the desired appliance group in the **Optimizer Policy Tree**.

### How do I view built-in Applications?

Built-in application library items can be found in **Library > Applications > Built-in**. You can view built-in applications, but not edit.

### To create a Custom Application in the Library

1. Go to **Library > Applications**.
2. Click **Create new application...**
3. In the **Name** section, key-in a name for the application. The name must be unique in the tenant.
4. In the **Definition** section, choose either the Layer 7 Signature, or select a combination of Network Object, Ports, DSCP and Protocols.
5. Click the **Create** button. The custom application will be added to the library list.

### To create a new Application directly in the Optimizer Policy Tree

1. Go to **Configured Appliances > (desired appliance group) > Applications**.
2. Click **Create new application...**
3. In the **Name** section, key-in a name for the application. The name must be unique in the tenant.
4. In the **Definition** section, choose either the Layer 7 Signature, or select a combination of Network Object, Ports, DSCP and Protocols.
5. Click the **Create** button. The custom Application is added to the Optimizer Policy tree for this current Appliance Group and is also saved to the Library.

### To add an Application from the library to an appliance group

1. Go to **Configured Appliances > (desired appliance group) > Applications**.
2. Click **Add Application from Library**.
3. Select one or more applications to apply to the current appliance group, and then click **application to <group\_name>**.

### To use a Custom Application in a policy definition

When creating a policy, you can use a custom Application to filter traffic to or from this application. The custom

Applications appear in the Application list within the **Filter** section. For more information, refer to [Policies](#) (page 430).

## Configuring application groups in the EMC

The Exinda Management Center provides a comprehensive set of built-in Application Groups for you to use, but you can also define Custom Application Groups.

Although an application can be a member of multiple application groups, to prevent conflict it can only be a member of the application group that is currently monitoring traffic. For example, Skype cannot be added to both the Voice group and the Messaging group because EMC can gather data from only one monitored application group for reporting.

When an Application Group is created in the library, it is applied tenant-wide and is therefore available in every appliance group within an Appliance Group section. If an application group is set for monitoring and/or being used in the optimizer tree, then this application group is pushed to the respective appliances within the push configuration.

If a custom application is added to an application group, where the group does not exist in the configuration of the appliances, the application is first added to the appliances and then the group is imported.

In addition, the following limitations should be noted:

- » If one application within a group is not supported by an appliance, then that application definition will not be sent to that specific appliance.
- » If you try to add an application to an appliance with a firmware version does not support the application, the EMC displays an error for the appliance and the application is not imported. However, the appliance does import the application group along with other settings.

## Where do I find Application Groups?

Application Group library items can be found in **Library > Application Group**. A lock icon in the Monitoring column indicates that an Application Group is in use and cannot, therefore, be deleted.

## To use an Application Group in a policy rule definition

When creating a policy rule, you can use an Application Group to filter traffic to or from the applications within the group. The Application Groups appear in the Application list within the **Filter** section. Refer to [Policies](#) (page 430). for details.

## How do I configure built-in Application Groups?

You can modify application groups either within the configured appliances or from the library.

1. Click the desired application group name to edit.

# Application Groups

Application groups can be used when defining a policy and to monitor groups of applications. An application can be a member of multiple application groups, but can only be a member of one application group that is monitored.

| Name ▲            | Applications |
|-------------------|--------------|
| Database Services | MS-SQL       |
|                   | MySQL        |
|                   | Oracle       |
|                   | PostgreSQL   |
|                   | TDS          |

2. You can configure monitoring status and add or remove applications within this group.

## Application Group

Application groups can be used when defining a policy and to monitor groups of applications. An application can be a member of multiple application groups, but can only be a member of one application group that is monitored.

> Name: Database Services

> Reporting: Include in application group reporting

> Applications:

> In Use: 6

[Update in Library](#) [Cancel](#)

3. You could also view which policies are currently using this application group under **In Use**.

# Application Group

Application groups can be used when defining a policy and to monitor groups of applications. An application can be a member of multiple application groups, but can only be a member of one application group.

> Name: Database Services

> Reporting: Include in application group reporting

> Applications:

▼ In Use: 6

| Type of use | Name                                            |
|-------------|-------------------------------------------------|
| Policy      | Database - Guarantee High 10%-100% - Accelerate |
| Policy      | Database - Guarantee Low 5%-100%                |
| Policy      | Database - Guarantee Med                        |
| Policy      | Database - Guarantee Med 8%-100%                |
| Policy      | Database - Guarantee Med 8%-100% - Accelerate   |
| Policy      | Database - Limit High 4%-70%                    |

[Update in Library](#) [Cancel](#)

How do I create an custom application group?

1. Go to **Library > Application Groups** and click **Create new application group in the library...**

## Application Group Library

Application groups can be used when defining a policy :  
An application can be a member of multiple application

**+ Create new application group in the library ...**

- Name ▲

- Database Services

2. Provide a name, configure monitoring, if you want this group to be monitored, and add the applications to be part of this group.

# Application Group

Application groups can be used when defining a policy and to monitor  
An application can be a member of multiple application groups, but ca

|                 |
|-----------------|
| > Name:         |
| > Reporting:    |
| > Applications: |

Create Cancel

How do I know which application groups are enabled for monitoring?

By default, all the built-in application groups are enabled for monitoring. On the main **Application Groups** page, you can view the specific groups that are set for monitoring.

## Application Groups

Application groups can be used when defining a policy and to monitor groups of applications.  
An application can be a member of multiple application groups, but can only be a member of one application group that is monitored.

| Name ▲            | Applications | Monitoring                          |
|-------------------|--------------|-------------------------------------|
| Database Services |              | <input checked="" type="checkbox"/> |
|                   | MS-SQL       |                                     |
|                   | MySQL        |                                     |
|                   | Oracle       |                                     |
|                   | PostgreSQL   |                                     |
|                   | TDS          |                                     |

| Applications | Monitoring                          |
|--------------|-------------------------------------|
|              | <input checked="" type="checkbox"/> |
| MS-SQL       |                                     |
| MySQL        |                                     |
| Oracle       |                                     |
| PostgreSQL   |                                     |
| TDS          |                                     |

You can always change the monitoring configuration by clicking on the application group name and changing it.

## Configuring schedules in the EMC

To create a Schedule in the Library

1. Go to **Library > Schedules**.
2. Click **Create new schedule in the library**.
3. In the **Name** section, key-in a name of the schedule item. The name must be unique in the tenant.
4. In the **Times** section, select the days of the week and the **Start Time** and **End Time** for this schedule.
  - To layer the time ranges, click **Add another time range**. For example, if you want to apply a schedule for Monday through to Friday from 9:00 to 17:00, but you need a different start and end time for weekends, you can add another range for Saturday and Sunday.
4. Click **Create**. The Schedule item is added to the Schedules Library category and is then available when defining **Policies** and **Virtual Circuits**.

### NOTE

By default, the **Always** schedule library item is pre-configured. You cannot change or edit it.

## Configuring APS for individual applications in the EMC

In the Exinda Management Center you can identify the applications that are important to the operation of the organization. As you identify these applications, you can assign thresholds against their performance scores that can trigger alerts. As the performance drops below the threshold, configured users then receive email alerts to draw attention to the issue.

The broad process of setting-up performance scoring is to:

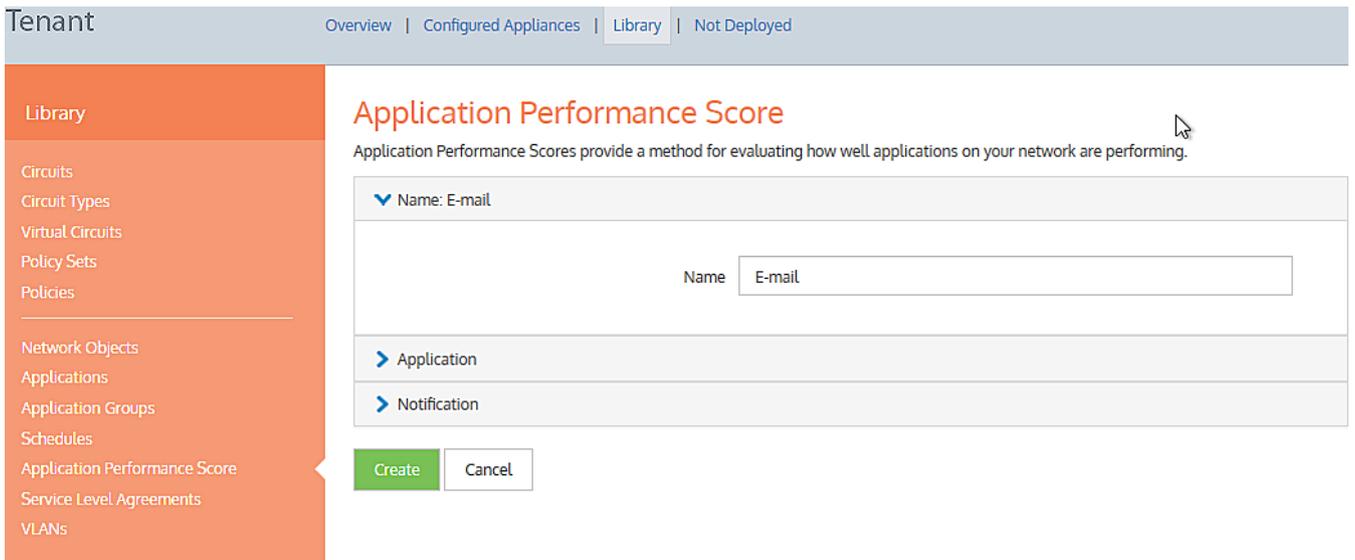
- » Firstly, identify the application, define the APS and threshold, and add it to the Configuration Library. For more information, refer to [Creating Application Performance Score \(APS\)](#) (page 446).
- » Secondly, assign the APS to the Configured Appliances, identify the network objects, and define a baseline period for determining the score. For more information, refer to [Adding APS to the appliances in the EMC](#) (page 448).

Go through these topics for more in-depth information:

### Creating Application Performance Score (APS)

Before you can start to monitor the performance of the important applications on your network, you need to identify those applications in the Configuration Library. The Configuration Library comes with definitions for a very large number of supported applications.

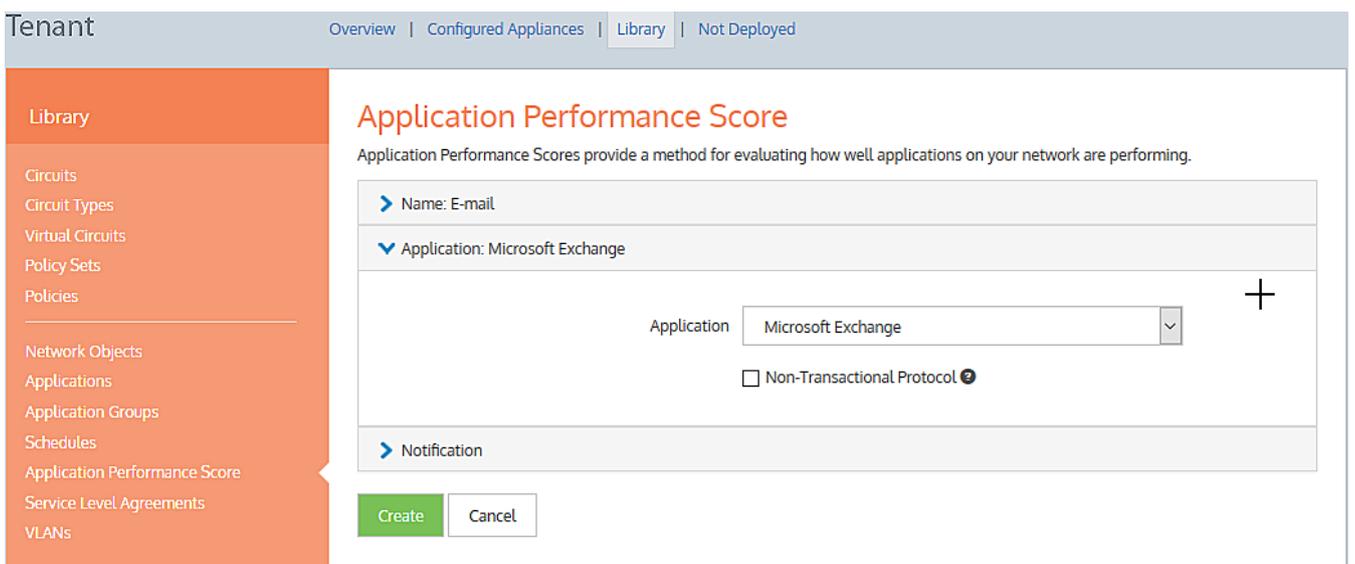
1. In the EMC interface, click **Library > Application Performance Scores**.
2. On the right side, select the **Create new application performance...** link to open the APS set up page.
3. Click **Name** to expand the section. Provide a meaningful name for the new performance score.



Screenshot 199: Setting up a performance score for email services.

4. Click **Application** section to expand.

5. Use the **Application** drop-down list to select the application you need.

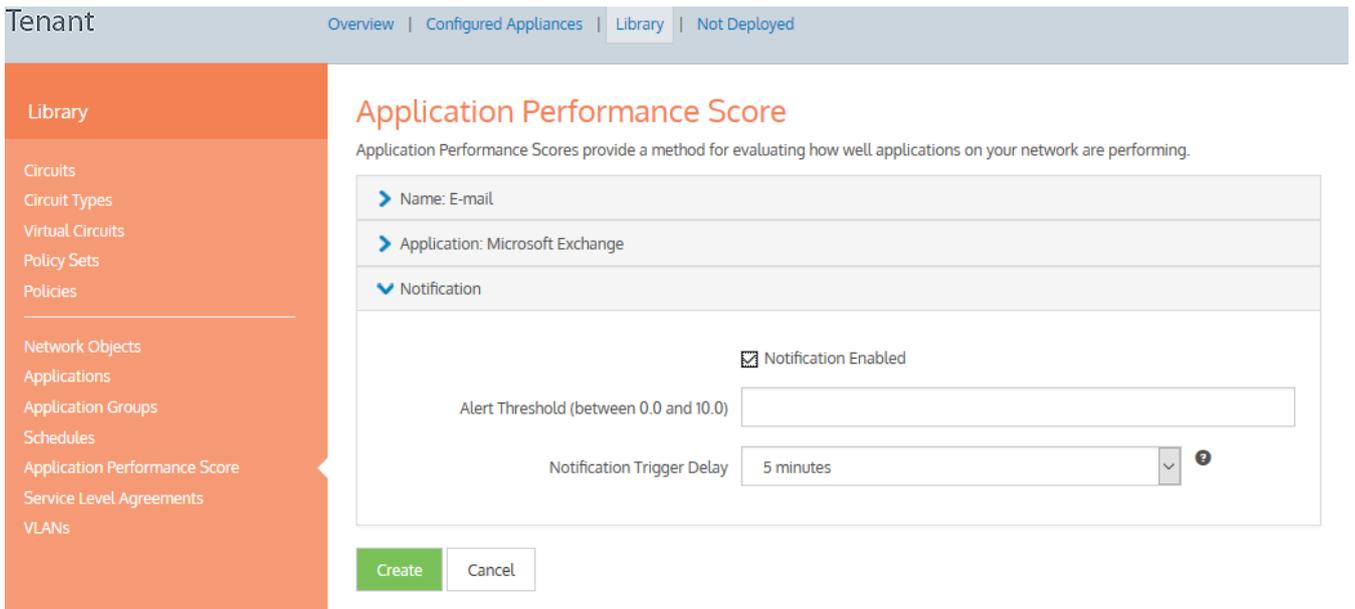


Screenshot 200: Selecting an application

### IMPORTANT

Before proceeding, you should consider the type of protocol used by the application. If it uses a non-transactional protocol, you should select the checkbox. For more information, click the ? button.

6. Click **Notification** to expand the section.



Screenshot 201: Configuring notification settings

7. Configure the following options:

| Option                     | Description                                                                                                                                                                                                                                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Notification Enabled       | Select if you want to be notified when the Alert Threshold is exceeded.                                                                                                                                                                                                          |
| Alert Threshold            | Type an alert level between 0.0 and 10.0. This is a measure of how important is the service the application provides. For example, an application that provides real-time communications would require a higher alert level than one whose responsiveness is not as critical.    |
| Notification Trigger Delay | click the down-arrow and select an acceptable delay period. The notification is only triggered if the condition persists for the entire delay period. Click the  button for more information. |

8. Click **Create**. The screen refreshes to show an entry for the new performance score.

#### NOTE

You can also create APS directly from the **(desired appliance group) > Application Performance Scores** interface. In this case, the APS is applied to the appliance group, and it is also saved in the configuration library.

### Adding APS to the appliances in the EMC

After completing the creation of the APS item definitions in the Configuration Library, you can then apply them to the appliances. When applying the scores, this is a global application; all appliances in the same appliance group receive the same configuration. After applying the scores, you must then push the configuration to the appliances in order to get any notifications.

1. Go to **Configured Appliances > Application Performance Scores**.
2. Click **Add application performance from the library...** link.
3. On the **Add Application Performance Score from Library** pop-up, expand the **Application Performance Score** section.
4. Use the **Application Performance Score** drop-down list to select the APS you need to apply.

## Add Application Performance Score from Library

Application Performance Score: E-mail

Network Object: Measure from All to All

Baseline: For next hour of pushing the configuration

Add Application Performance Score to 'Configured Appliances' Cancel

Screenshot 202: Selecting APS to apply

### NOTE

If you have appliances that do not have the most recent firmware, you may notice the interface reference to the application name is highlighted in red, when you try to add an APS to the configured appliances. You can hover over  to view details about the issue on a tooltip.

5. Expand the **Network Object...** section.

## Add Application Performance Score from Library

Application Performance Score: E-mail

Network Object: Measure from All to All

Internal Network Object: All

External Network Object: All

Baseline: For next hour after pushing the configuration

Add Application Performance Score to 'Configured Appliances' Cancel

Screenshot 203: Selecting internal and external network objects

- Use the **Internal Network Object** drop-down list to select the object that fits your needs.

### TIP

The internal objects are the [Local Network Objects](#) defined in the Configuration Library. The default is A.1.1, but you can choose from any of the others that have been defined.

- Use the **External Network Object** drop-down list to select the previously Network Object that fits your needs.

When complete, the title of the Network Object section changes to summarize the settings you have made.

6. Expand the **Baseline** section and choose the duration of the baseline calculation.

#### NOTE

To establish a baseline for the performance of an application, its performance in the network must be monitored for period of time. The Baseline Length you define is the initial monitoring period, but if the baselining fails, the initial monitoring is automatically extended. The calculation will start at the beginning of the hour the configuration is pushed.

For example, if it is pushed at 4:37, the baseline calculation will start at 4:00. If there is no data for that period, it will extend the duration to the next longest period.

## Add Application Performance Score from Library

> Application Performance Score: E-mail

> Network Object: Measure from All to All

▼ Baseline: For next hour after pushing the configuration

Baseline will be performed for the selected amount of time after the configuration is pushed to the appliance. If baseline does not succeed within the time period the next biggest time period is started.

Baseline Length

Add Application Performance Score to 'Configured Appliances' Cancel

Screenshot 204: Choose the duration of the baseline calculation

7. When finished, click **Add Application Performance Score to 'Configured Appliance'**. The screen refreshes to show an entry for the new performance score.

8. Repeat each of the preceding steps to apply all of the APS definitions you need.

#### IMPORTANT

After applying all the APS definitions, you must push the configuration to all or any particular appliance group. For more information, refer to [Sending configuration changes to the appliances](#) (page 176).

### Removing APS from the Appliances

If any existing APS is no longer required, you can remove it from the appliance configuration. Removing an APS in this way does not remove it from the Configuration Library; the APS item remains there for future reuse. If you do want to remove it from the library, see [Maintaining APS in the Configuration Library](#). But you must remove an APS item from the appliance configuration before deleting it from the Configuration Library.

1. Go to **Configured Appliances > Application Performance Scores**. The currently assigned APS items appear on the page.

2. In the **Name** column, locate the APS item and use the **X** icon on the right to delete.

Screenshot 205: Confirming removal

3. Confirm and click **Remove**.

Are you sure you want to remove this application performance score from the appliance group?

Note: It will still be available in the application performance score library.

Don't ask when removing application performance scores again.

Remove Don't Remove

### Maintaining APS in the EMC configuration library

Over time, it may become necessary to modify the Application Performance Score items stored in the Configuration Library. When modifying an APS item, you are modifying its use wherever it has been applied.

1. In the EMC interface, click **Library > Application Performance Scores**. A listing of the currently defined performance scores appears on the right.
2. In the **Name** column, click on the name of the APS item you need to modify.

---

## Application Performance Score (APS) Library

Application Performance Scores provide a method for evaluating how well applications on your network are performing.

[+ Create new application performance score in the library ...](#)

| Name                   | Application        | Notification                                                                            |
|------------------------|--------------------|-----------------------------------------------------------------------------------------|
| <a href="#">E-Mail</a> | Microsoft Exchange | ✓  |

3. Do the following, as needed:

## Application Performance Score

Application Performance Scores provide a method for evaluating how well applications on your network are performing.

▼ Name: E-Mail

Name

▶ Application: Microsoft Exchange

▶ Notification: Notify when the application's performance score goes below 6 for 5 minutes

▶ In Use: 1

- Modify the APS **Name**.
- Change the **Application**
- Update the **Notification** definition

4. Before completing the modification, expand the **In Use** section to see how the APS items is currently being used.

5. Click **Update in Library**.

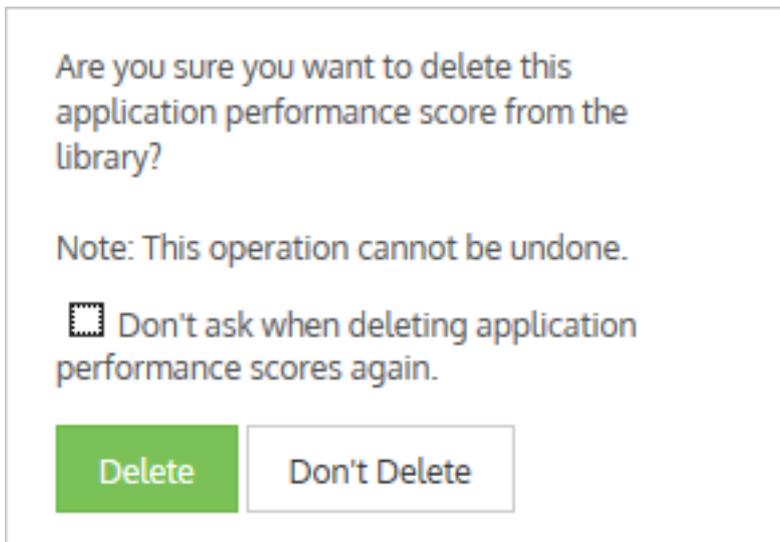
### Deleting APS from the Configuration Library

Over the time, it may become necessary to delete an Application Performance Score item stored in the Configuration Library.

1. Go to **Library > Application Performance Scores**.
2. Locate the APS item to delete from the list and click the  icon.

#### TIP

You can only delete those APS items that are not currently in use. If a  icon appears in the row of an APS item you wish to delete, it means that the item cannot be deleted because it is in use. If you still want to delete it, you first need to unlink it from the appliances. For more information, refer to [Maintaining APS in the EMC configuration library](#) (page 451).



Screenshot 206: Confirming deletion of APS from the Configuration Library

3. Confirm and click **Delete**.

### Configuring service level agreement objects in the EMC

The Service Level Agreement (SLA) library objects are used to monitor the availability of particular IP addresses. By creating an SLA object, you identify the IP address to monitor. The Exinda Management Center then sends one ICMP ping every 10 seconds to the IP address. You can specify the ping packet size to use. You can also specify when an alert is triggered by defining the ping latency threshold and the duration by which the ping latency threshold was exceeded. An alert is triggered when the latency of the SLA site exceeds the latency threshold for longer than the specified duration.

### Where do I find Service Level Agreements?

The Service Level Agreements library items can be found in:

**Library > Service Level Agreements.**

Or,

**Configured Appliances > Service Level Agreements.**

### Setting up a Service Level Agreement in the Configuration Library

1. Go to **Library > Service Level Agreements**.
2. Click **Create new service level agreement in library...**
3. In the **Name** field, key-in a meaningful name for the new Service Level Agreement.

▼ Name

Name

4. Click **Service Agreement** and do each of the following:

- In the **Destination** field, key-in the IP address of the server whose availability you need to monitor.
- In the **Latency Threshold (ms)** field, key-in a value for the response time.
- In the **Ping Size (bytes)** field, key-in a packet size, for example, 1024.
- Select the **Enable Ping** checkbox.

▼ Service Agreement: Ping 192.168.0.25 with ping size of 1024 bytes, allowing for 25 ms of delay. Ping is enabled.

|                        |                                           |
|------------------------|-------------------------------------------|
| Destination IP         | <input type="text" value="192.168.0.25"/> |
| Latency Threshold (ms) | <input type="text" value="25"/>           |
| Ping Size (bytes)      | <input type="text" value="1024"/>         |

Enable Ping

4. Click **Notification**, and select a delay period from the drop-down list. The options are:

- 0 – Disabled – this disables the alert
- 30 seconds
- 60 seconds
- 5 minutes
- 30 minutes
- 1 hour – the Default setting

**TIP**

The default delay is 1 hour. If this setting fits your needs, you do not need to change anything here.

▼ Notification: Notify if threshold is met for 1 hour.

An alert will be sent out after the threshold has been surpassed for the specified amount of time.

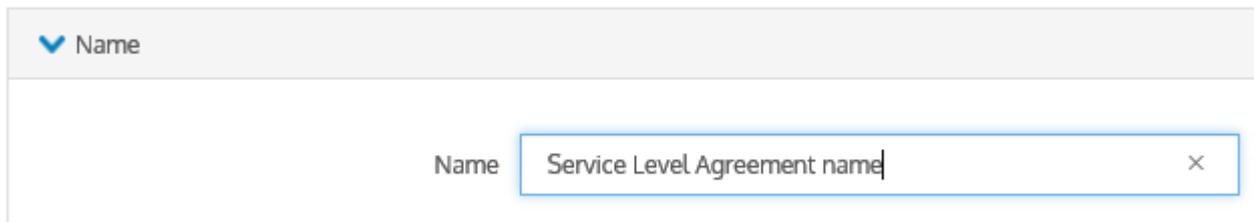
Notification

7. Click **Create**.

### Setting up a Service Level Agreement directly in the Appliances

1. Go to **(desired appliance group) > Service Level Agreements**.
2. Click **Create new service level agreement...**

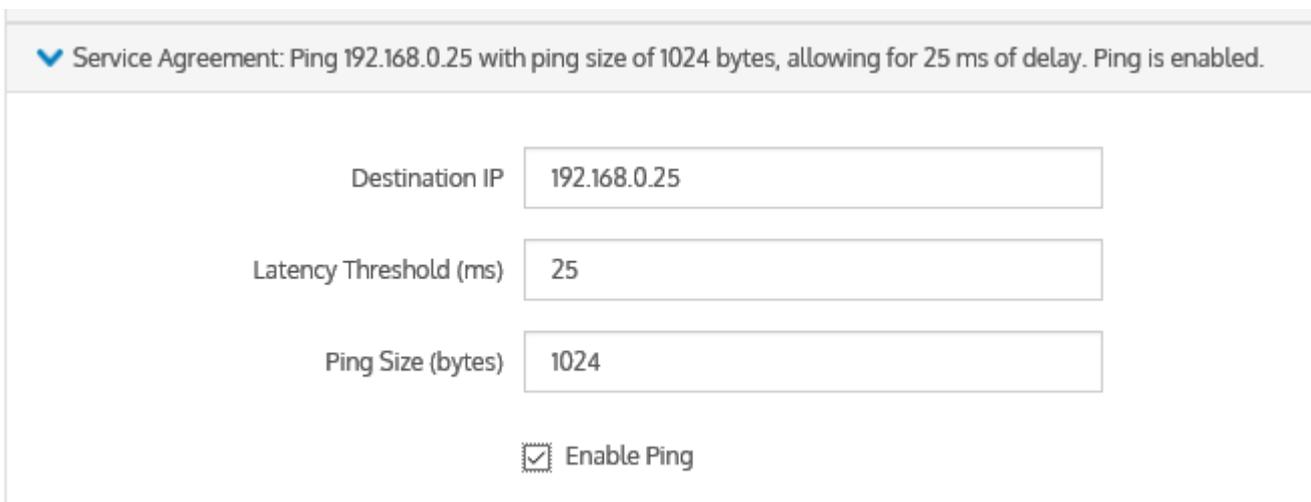
3. In the **Name** field, key-in a meaningful name for the new Service Level Agreement.



A screenshot of a configuration form. At the top, there is a section header "Name" with a downward-pointing chevron icon. Below this, there is a text input field with the label "Name" to its left. The input field contains the text "Service Level Agreement name" and has a small "x" icon in the top right corner, indicating it is a clearable field.

4. Click **Service Agreement** and do each of the following:

- In the **Destination** field, key-in the IP address of the server whose availability you need to monitor.
- In the **Latency Threshold (ms)** field, key-in a value for the response time.
- In the **Ping Size (bytes)** field, key-in a packet size, for example, 1024.
- Select the **Enable Ping** checkbox.



A screenshot of a configuration form for a Service Agreement. At the top, there is a section header "Service Agreement: Ping 192.168.0.25 with ping size of 1024 bytes, allowing for 25 ms of delay. Ping is enabled." with a downward-pointing chevron icon. Below this, there are three text input fields and one checkbox. The first field is labeled "Destination IP" and contains "192.168.0.25". The second field is labeled "Latency Threshold (ms)" and contains "25". The third field is labeled "Ping Size (bytes)" and contains "1024". Below these fields is a checkbox labeled "Enable Ping" which is checked.

4. Click **Notification**, and select a delay period from the drop-down list. The options are:

- 0 – Disabled – this disables the alert
- 30 seconds
- 60 seconds
- 5 minutes
- 30 minutes
- 1 hour – the Default setting

#### TIP

The default delay is 1 hour. If this setting fits your needs, you do not need to change anything here.

▼ Notification: Notify if threshold is met for 1 hour.

An alert will be sent out after the threshold has been surpassed for the specified amount of time.

Notification

1 hour



7. Click **Create**. The Service Level Agreement is added to the appliances and is also saved to the configuration library.

### Modifying a Service Level Agreement

You can modify and update all the properties of the Service Level Agreement except its name.

### Deleting a Service Level Agreement

#### NOTE

You can only delete those SLA items that are not currently in use. If a  icon appears in the row of a SLA item you wish to delete, it means that the item cannot be deleted because it is in use. If you still want to delete it, you first need to unlink it from the appliances.

1. Go to **Library > Service Level Agreements**.

2. In the list of SLA library objects, locate the SLA item you need to delete.

## Service Level Agreement

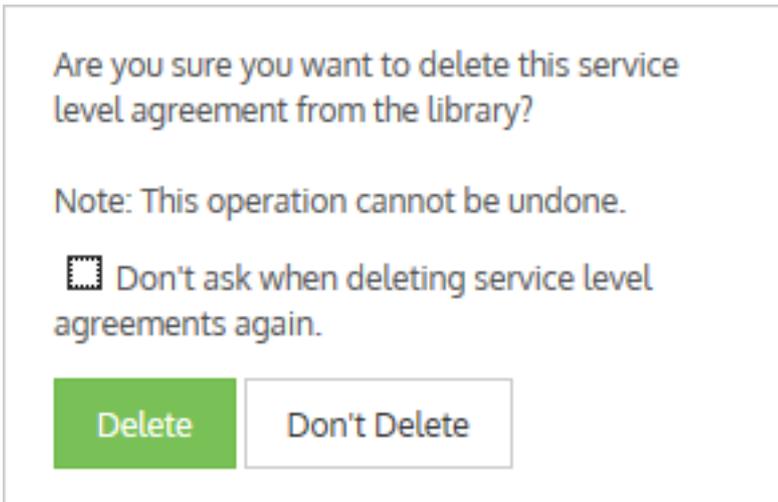
Service level agreements (SLA) provide a method for evaluating how well a server is running.

 Create new service level agreement in the library ...

| Name ▲                       | Destination IP | Threshold | Ping Size (bytes) | Notification | Enabled |                                                                                       |
|------------------------------|----------------|-----------|-------------------|--------------|---------|---------------------------------------------------------------------------------------|
| Service Level Agreement name | 10.10.9.55     | 5 ms      | 1024              | 5 minutes    | ✓       |  |

Delete

3. In the SLA item row, click .



4. Confirm and click **Delete**.

### Configuring VLANs in the EMC

Virtual LAN (VLAN) Objects are used to logically separate hosts (or groups of hosts) on a functional basis rather than on a physical basis. Once VLAN Objects are defined, they can be used in Optimizer policies to filter traffic. By default, the Exinda Appliance has a single VLAN defined called **ALL**, which matches all traffic (regardless if that traffic is part of a VLAN or not). The **ALL** VLAN object is protected and cannot be edited or deleted. Additional VLAN Objects can easily be added.

#### To add a new VLAN object

1. Go to **Library > VLANs**.
2. Click **Create new VLAN in the library...**
3. In the **Name** field, key-in a meaningful name for the VLAN.

## VLAN

Define 802.1Q VLANs to logically group network nodes. These VLANs can then be used in optimizer policies to filter traffic.

A screenshot of a web form for creating a new VLAN. At the top, there is a dropdown menu labeled "Name: VLAN name" with a blue downward arrow. Below this, the "Name" field is highlighted with a blue glow and contains the text "VLAN name".

4. Expand the **Definition** section:

- a. In the VLAN ID **Start** and **End** fields, key-in the range of VLAN IDs that must appear in the VLAN.

#### NOTE

The absolute range is 0–4094. This would equate to all VLAN IDs being in the VLAN. Leaving both fields blank would give the same result. To define a lesser range, type a range somewhere within the absolute range. To isolate one VLAN, type its ID value in both the **Start** and **End** fields.

b. In the VLAN Priority **Start** and **End** fields, key-in the range of values for this VLAN.

#### NOTE

You can define priorities within a maximum range of 0–7. This would equate to all priorities being assigned to the VLAN. Leaving both fields blank would give the same result. To define a lesser range, type a range somewhere between the maximum range. To define just one priority, type the same value in both the Start and End fields.

## VLAN

Define 802.1Q VLANs to logically group network nodes. These VLANs can then be used in optimizer policies to filter traffic.

> Name: VLAN name

▼ Definition

VLAN ID (0 - 4094)

Start:  End:

VLAN Priority (0 - 7)

Start:  End:

4. Click **Create**.

### Deleting a VLAN

#### NOTE

You can only delete those VLAN objects that are not currently in use. If a  icon appears in the row of a VLAN object that you wish to delete, it means that the object cannot be deleted because it is in use. If you still want to delete it, you first need to unlink it from the appliances.

Also the ALL VLAN object is protected and cannot be deleted.

**The "All" VLAN is available by default and cannot be deleted**

1. Go to **Library > VLANs**.

2. In the list of VLAN library objects, find the VLAN you need to delete.

## VLAN Library

Define 802.1Q VLANs to logically group network nodes. These VLANs can then be used in optimizer policies to filter traffic.

+ Create new VLAN in the library ...

| Name ▲            | VLAN ID   | VLAN Priority |        |
|-------------------|-----------|---------------|--------|
| All               | 0 - 4094  | 0 - 7         |        |
| Database          | 0 - 25    | 2 - 4         |        |
| Social Networking | 100 - 500 | 0             | Delete |

3. In the VLAN object row, click

Are you sure you want to delete this VLAN from the library?

Note: This operation cannot be undone.

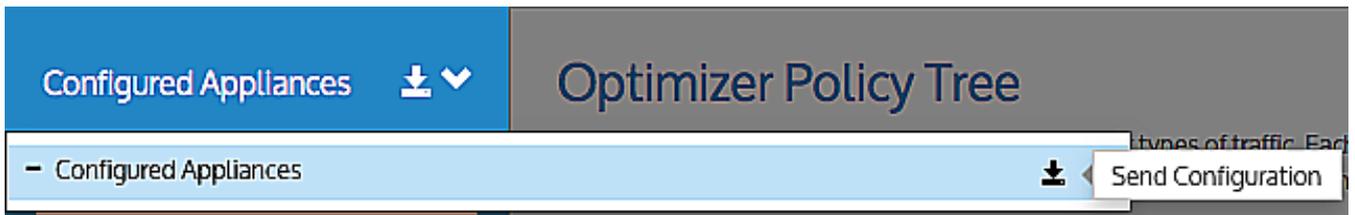
Don't ask when deleting VLANs again.

4. Confirm and click **Delete**.

### 0.15.4 Configuring your Appliances through the CLI

You can use the Command Line Interface (CLI) to push specific commands to the current Appliance Group for instances where the Exinda Management Center User Interface does not support such configuration (e.g., adding VLAN to appliances).

However, if you are sending commands that are sensitive to the order in which they are executed, you should click the **Send Configuration** icon after each command to ensure the correct order of operations is followed. For example, configure the Optimizer Policy Tree and send the configuration to the appliances.

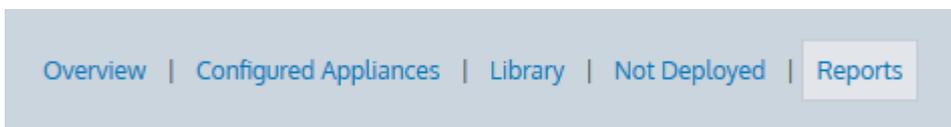


Then send the CLI commands to restart the optimizer and send that to the appliances. Note that you do not need to wait for the appliances to receive the configuration before issuing the next set of configuration or commands.

For more information, refer to [Exinda Command Line Interface \(CLI\)](#) (page 740).

## 0.15.5 Reports in EMC

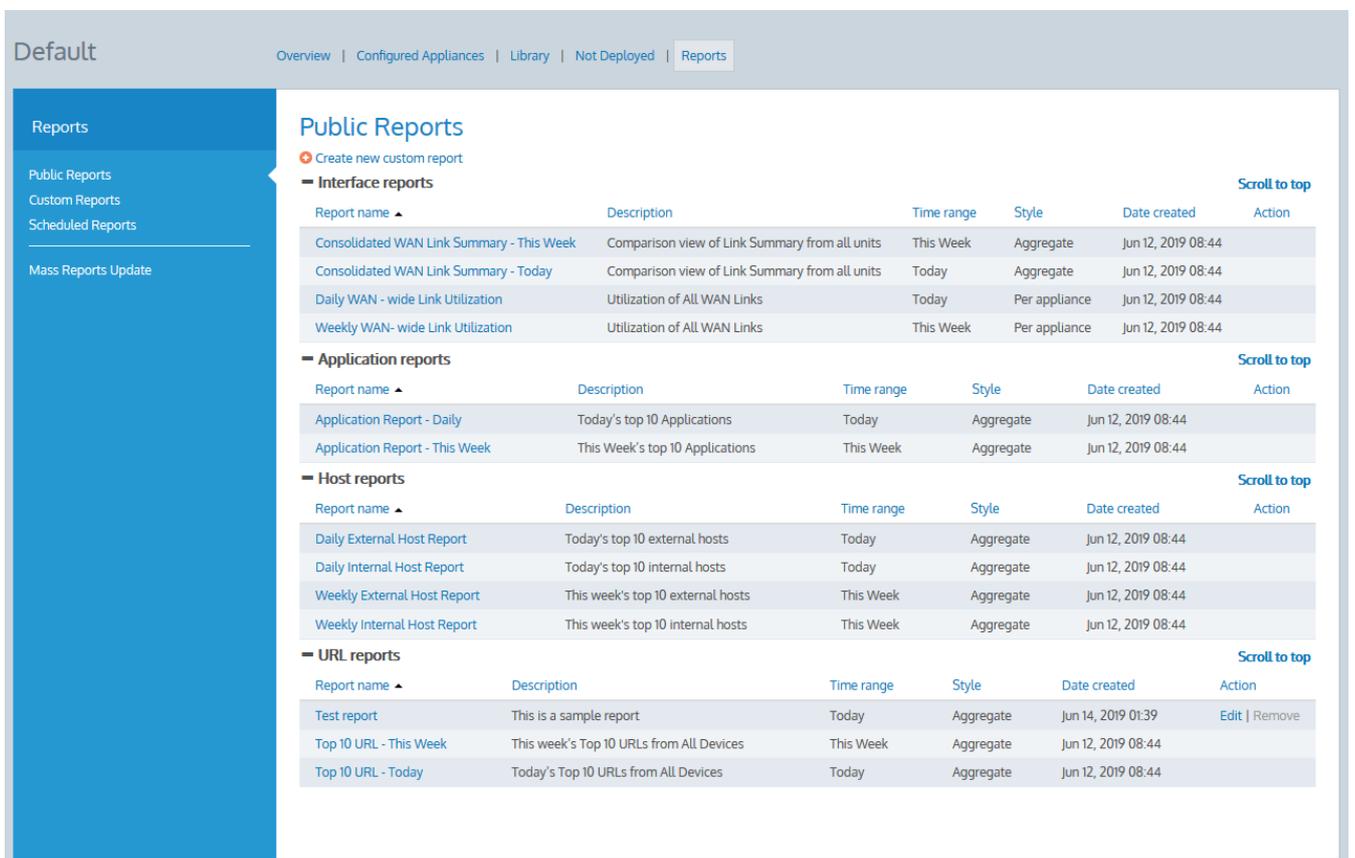
In tenant view, click **Reports** on the top menu to view, edit, or delete existing reports, create new ones, and schedule reports. Besides, the EMC supports updating reports in bulk.



Screenshot 207: The Report option

All reports in the EMC are split by report accessibility. *Public* reports are available for all users. *Private* reports are available only for the user who created them. *Custom* reports which you create can be either public or private.

The EMC also includes a set of predefined public reports. You cannot edit or delete these reports.



Screenshot 208: The Public Reports page

## Time ranges for reports

- » **This hour:** From the time now (rounded to 0 decimal places) to +60 minutes
- » **Last hour:** From the time now (rounded to 0 decimal places) to -60 minutes
- » **Today:** Today from 00h00m00s to 23h59m59s
- » **Yesterday:** Yesterday from 00h00m00s to 23h59m59s
- » **This Week:** From Monday of the current week, 00h00m00s to Sunday of the current week, 23h59m59s
- » **Last Week:** From Monday of the previous week, 00h00m00s to Sunday of the previous week, 23h59m59s
- » **This Month:** From the start of the current month, 00h00m00s to the end of the current month, 23h59m59s
- » **Last Month:** From the start of the previous month, 00h00m00s to the end of the previous month, 23h59m59s
- » **This Year:** From the start of the current year, 00h00m00s to the end of the current year, 23h59m59s
- » **Last Year:** From the start of the previous year, 00h00m00s to the end of the previous year, 23h59m59s

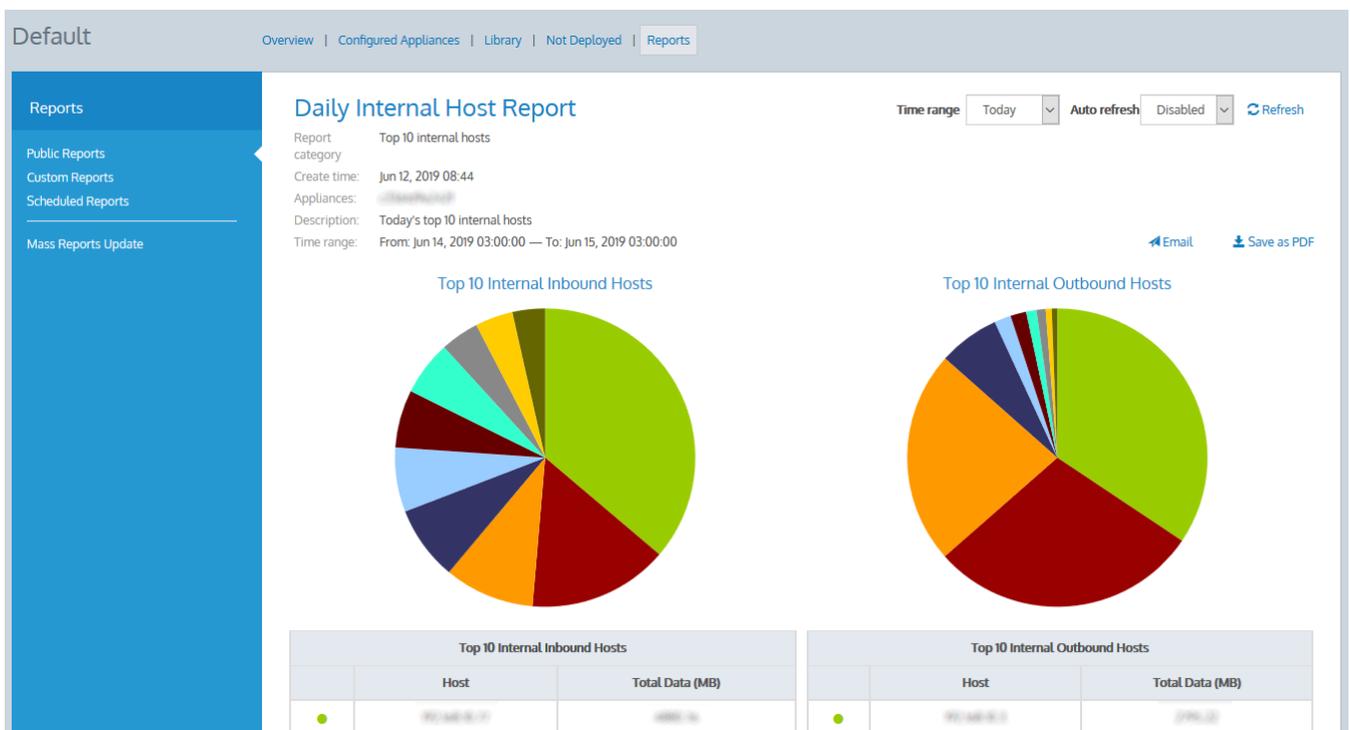
## Viewing reports

You can view any report in your browser, then export the results into PDF and send it by email. Additionally, you can personalize the PDF by adding your company's logo.

To view a report

1. Go to the tenant for which you want to view reports.
2. Go to **Reports > Public Reports** or **Reports > Custom Reports**.
3. Select a report that you want to view from the report list. All reports are listed by category. You can also click on any column to sort records by the selected parameter.

The report page opens with the generated report.



Screenshot 209: A report example

To download personalized report in PDF

1. On the report page, click **Save as PDF**.
2. Specify the report title.
3. Optionally, upload a logo image. The EMC supports jpg, gif, png and bmp image file formats. The maximum image size is 1MB. The image resolution must be 1024x450.
4. Click **Download**.

To generate and email personalized report

1. On the report page, click **Email**.
2. Specify one or more emails (separated by commas) to send the report to.
3. Specify the report title.
4. Optionally, upload a logo image. The EMC supports jpg, gif, png and bmp image file formats. The maximum image size is 1 MB. The image resolution must be 1024x450 px.
5. Click **Send**.

### Creating and editing custom reports

You can create, edit, and delete your own reports. Additionally, you can change visibility, time range, and add or remove appliances for multiple reports in bulk.

To create a custom report

1. Go to the tenant for which you want to create reports.
2. Go to **Reports > Custom Reports** or **Reports > Public Reports**.
3. Click  Create new custom report.
4. Enter a **Report name** and **Report description** for the new report.
5. Select a **Report category**.
6. Choose **Report style** as **Aggregate** to include consolidated data of all the appliances you select, or choose **Per appliance** to include separate data for each selected appliance.
7. Set **Report accessibility** of the report as either **Public** or **Private** as required.
8. Select a **Time range** for compiling data in the report. This is based on your time zone settings.
9. Select the **Appliances** you want to include the report.
10. Click **Create**.

To edit a report

1. Go to the tenant for which you want to edit the report.
2. Locate the report on the **Reports > Custom Reports** or **Reports > Public Reports** page. You can edit public reports and private reports which you created.
3. Click **Edit** next to the report.
4. Change the report settings as required.
5. Click **Save**.

To update multiple reports at once

1. Go to the tenant for which you want to update reports.
2. Go to **Reports > Mass Reports Update**.
3. Under **Public Reports** and **Custom Reports**, select the reports for update.
4. Select appliances to include in this report.
5. Select if you want to **Add** or **Remove** the appliances.
6. Under **Report accessibility**, select if you want to make the reports public or private.
7. Under **Time Range**, specify the report frequency.
8. Click **Save**.

To delete a report

1. Go to the tenant for which you want to delete the report.
2. Locate the report on the **Reports > Custom Reports** or **Reports > Public Reports** page. You can delete public reports and private reports which you created.
3. Click **Delete** next to the report.

### Scheduling reports

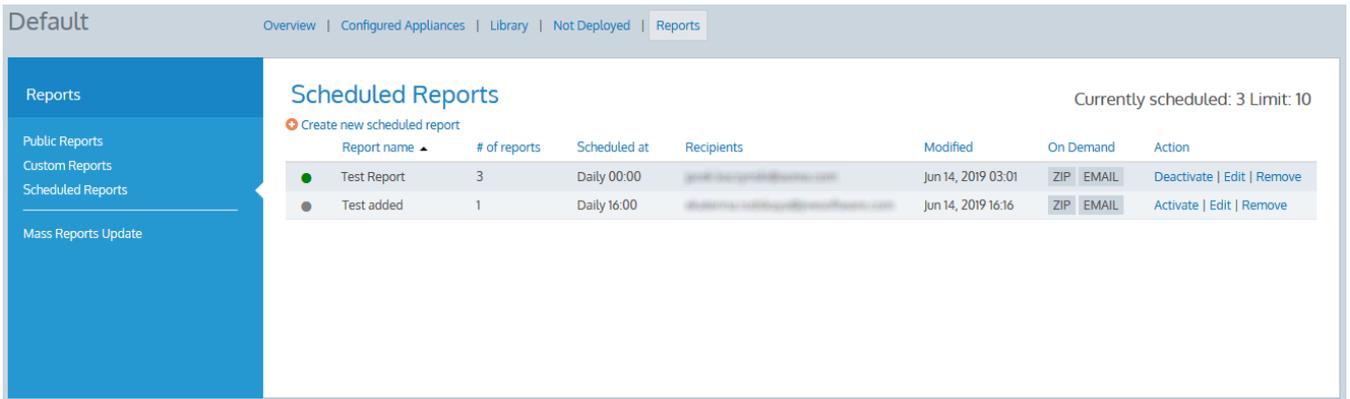
You can schedule reports to run daily, weekly, monthly, or yearly. A PDF with the report is then sent by email to users in your organization. The schedules can be updated if required. You can also resend or download the latest report manually by clicking **ZIP** or **Email** in the **On Demand** column.

Note that you can schedule a maximum of 10 *active* reports, that is, the reports that are generated and sent on schedule. The unused reports can be deactivated.

To schedule a report

1. Go to the tenant for which you want to create reports.
2. Go to **Reports > Scheduled Reports**.
3. Click  **Create new scheduled report**.
4. Enter a **Schedule name** for the scheduled report.
5. Specify email addresses of the recipients to whom you want this report to be sent.
6. Select the **Report frequency** and specify the time when the report is created. The report frequency determines how often you would like this report to be generated and sent. If you select a daily frequency, you're also provided with the list of the days in a week, to include or exclude from.
7. Under **Scheduled state**, **Enable** or **Disable** the scheduler as required.
8. Under **Public Reports** and **Custom Reports**, select the reports to include in this scheduler.
9. Click **Create**.

The report is created and set to **Active** status.



Screenshot 210: The Scheduled Reports page

To activate or deactivate a scheduled report

1. Go to the tenant for which you want to create reports.
2. Go to **Reports > Scheduled Reports**.
3. Select **Deactivate** next to the report to put the it on hold.
4. Select **Activate** next to the report to make it active.

## 0.16 Service Delivery Point (SDP)

The Exinda Service Delivery Point (SDP) is a high performance add-on designed for enterprise network environments looking to centrally manage multi-box Exinda deployments. It is available as both a hosted service and a virtual appliance.

SDP simplifies the tasks of installing, configuring, monitoring and reporting WAN optimization appliances. It is a key differentiator in the traffic shaping & WAN optimization space. A fundamental component of Exinda's Unified Performance Management solution, it rounds out the Exinda product line and makes it the most comprehensive and effective solution for achieving peak application performance.

With secure access via a Web browser, SDP subscribers gain full visibility into network usage and control over applications at any WAN site. SDP helps IT managers identify and control the underlying causes of poor network performance, whether it be unwanted recreational peer-to-peer traffic or a misconfigured server.

Custom reports provide a granular analysis of network usage, top applications and top URLs. This information is critical in setting an optimal network policy, throttling back applications and for future capacity planning.

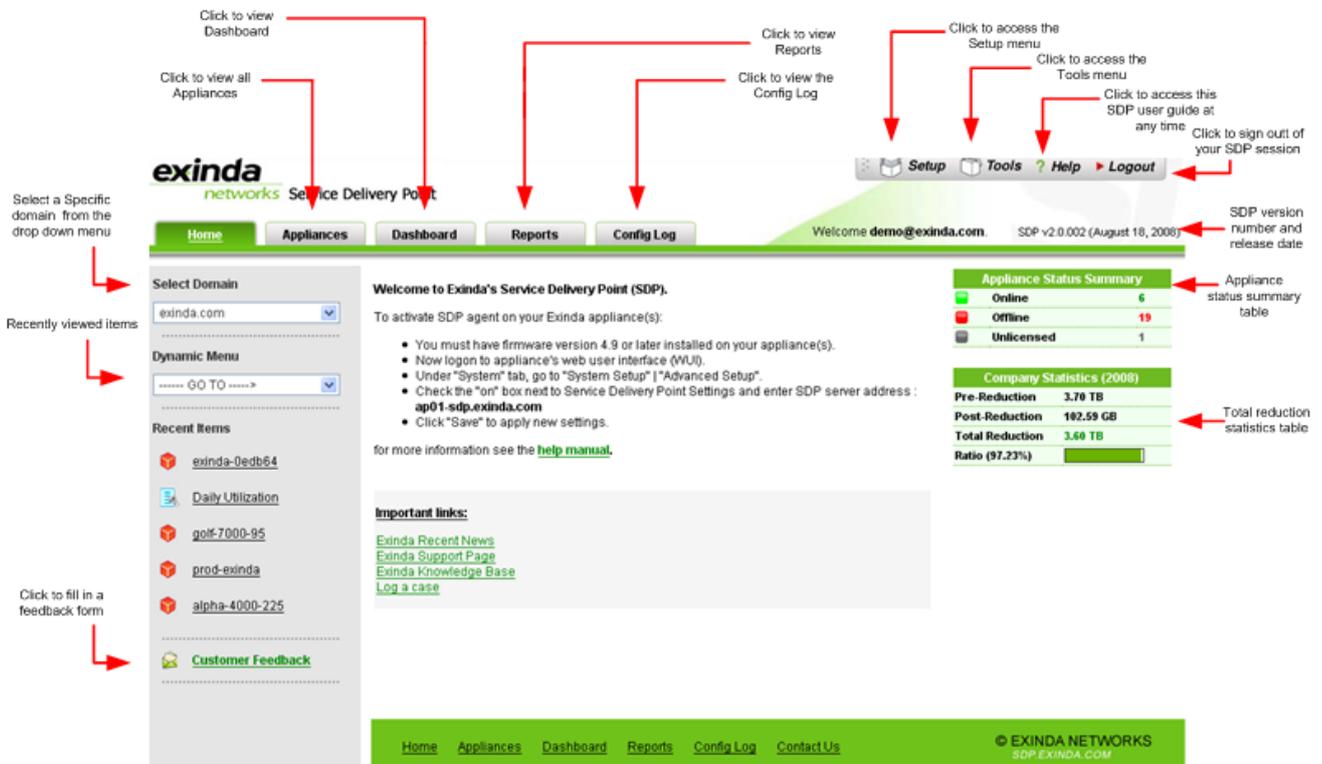
For more information, refer to [SDP Web User Interface \(WUI\) and Features](#) (page 465).

Refer to these topics for a more in depth-understanding on how to set up and use SDP:



## 0.16.1 SDP Web User Interface (WUI) and Features

### WUI



Screenshot 211: Service Delivery Point (SDP) WUI main screen

### Central device manager

- » Displays all devices in a customizable list and provides access to device details and configuration file.
- » Change individual or mass device config.
- » Maintains an archive system for previous device configuration files.
- » Supports manual change of a device config.
- » Deploying policies to a single or multiple device(s).

### Central firmware manager

- » Displays all available firmware updates.
- » Update firmware on a single or multiple device(s).
- » Scheduling future firmware update.

### Central reporting manager

- » Provides access to predefined reports.
- » Custom report wizard enables user creating new reports as well as editing, saving and deleting them.

## Dashboard

- » Dashboard presents a graphical view of overall daily throughput and traffic analysis statistics from a selected device.
- » Drop down menu enables to change dashboard view by selecting different devices.

## User profile setup

- » Change password option.
- » Access to administration tools. (only available to admin user)

## Tools

- » Searching facility to find subscription expiry details of a device.
- » Searching facility to find a purchase order details.

### 0.16.2 Getting started with SDP

Users can register with Exinda SDP via the Exinda web site. Your email address must be the registered 'End User' contact in our database for an Exinda product. The registered email address can be found on the maintenance certificate sent to the end user. You can also request this by sending your serial number or hardware key to [support@exinda.com](mailto:support@exinda.com).

For more information, refer to [SDP appliance system settings](#) (page 472).

## Registering

1. Go to <https://sdp-auth.exinda.com/sdp/registration.php>.
2. Enter your email address
3. Click **Submit**. You will receive a registration confirmation email from Exinda with your default credentials.

## Logging in

1. Go to <https://sdp-auth.exinda.com/sdp/sdp.php>.
2. Enter your email address
3. Enter your password. Your initial password is included in the confirmation email. We recommended that you change your password the first time you login.

### NOTE

If you have an SDP appliance, then go to `https://<ip-address>` to access the SDP host. Where `<ip-address>` is the IP address of the SDP appliance. For further information see [SDP Admin](#) and [SDP Webmin](#).

**Login Details ::**

Email Address ::

Password ::

Screenshot 212: Logging in via the SDP interface.

For more information, refer to [Configuring appliance settings in SDP](#) (page 467).

## Logging out

Click Logout on the menu located at the top right corner to exit your SDP session.



## Password recovery

1. Go to <https://sdp-auth.exinda.com/sdp/forgotpassword.php>.
2. Enter your registered email address.
3. Click **Submit**. An email with a new password will be sent to your email address.

## Configuring appliance settings in SDP

The SDP admin menu allows you to configure your SDP appliance settings. To access the SDP admin menu go to <https://<ip-address>/admin> and log in. The default username and password is `admin`.

You can now view the admin menu. You have the option to add/view users and allocate appliances to them. You can also change the default password and the admin email address, upload a new firmware for later installation and replace appliance in case of hardware replacement.

**SDP Admin Menu** [Help](#) | [Logout](#)

- [Register New User](#)
- [View Users](#)
- [Allocate Appliance](#)
- [Deallocate Appliance](#)
- [Replace Appliance](#)
- [Change Admin Email](#)
- [Change Admin Password](#)
- [Network Settings](#)
- [Add Firmware](#)
- [Upload Firmware](#)
- [Firmware List](#)

**Email Notification Timer Service**

[Start Timer](#) [Stop Timer](#) [Timer Status](#)

To register a new user

1. Click **Register New User**.
2. Fill in all fields.
3. Click **Save**.
4. An email notification with account details will be sent to the new user.

**NOTE**

Multiple entries of the same user is not allowed.

**Register New User**

First Name :  \*

Last Name :  \*

Email : (unique)  \*

[Save](#) | [Back to menu](#)

To view or delete users

1. Click **View Users**. The users table shows all registered users and the appliances that have been allocated to them.
2. You can delete a user by clicking **Delete**.

| SDP Primary User List |           |                        |                                                                                                 |                        |
|-----------------------|-----------|------------------------|-------------------------------------------------------------------------------------------------|------------------------|
| First name            | Last name | Email                  | Appliance(s)                                                                                    | Action                 |
| demo                  | me        | demo@exinda.com        | 0145438<br>0145964<br>1013510<br>1101353<br>1111192<br>1313172<br>3076441<br>3076444<br>3415011 | <a href="#">Delete</a> |
| James                 | Smith     | james.smith@exinda.com | 8713310                                                                                         | <a href="#">Delete</a> |

[Back to menu](#)

To Allocate an appliance

1. Click **Allocate Appliance**.
2. Select the appliance you would like to allocate to a user.
3. Click **Allocate**.
4. Select a user from the drop down menu.
5. Click **Confirm**.

#### NOTE

You cannot allocate an appliance to multiple users.

| Unallocated Appliance List                |              |               |                     | <a href="#">Allocate</a>   <a href="#">Back to menu</a> |
|-------------------------------------------|--------------|---------------|---------------------|---------------------------------------------------------|
| <input type="checkbox"/> Hostname         | Host ID      | IP Address    | Region              |                                                         |
| <input type="checkbox"/> 104              | 0671144      | 172.14.1.104  | Australia/Melbourne |                                                         |
| <input type="checkbox"/> 105-490rc4       | 2821312      | 172.14.1.105  | America/Detroit     |                                                         |
| <input type="checkbox"/> 108              | 0060e0e22bb4 | 192.168.0.108 | Etc/GMT             |                                                         |
| <input type="checkbox"/> 4800             | 0128034      | 192.168.1.58  | Pacific/Auckland    |                                                         |
| <input type="checkbox"/> 491-Server       | 9113212      | 192.168.0.156 | Australia/Hobart    |                                                         |
| <input type="checkbox"/> 6000-client-66   | 0010f30d4964 | 192.168.5.253 | Australia/Melbourne |                                                         |
| <input type="checkbox"/> 6800_EVAL        | 1537139      | 0.0.0.0       | Europe/London       |                                                         |
| <input type="checkbox"/> arm-4800         | 0992012      | 192.168.0.206 | Australia/Melbourne |                                                         |
| <input type="checkbox"/> charlie-4000-155 | 087112008711 | 192.168.0.155 | Etc/GMT             |                                                         |
| <input type="checkbox"/> charlie-4000-178 | 307645430764 | 192.168.0.178 | Australia/Melbourne |                                                         |

| Allocate Appliance |                                                                                                            | <a href="#">Confirm</a>   <a href="#">Back to menu</a> |
|--------------------|------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| Select User :      | <input type="text" value="demo@exinda.com"/>                                                               |                                                        |
| Appliance(s) :     | [104] [105-490rc4] [108] [4800] [491-Server] [6000-client-66]<br>[6800_EVAL] [arm-4800] [charlie-4000-155] |                                                        |

[Confirm](#) | [Back to menu](#)

To Deallocate an appliance

1. Click **Deallocate Appliance**.

2. Select a registered user from the drop down menu.
3. Select the appliance(s) you would like to remove from the user's account.
4. Click **Deallocate**.

**Deallocate Appliance**
[Deallocate](#) | [Back to menu](#)

**Select User :**

**Appliance(s) :**

| <input type="checkbox"/> <u>Hostname</u> | <u>Host ID</u> | <u>IP Address</u> | <u>Region</u>       |
|------------------------------------------|----------------|-------------------|---------------------|
| <input type="checkbox"/> 112             | 221214822121   | 192.168.0.112     | Etc/GMT             |
| <input type="checkbox"/> 146             | 0010f30991b0   | 192.168.0.146     | Australia/Melbourne |
| <input type="checkbox"/> 155             | 0010f3087b20   | 192.168.0.155     | Australia/Melbourne |
| <input type="checkbox"/> 223             | 0010f30fdc2e   | 192.168.0.223     | Etc/GMT             |
| <input type="checkbox"/> 4000-server-164 | 0010f30592bc   | 192.168.1.253     | Australia/Melbourne |
| <input type="checkbox"/> 6000-server-64  | 0010f3126516   | 192.168.4.253     | Australia/Melbourne |
| <input type="checkbox"/> alpha-4000-154  | 0010f30b1a16   | 192.168.0.154     | Etc/GMT             |

To change the Admin email

1. Click **Change Admin Email**.
2. Enter the new email address.
3. Click Save.
4. An email notification will be sent to the new address.

**NOTE**

Default admin email address is *sdp@exinda.com*.

To change the Admin password

1. Click **Change Admin Password**.
2. Enter your current password.
3. Enter your new password.
4. Confirm your new password.
5. Click **Save**.

To add a new firmware version

1. Click **Add Firmware**.
2. Enter the new Firmware name.
3. Enter the File name, size and release date for the new firmware version. This information can be found on the support section of the Exinda Website.
4. Enter the release code. Please email on *sdp@exinda.com* to obtain the release code.
5. Enter a link to the release notes of the new firmware. The release notes can be found on the support section of the Exinda Website.
6. Click **Save**.

### Add Firmware

|                 |                                                                                                             |   |
|-----------------|-------------------------------------------------------------------------------------------------------------|---|
| Firmware Name : | 490                                                                                                         | * |
| File Name :     | 470-upgradeto490                                                                                            | * |
| File Size:      | 14.5MB                                                                                                      | * |
| Release Date:   | 29/05/2007                                                                                                  | * |
| Release Code :  | NU6789                                                                                                      | * |
| Release Note:   | <a href="http://www.exinda.com/public/support/knowledge">http://www.exinda.com/public/support/knowledge</a> | * |

[Save](#) | 
 [Back to menu](#)

To upload a new firmware file

1. Click **Upload Firmware**.
2. Click **Browse** and select the firmware file from your local drive.
3. Click **Save** to start the upload progress.
4. Once the upload is completed navigate back to main menu and select the Firmware list to view the available added or uploaded firmware versions.

#### NOTE

The SDP server will use its local IP address when generating the firmware URL. To assign an external IP address and this alter the firmware URLs navigate to Network Settings, enter the new IP and click Save. The firmware links will be automatically updated.

To replace an existing exinda hardware

1. You can use this option in case of a hardware failure, where an existing appliance will be replaced by a new one. By following this process all the SDP information of the old unit will be transferred to the new one. To start click on Replace Appliance.
2. Enter the Host ID of the old appliance and the Host ID and serial number of the new one.
3. Click **Save**.
4. Then follow the Allocate an appliance steps to add the new appliance to a user account. The allocate process should NOT be performed prior to the replace hardware process.

| Appliance Replacement (RMA)                         |                                           |
|-----------------------------------------------------|-------------------------------------------|
| Current Host ID                                     | <input type="text" value="0010f30b1b92"/> |
| New Host ID                                         | <input type="text" value="0010f3126516"/> |
| New Serial Number                                   | <input type="text" value="EX09107812"/>   |
| <a href="#">Save</a>   <a href="#">Back to menu</a> |                                           |

## SDP appliance system settings

The SDP webmin menu allows you to set up the SDP appliance. To access the webmin menu go to default address and login. The default username and password is admin. If the IP address has been changed then use this link format `https://<ipaddress>:10000`. In order to access SDP webmin on the default IP address, ensure that your network allows the access. Use the settings below to set up the default network configuration of the SDP appliance if required - IP address - 192.168.0.199/Netmask - 255.255.255.0/Default Gateway - 192.168.0.1.

The home page provides a summary of the appliance status.

|                          |                                                                                      |
|--------------------------|--------------------------------------------------------------------------------------|
| <b>System hostname</b>   | sdp64                                                                                |
| <b>Operating system</b>  | Ubuntu Linux 6.06.1                                                                  |
| <b>Webmin version</b>    | 1.360                                                                                |
| <b>Time on server</b>    | Wed Sep 26 13:47:45 2007                                                             |
| <b>CPU load averages</b> | 0.00 (1 min) 0.00 (5 mins) 0.00 (15 mins)                                            |
| <b>Real memory</b>       | 3.86 GB total, 1.61 GB used                                                          |
|                          |  |
| <b>Virtual memory</b>    | 3.05 GB total, 0 bytes used                                                          |
|                          |  |
| <b>Local disk space</b>  | 70.36 GB total, 4.91 GB used                                                         |
|                          |  |

To change time zone

1. Go to **Hardware > System**.
2. Select a time zone from the drop down menu.
3. Click **Save**.
4. Ensure that the System time is correct. If the System time is different to the Hardware time, then click **Set system time to hardware time** or **Set hardware time to system time** button.

To configure the postfix mail server - 1st Option

1. Go to **Servers > Postfix configuration**.
2. Click **Edit Config Files**.
3. Scroll down and replace exinda.com with the customer's domain name.

4. Click **Save**.

To configure the postfix mail server - 2nd Option

1. Go to **Servers > Postfix configuration**.
2. Click **General Options**.
3. Change What domain to use in outbound mail to appropriate value.
4. Change Local internet domain name to appropriate value.
5. Click **Save** and **Apply**.

Network Configuration

1. Go to **Networking > Network Configuration**.
2. Configure the following settings:
  - a. Click **Hostname and DNS client** to change the server's hostname and DNS IP address. Click **Save**.
  - b. Click **Host addresses** to add hosts to the server. If you have changed the hostname, make sure you have an entry for the IP address mapped to new hostname. You should delete any IP addresses to hostname values that are not relevant.
  - c. Click **Routing and Gateways** to alter the default route and change ONLY gateway IP address under Routing configuration activated at boot time.
  - d. Click **Network Interfaces** to change the IP address and click on the ethernet name which has the IP address 192.168.0.199 under Interfaces Activated at Boot Time. Set Static IP address, Netmask and Broadcast fields to the appropriate values and ensure that the Activate at boot time is set to Yes Note: DO NOT change anything in the Interface active now section. Click **Save**. DO NOT click **Save and Apply** button as the system will become unresponsive.

#### NOTE

The server needs to be rebooted for the configuration changes to be applied.

To change password

1. Go to **System > Change passwords**.
2. Enter your old password.
3. Enter your new password and confirm.
4. Click **Change**.

To reboot the server

1. Go to **System > Bootup and Shutdown**.
2. Click **Reboot System**.

#### NOTE

The server needs to be rebooted for the configuration changes to be applied.

To update SDP firmware

1. Download relevant SDP file from the link provided. (An email is sent to customers with the activation key and

download link). If it is to update SDP software filename will be SDP.ear and if it is to update Admin portal, file name will be SDP-admin.ear.

2. Login to webmin admin console via `https://<ip>:10000`.

3. Go to **Others > Upload and Download**.

4. In **Upload and Download** window, specify file to upload by clicking on Browse button.

5. Leave all other values as default.

6. Click **upload** button to upload the file.

7. Then go to **Others > Custom commands**.

8. In the custom commands window, select update-sdp or update-sdp-admin according to uploaded file. (If you uploaded `SDP.ear`, you should select **update-sdp**. If you uploaded `SDP-admin.ear`, then you should select **update-sdp-admin**).

9. Specify activation key provided and click the selected command button.

To perform a factory default

1. Navigate to **Others > Custom commands**.

2. Click **Factory Default**. This action will restore the server to its factory default settings. Any existing data will be erased.

To update your SDP license

1. Navigate to **Others > Custom commands**.

2. Enter the SDP Activation key and number of licenses. This information can be found in your SDP software certificate.

To view the Host ID

1. Navigate to **Others > Custom commands**.

2. Click **Display Host ID**. The host id will be displayed in the next window.

### 0.16.3 Changing the default view settings in SDP WUI

The Exinda Service Delivery Point (SDP) is a high performance add-on designed for enterprise network environments looking to centrally manage multi-box Exinda deployments. It is available as both a hosted service and a virtual appliance.



Use the **Setup** section to configure the view of Exinda Solution Center:

To change the default view settings

1. Click **Setup**.

2. Click **Default View Settings**.

3. Choose between List and Group view.

4. Click **Save**.

To add/remove gadgets from your Dashboard

1. Click **Setup**.
2. Click **Gadget Settings**.
3. Select the gadgets you would like to include in your dashboard.
4. Click **Save**.

To change your password

1. Click **Setup**.
2. Click **Change Password**.
3. Enter your current password .
4. Enter your new password and confirm .
5. Click **Save**.

|                          |       |
|--------------------------|-------|
| Enter Current Password : | ***** |
| Enter New Password :     | ***** |
| Confirm New Password :   | ***** |

To set the time zone:

1. Click **Setup**.
2. Click **Time Zone Settings**.
3. Select your Time Zone from the drop down menu
4. Click **Save**.

To subscribe to daily email updates

1. Click **Setup**.
2. Click **Email Subscription**.
3. Check **Subscribe to daily updates** and **Notify on my actions** options as required.
4. Select a schedule time from the drop down. This defines the time that you will be receiving the daily updates.
5. Click **Save**.

**Subscribe to daily updates**

The daily updates email contains a summary of all changes made to the appliances.

**Notify on my actions**

By default, you are not sent an email when you make the changes. Check this option to be notified when this occurs.

**Scheduled at :** 9.00 AM

To view new features

Go to **Setup** and click **New features**. All new SDP features are listed in this page as they are added to your account. You

can also view the SDP release notes from this page.

If you are the owner of the device you'll have the option to create user accounts, user groups and devices groups. For more information, refer to [Users and groups administration](#) (page 476).

## Users and groups administration

If you are the owner of the device you'll have the option to create user accounts, user groups and devices groups.

Click **Setup** on the menu located at the top right corner.



Then click **Users and Groups** to create, view or edit a user, user group or device group.

| User Setup                      | User Group Setup                      | Appliance Group Setup                      |
|---------------------------------|---------------------------------------|--------------------------------------------|
| <a href="#">Create New User</a> | <a href="#">Create New User Group</a> | <a href="#">Create New Appliance Group</a> |
| <a href="#">View Users</a>      | <a href="#">View User Groups</a>      | <a href="#">View Appliance Groups</a>      |

To create a new user account (under your profile)

1. Click **Setup**.
2. Click **Users and Groups**.
3. Click **Create New User**.
4. Enter the new user's first and last name
5. Enter the new user's email address
6. Click **Save**.

Create User :

|                  |                                                     |   |
|------------------|-----------------------------------------------------|---|
| First Name :     | <input type="text" value="James"/>                  | * |
| Last Name :      | <input type="text" value="Smith"/>                  | * |
| Email : (unique) | <input type="text" value="james.smith@exinda.com"/> | * |

After creating a new user account, you need to set users privileges. You can select which Exinda devices the user will have access to when logging in and whether read or write access is required.

To set user privileges

1. Click **Setup**.
2. Click **Users and Groups**.
3. Click **View Users**.

|                    |         |         |                     |                                                                               |
|--------------------|---------|---------|---------------------|-------------------------------------------------------------------------------|
| jenny@exinda.com   | Jenny   | Thomson | 22/03/2007 02:12 PM | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |
| alex@exinda.com    | Alex    | Simmons | 22/03/2007 02:12 PM | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |
| anthony@exinda.com | Anthony | Kal     | 22/03/2007 02:13 PM | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |
| gin@exinda.com     | Gin     | Tan     | 22/03/2007 02:13 PM | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |
| james@exinda.com   | James   | Smith   | 22/03/2007 02:16 PM | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |

4. Select a user account and click **Set Privilege**.
5. Select the device(s) or device group(s) you wish to allow access for the specific user account.
6. Select **R** for read access and **W** for write access.
7. Click Save.

**Select DeviceGroup and Privilege:**

Sydney Devices  R  W

Melbourne Devices  R  W

**Select Device and Privilege:**

DEFAULT  R  W

EX09101183  R  W

105-490rc4  R  W

ex-production  R  W

**NOTE**

A user's email address cannot be edited.

To create a new user group (under your profile)

Click **Setup**.

Click **Users and Groups**.

Click **Create New User Group**.

Enter a Group Name.

Select the users you wish to include in the group. Hold the ctrl key down for multiple selections.

Click **Save**.

Group Name :  \*

Select User:

jenny@exinda.com

alex@exinda.com

anthony@exinda.com

gin@exinda.com

james@exinda.com

To set user group privileges

1. Click **Setup**.
2. Click **Users and Groups**.
3. Click **View User Groups**.

| Group Name     | Date Created        | Group Member                                          |                                                                               |
|----------------|---------------------|-------------------------------------------------------|-------------------------------------------------------------------------------|
| Melbourne_Team | 22/03/2007 03:24 PM | jenny@xinda.com<br>gin@xinda.com<br>anthony@xinda.com | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |
| Sydney_Team    | 22/03/2007 03:25 PM | alex@xinda.com<br>james@xinda.com                     | <a href="#">Edit</a>   <a href="#">Delete</a>   <a href="#">Set Privilege</a> |

4. Select a group and click **Set Privilege**.
5. Select the device(s) you wish to allow access for the specific group
6. Select **R** for read access and **W** for write access.
7. Click **Save**.

#### NOTE

Any single user can be a member of multiple groups. Those groups may have different privileges on a single device. In that case that particular user will have the highest permission assigned to him individually or in any group for that device.

To create a new appliance group

1. Click **Setup**.
2. Click **Users and Groups**.
3. Click **Create New Appliance Group**.
4. Enter a name for the Appliance Group
5. Select the device(s) you wish to include in this group. Hold the ctrl key down for multiple selections.
6. Click **Save**.

Group Name :

Select User:

DEFAULT

EX09101183

105-490rc4

ex-production

### Configuring SDP policy service

SDP Policy Service provides functionality to minimize user input and to ensure that the network maintains pre-defined service levels. SDP Auto-Pilot and Semi-Assist modes are available for hosted service subscribers only.



To change the Policy Service mode

1. Click **Setup**.
2. Click **Users and Groups**.
3. Choose a mode:

- **Autopilot Mode:** In auto-pilot mode, the Policy Service Manager makes steering changes to the appliance without any form of user intervention. SDP will monitor the results of changes to ensure that service levels are being met.
- **Semi Assist Mode:** In semi-assist mode, the Policy Service Manager recommends configuration changes based on current data analysis. SDP applies recommended configuration changes to the appliance upon user approval.
- **Manual Mode:** Default mode. SDP provides all the data but analysis and applying changes are the responsibility of the user.

4. Click **Save**.

## 0.16.4 Tools

The Tools menu gives the option to search a device's details, such as software subscription (SS) expiration date, firmware version and product/hardware/activation keys. This feature is only available on hosted Exinda SDP. You can also log a configuration recommendation request to Exinda TAC via this page.

Click **Tools** on the menu located at the top right corner.



To search a device

1. Click **Tools**.
2. Click **Expiry Search**.
3. Enter the device's serial number (S/N) or hardware key (H/K)
4. Click **Search**.

The following information is shown when a valid S/N or H/K is entered.

|                             |                                        |
|-----------------------------|----------------------------------------|
| <b>Serial Number :</b>      | <b>EX09101967</b>                      |
| <b>Hardware Key :</b>       | <b>0621312</b>                         |
| <b>Product :</b>            | <b>Exinda 4700</b>                     |
| <b>Licensed Bandwidth :</b> | <b>2</b>                               |
| <b>Product Key :</b>        | <b>CA53-123E1-7170-12B79-D15B-4DAF</b> |
| <b>Activation Key :</b>     | <b>295A4-909B7-2CC2E-6CF75-593CD</b>   |
| <b>Firmware Version :</b>   | <b>450-hf2</b>                         |
| <b>Expiry Date (SS):</b>    | <b>2005-03-20</b>                      |
| <b>Expiry Date (HM):</b>    | <b>0000-00-00</b>                      |
| <b>Expiry Date (PS):</b>    | <b>0000-00-00</b>                      |

### NOTE

The above appliance information is only available on hosted Exinda SDP. The Exinda SDP Appliance offers a limited "Expiry Search" option. The Serial number and SS/HM/PS expiry dates are not available.

To search an order

1. Click **Tools**.
2. Click **Order Search**.
3. Enter the order number
4. Click **Search**.

The following information is shown when a valid order number is entered.

**Order Details**

Purchase Order No: 2952                      Order Date: 2006-12-07

**Billing Address:**                                      **Shipping Address:**

Your address                                              Your address

**Order Items:**

| Part Number / Product Name | Quantity |
|----------------------------|----------|
| EX-4700-15                 | 2        |
| EX-4700-15-HM              | 2        |

**Order Status:**

|                                     |               |                                              |
|-------------------------------------|---------------|----------------------------------------------|
| <input checked="" type="checkbox"/> | Received      | Order received                               |
| <input checked="" type="checkbox"/> | Sales         | Salesforce entered (Acc,Cont,Opp,Partners)   |
| <input checked="" type="checkbox"/> | Accounts pre  | Invoice created                              |
| <input checked="" type="checkbox"/> | Build         | Products set to 'Allocated/Unsent' in Lic DB |
| <input checked="" type="checkbox"/> | Send certs    | Expiration dates in Lic DB updated           |
| <input type="checkbox"/>            | Shipping      | Freight booked                               |
| <input checked="" type="checkbox"/> | Accounts post | Invoice sent                                 |
| <input checked="" type="checkbox"/> | Shipping      | All products shipped                         |

**Serial Numbers:**

| Serial Number | Product     | License | Status          |
|---------------|-------------|---------|-----------------|
| EX09110675    | Exinda 4700 | 15      | Product Shipped |
| EX09110591    | Exinda 4700 | 15      | Product Shipped |

To log a config recommendation request

1. Click **Tools**.
2. Click **Config Recommendation**.
3. Select the appliance for which you would like a config recommendation.
4. Input the Inbound/Outbound bandwidth and latency information.
5. Upload your network information and any other files you feel will be useful for an accurate recommendation. You can upload up to 3 files of 1MB maximum size.
6. Give a detailed explanation of what you are trying to achieve with the exinda.
7. Click **Submit**.
8. Exinda TAC will send you recommendations, based on the information you have submitted, within 48 hours.

## 0.16.5 Managing appliances in SDP

Click **Appliances** on the top menu to enter the device management view. This page lists all the devices owned by the current user. You can also create a custom view. For more information, refer to [Creating a custom view](#) (page 482).

Click on any header to sort devices by the selected parameter. The orange triangle indicates the parameter that the devices are currently sorted by.

If the appliance status is offline it means that it cannot contact the SDP or the SDP option is currently OFF

| HostName ▲                     | Serial No                  | IP Address                     | Status     |                                                                         |
|--------------------------------|----------------------------|--------------------------------|------------|-------------------------------------------------------------------------|
| <a href="#">alpha-4000-141</a> | <a href="#">EX09103143</a> | <a href="#">192.168.0.141</a>  | Offline    | <a href="#">WebUI</a>                                                   |
| <a href="#">alpha-4000-154</a> | <a href="#">EX09112397</a> | <a href="#">0.0.0.0</a>        | Online     | <a href="#">WebUI</a>   <a href="#">Config</a>   <a href="#">Update</a> |
| <a href="#">alpha-4000-225</a> | <a href="#">EX09107819</a> | <a href="#">0.0.0.0</a>        | Online     | <a href="#">WebUI</a>   <a href="#">Config</a>   <a href="#">Update</a> |
| <a href="#">alpha-4x00-146</a> | <a href="#">EX09105565</a> | <a href="#">192.168.0.146</a>  | Offline    | <a href="#">WebUI</a>                                                   |
| <a href="#">Alpha-6000-64</a>  | <a href="#">EX09107812</a> | <a href="#">192.168.10.253</a> | Offline    | <a href="#">WebUI</a>                                                   |
| <a href="#">Alpha-6000-65</a>  | -                          | <a href="#">192.168.50.253</a> | Unlicensed | <a href="#">WebUI</a>                                                   |

Clicking on WebUI will open a new window to the appliance's Web console

When the device is online you have the option to perform firmware upgrades or config alterations.

Click to apply available firmware updates. Valid software subscription is required.

Click config to Access all the configuration options. You can manually change the config, retrieve and save a config or run the Optimizer Wizard

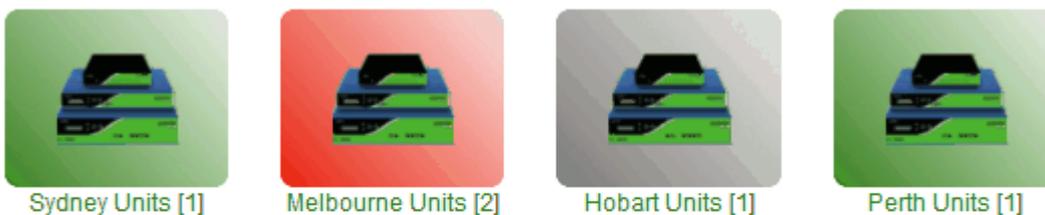
Click on any device to view detailed information. If the device is online you will have the option to alter the config and apply firmware updates.

If the appliance status is Unlicensed it means that it is not a SDP licensed appliance or the limit of licensed appliances has been exceeded

The above picture is a list view of the appliances. A group view is also available.

To create a new appliance group:

1. Click **Appliances** on the top menu to enter the device management view
2. Click **Group View**.
3. Click **Create New Group**.
4. Enter a group name.
5. Select the appliances you wish to include in the group.
6. Click **Save**.



- Empty Group - No appliance or unlicensed appliances.
- Offline Group - One or more appliances currently offline.
- Online Group - One or more appliances online, none offline.

The default view is list. If you would to make group the default view, navigate **Setup > Default view** settings and select **Group view**. This change will be effective next time you log into your account. To delete an appliance group navigate to **Setup > User and Groups Settings > View Appliance Groups**.

When selecting a device the following information is displayed. If the device is online you can view the device's config, edit it or upgrade it. For more information see [Edit Config](#) and [Upgrade Firmware](#).

### Appliance Details

|                       |              |                            |            |
|-----------------------|--------------|----------------------------|------------|
| <b>Host ID</b>        | 0010f30edb64 | <b>Privilege</b>           | W          |
| <b>Serial No</b>      | EX09114721   | <b>Owner</b>               | Y          |
| <b>Version</b>        | 6.0.0.13941  | <b>Expiry Date (SS)</b>    | 2010-09-30 |
| <b>Product Key</b>    |              | <b>Expiry Date (HM)</b>    | 0000-00-00 |
| <b>Activation Key</b> |              | <b>Expiry Date (PS)</b>    | 0000-00-00 |
|                       |              | <b>Expiry Date (SDP)</b>   | 2010-10-01 |
|                       |              | <b>License Expiry (V5)</b> | 2010-02-28 |

### Optional Settings

|                           |                                    |      |                      |                                                                                           |                                       |
|---------------------------|------------------------------------|------|----------------------|-------------------------------------------------------------------------------------------|---------------------------------------|
| <b>Inbound Bandwidth</b>  | <input type="text" value="10000"/> | kbps | <b>Local Subnets</b> | <input type="text" value="192.168.0.0/16"/><br><input type="text" value="172.16.0.0/16"/> | Enter one subnet (a.b.c.d/n) per line |
| <b>Outbound Bandwidth</b> | <input type="text" value="10000"/> | kbps |                      |                                                                                           |                                       |

[Save](#)

[Show Config](#) | [Show Diagnostics File](#) | [Change Config](#) | [Update Firmware](#) | [Config Recommendation](#)

This page also allows you to save and view the diagnostics file. Click **Show Diagnostics** file and then **Get Diagnostics** file. You'll be directed to the Config Log, which will show your request as pending. After 5 minutes go back to the "Show Diagnostics file" page. The diagnostics file will now be available. Send the diag file to Exinda TAC upon request.

Inbound Bandwidth and Outbound Bandwidth are optional settings. These values are stored on the SDP server and can help you manage multiple appliances. When making a Manual Configuration Change, you can specify {inbound\_bandwidth} or {outbound\_bandwidth}, SDP will automatically replace these template variables with the values you have entered here. In addition, the Optimizer Wizard will use these values to help you complete the Wizard.

Local Subnets is also an optional setting. You can specify {create\_local\_networkobjects} in a Manual Configuration Change to create a Network Object called 'Local' containing the subnets you have entered here.

### Related topics

[Create Custom View](#)

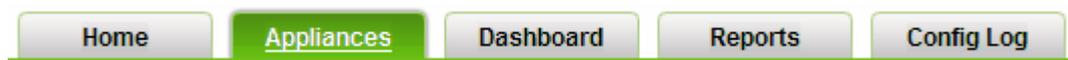
[Edit Config](#)

[Upgrade Firmware](#)

[Configuration Templates](#)

### Creating a custom view

The Appliances page displays a sortable list. You can create your own custom view and select the information you wish to display, set data filtering and sorting rules.



To create a device custom view

1. Click **Appliances** on the top menu to enter the device management view.
2. Click **Create Custom View**.
3. Enter a **Name** for the custom view.
4. Set the accessibility of the custom view to either **Private** or **Public**.

5. (Optional) - Choose the details to be displayed in this custom view.
6. (Optional) - Set search criteria for advance data filtering. Select the parameter you want to filter, set the filter criteria and enter the required value for filtering.
7. (Optional) - Choose a parameter to define how your records are sorted.
8. (Optional) - Choose how you would like to group the records in the view.
9. Click **Save**.

The new custom view will be visible in the **View** drop down menu. You can also edit or delete an existing view.

View:   [Edit View](#) | [Create Custom View](#)

To edit a device custom view

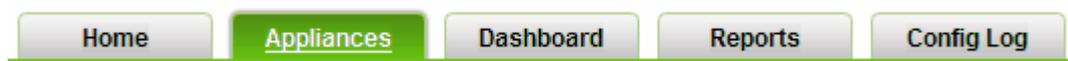
1. Click **Appliances** on the top menu to enter the device management view.
2. Select a view form the drop down menu.
3. Click **Edit View**.
4. Make necessary changes.
5. Click **Save**.

To delete a device custom view

1. Click **Appliances** on the top menu to enter the device management view.
2. Select a view form the drop down menu.
3. Click **Edit View**.
4. Click **Delete**.

### Editing appliance configurations

When a device is online you can view the config and perform any required changes. The available options depend on the firmware version currently installed in your exinda. The options available on V5 appliances differ to the ones available on v4 appliances.



To view an appliance config

1. Go to the **Appliances** page.
2. Select a device. The device status must be online.
3. Click **Show Config** to view the text based configuration.

To edit an appliance config - Option 1 Manual Edit

1. Go to the **Appliances** page.
2. Select a device. The device status must be online.
3. Click **Show Config** to view the text based configuration.
4. Click **Edit Config**.
5. Make necessary changes.

6. Click **Save**.

To edit an appliance config - Option 2 Optimizer Jump Start

1. Go to the **Appliances** page.
2. Select a device. The device status must be online.
3. Click **Change Config**.
4. Select Optimizer Jumpstart and click **Next**.
5. Follow the Optimizer wizard.

To edit the config for multiple devices

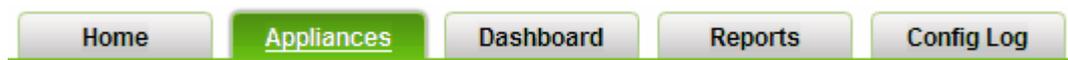
1. Go to the **Appliances** page.
2. Click **Mass Config**.
3. Select the devices you wish to update or select **All**.
4. Choose one of the edit options described above. The available options depend on the firmware version installed on the appliances.
5. Review changes and click **Confirm**.
6. A list off all completed actions will be displayed. You can select to view all or pending only action from the drop down menu. Also, see **Status**.

To restore a previous config

1. Go to the **Appliances** page.
2. Select a device. The device status must be online.
3. Click **Show Config** to view the text based configuration.
4. Click **Edit Config**.
5. Select one of the previous configuration files. You can view and/or restore up to 5 previous config files. They are sorted by date.
6. Copy the text.
7. Return to the **Edit Config** page and paste the copied text.
8. Click **Save**.

## Upgrading appliance firmware

If an appliance is online, you can upgrade to the latest firmware.



To update a device:

1. Go to the **Appliances** page.
2. Select a device. The device status must be online.
3. Click **Update Firmware**.
4. Select a firmware upgrade version.

5. If you want the upgrade to take place in the future, choose a desired date and time. Otherwise click **Next**.

6. Review all information and click **Confirm**.

Available Firmwares :

[Next](#) | [Cancel](#)

| Firmware Updates                                                                                             | Release Date | Update Name          | Size    |
|--------------------------------------------------------------------------------------------------------------|--------------|----------------------|---------|
|  Exinda Firmware v4.9 (rc5) | 05/04/2007   | 490-upgradeto490_rc5 | 14.4 MB |

**Scheduled Firmware Updates**

2007 Sep 7 18 : 00

#### NOTE

Valid Software Subscription is required to update a device.

#### NOTE

The time and date are based on your time zone settings. To configure your time zone settings navigate to **Setup > Timezone settings**.

To update multiple devices:

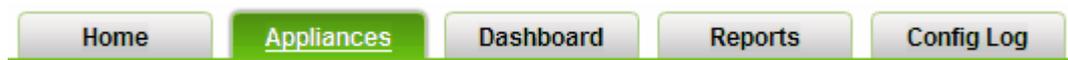
1. Go to the **Appliances** page.
2. Click **Mass Update**.
3. Select the devices you wish to update or select **All**. Click **Next**.
4. Select a firmware upgrade version.
5. If you want the upgrade to take a place in the future, choose a desired date and time. Otherwise click **Next**.
6. Review all information and click **Confirm**.

#### NOTE

The mass update feature is only available for appliances running on firmware version 5.

## Creating and applying configuration templates in SDP

If your appliance is running on V5, then you'll have the option to create and apply configuration templates. Configuration templates allow for easier and quicker configuration of one or multiple appliances. Templates can be saved and reused at any time.



To create a configuration template:

1. Go to the **Appliances** page.
2. Select the **Configuration templates** for the V5 tools
3. Click **Create New configuration template**.

4. Enter a **Name** for the new configuration template.
5. Enter a suitable **Description** to describe your template.
6. Check the **Configuration Option(s)** you want to define in your template and in the pop-up window that appears define its desired parameters.
7. Click **Save**.

To apply a configuration template:

1. Go to the **Appliances** page
2. Select the **Configuration templates** for the V5 tools.
3. Select the template you'd like to apply and click **Apply**.
4. Select the appliance(s) you wish to configure.
5. Click **Next** and confirm the details.

### 0.16.6 The SDP dashboard

The SDP Dashboard is fully customizable. It allows you to create a number of custom views based on predefined reports. All dashboard items are aggregate reports, but you can view appliance individual reports by clicking on the graphs.



To create a Dashboard view

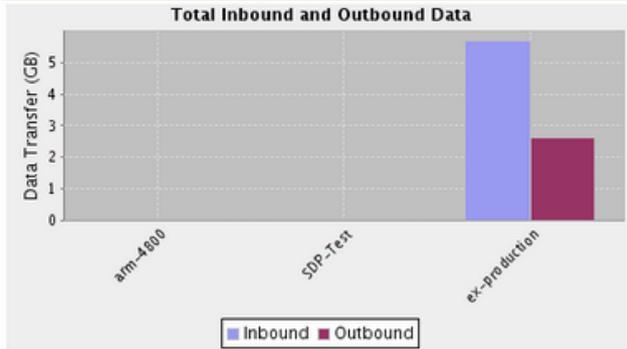
1. Click Dashboard on the top menu to enter to the dashboard view.
2. Click **Dashboard List**.
3. Click **Create New dashboard**.
4. Enter a **Title** for the new dashboard.
5. Enter a **Description** for the new dashboard.
6. Select a layout style.
7. Set accessibility of the dashboard to either **Private** or **Public**.
8. Click **Save**.
9. Go to **Dashboard List**.
10. Click on the dashboard you just created.
11. Click **Edit**.
12. Click **Add Component**.
13. Enter a **Title** for the new component.
14. Enter **Description** to describe the new component.
15. Select a predefined **Source Report**.
16. Click **Save**.

You can repeat steps 9-16 to include additional components in your dashboard. Use the arrows  to arrange all components on your dashboard. Click on the graphs to view complete individual reports.

[Add Component](#)

WAN view - This Week

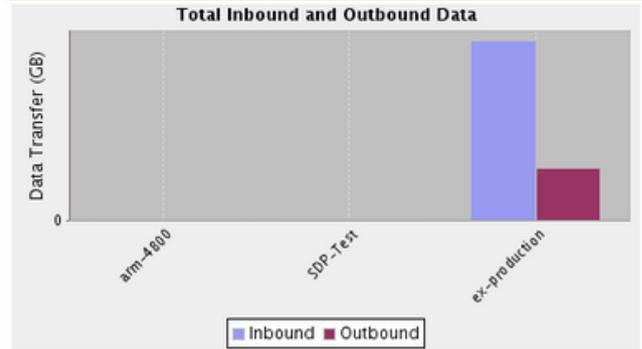
[Refresh](#) [Zoom](#) [Fullscreen](#) [Edit](#) | [Delete](#)



[Add Component](#)

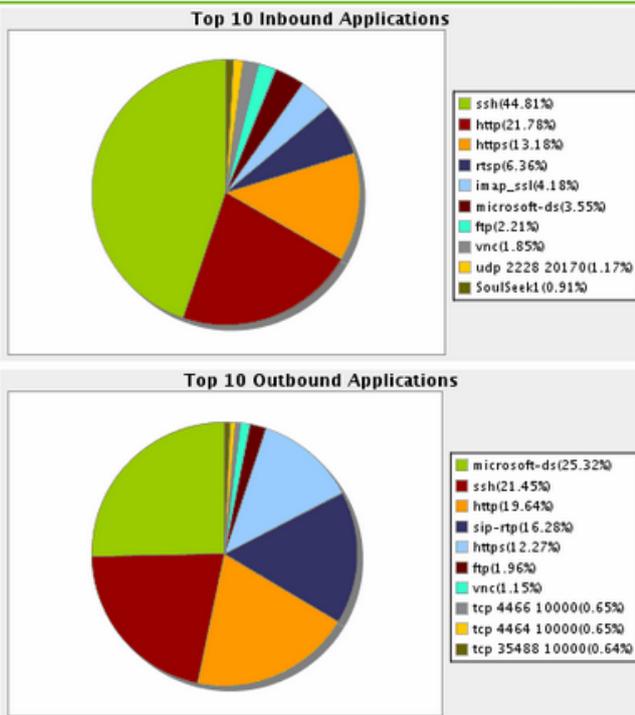
WAN view - Today

[Refresh](#) [Zoom](#) [Fullscreen](#) [Edit](#) | [Delete](#)



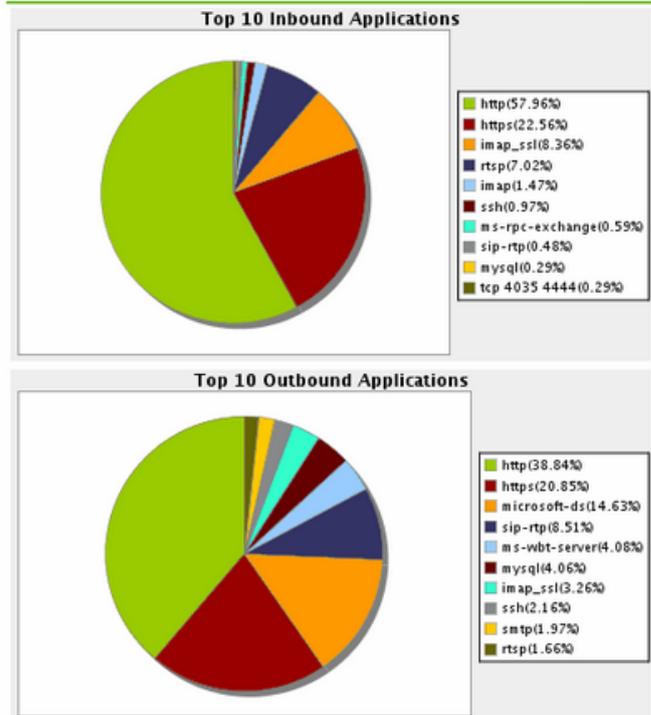
Applications - This Week

[Refresh](#) [Zoom](#) [Fullscreen](#) [Edit](#) | [Delete](#)



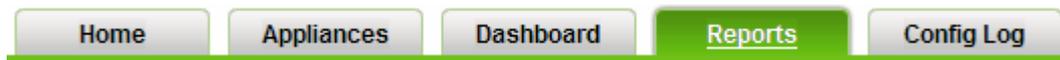
Applications - Today

[Refresh](#) [Zoom](#) [Fullscreen](#) [Edit](#) | [Delete](#)



### 0.16.7 Viewing reports in SDP

Click **Reports** on the top menu to view or edit existing reports, create new ones and set report schedules.



To create a custom report:

1. Click **Reports** on the top menu to enter the device management view
2. Click **Create Custom Report**.
3. Enter a **Name** and **Description** for the new report.

4. Select the **Report Category**.
5. Choose **Report Style** as **Aggregate** to include consolidated data of all the appliances you select, or choose **Per appliance** to include separate data for each selected appliance.
6. Set **Accessibility** of the report as either **Public** or **Private** as per the need.
7. Select a **Date Range** for compiling data in the report. This is based on your time zone settings.
8. Select the **Appliances** you want to include the report.
9. Click **Save**.

To schedule a custom report:

1. Click **Reports** on the top menu to enter the device management view.
2. Click **Schedule a Report**.
3. Enter a **Name** for the scheduled report.
4. Specify email addresses of the recipients to whom you want this report to be sent.
5. Select the **Report Frequency**. This determines how often you would like this report to generated and sent. If you select a daily frequency, you're also provided with the list of the days in a week, to include or exclude from.
6. **Enable** or **Disable** the scheduler as required.
7. Select the **Reports** to include in this scheduler. The maximum reports you can include at a time is 10.
8. Click **Save**.

#### NOTE

To view all your schedules navigate to Reports > Report Central. You can also enable/disable, delete or edit a report schedule from this page.

| Display Name                                               | Number of Reports | Popularity | Scheduled At | Recipient(s)   | Last Modified       | On Demand | Action                                                                  |
|------------------------------------------------------------|-------------------|------------|--------------|----------------|---------------------|-----------|-------------------------------------------------------------------------|
| <input checked="" type="checkbox"/> Monthly Network Report | 4                 | Monthly    | 16.00 PM     | demo@xinda.com | 18/05/2009 01:58 PM |           | <a href="#">Disable</a>   <a href="#">Edit</a>   <a href="#">Delete</a> |
| <input checked="" type="checkbox"/> Daily Interface Report | 1                 | Daily      | 15.00 PM     | demo@xinda.com | 18/05/2009 01:57 PM |           | <a href="#">Disable</a>   <a href="#">Edit</a>   <a href="#">Delete</a> |
| <input checked="" type="checkbox"/> Weekly SLA Report      | 2                 | Weekly     | 15.00 PM     | demo@xinda.com | 18/05/2009 01:58 PM |           | <a href="#">Enable</a>   <a href="#">Edit</a>   <a href="#">Delete</a>  |

To view a report:

1. Click **Reports** on the top menu.
2. Select a report that you want to view from the report list.
3. All reports are listed by category. You can also click on any column to sort records by the selected parameter.

| All Reports                                               |                                                                            |            |               |                     |                                               |  |
|-----------------------------------------------------------|----------------------------------------------------------------------------|------------|---------------|---------------------|-----------------------------------------------|--|
| Interface Reports                                         |                                                                            |            |               |                     |                                               |  |
| Report Name                                               | Description                                                                | Time Range | Style         | Date Created        | Action                                        |  |
| <a href="#">Consolidated WAN Link Summary</a>             | <a href="#">Comparison view of Link Summary from all units</a>             | Today      | Aggregate     | 21/09/2007 04:37 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">wan link</a>                                  | <a href="#">test wan link</a>                                              | Last Week  | Aggregate     | 20/07/2007 12:36 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Consolidated WAN Link Summary - Last Week</a> | <a href="#">Last week's comparison view of link summary from all units</a> | Last Week  | Aggregate     | 19/09/2007 10:46 AM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Daily WAN-wide Link Utilization</a>           | <a href="#">Utilization of all WAN links</a>                               | Today      | Per Appliance | 01/02/2007 11:11 AM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Daily Utilization</a>                         | <a href="#">Today's link utilization</a>                                   | Today      | Per Appliance | 21/09/2007 05:11 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Weekly Utilization</a>                        | <a href="#">Last week's link utilization</a>                               | Last Week  | Aggregate     | 20/07/2007 05:44 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| Application Reports                                       |                                                                            |            |               |                     |                                               |  |
| Report Name                                               | Description                                                                | Time Range | Style         | Date Created        | Action                                        |  |
| <a href="#">WAN-wide Applications</a>                     | <a href="#">Top Applications</a>                                           | Today      | Aggregate     | 23/01/2007 10:31 AM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Weekly Application Report</a>                 | <a href="#">Last week's top 10 applications</a>                            | This Week  | Per Appliance | 17/07/2007 02:42 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Daily Application Report</a>                  | <a href="#">Today's top 10 applications</a>                                | Last Week  | Per Appliance | 22/05/2007 04:28 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| Host Reports                                              |                                                                            |            |               |                     |                                               |  |
| Report Name                                               | Description                                                                | Time Range | Style         | Date Created        | Action                                        |  |
| <a href="#">Daily Internal Host Report</a>                | <a href="#">Today's top 10 internal hosts</a>                              | Last Hour  | Aggregate     | 20/07/2007 03:01 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Weekly Internal Host Report</a>               | <a href="#">Last week's top 10 internal hosts</a>                          | Last Week  | Aggregate     | 22/05/2007 04:32 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Weekly External Host Report</a>               | <a href="#">Last week's top 10 external hosts</a>                          | Last Week  | Aggregate     | 22/05/2007 04:32 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Daily External Host Report</a>                | <a href="#">Today's top 10 external hosts</a>                              | Last Week  | Per Appliance | 22/05/2007 04:29 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| URL Reports                                               |                                                                            |            |               |                     |                                               |  |
| Report Name                                               | Description                                                                | Time Range | Style         | Date Created        | Action                                        |  |
| <a href="#">Top10URL</a>                                  | <a href="#">Today's Top10 URL, All Devices</a>                             | Today      | Aggregate     | 21/02/2007 01:11 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Weekly URL Report</a>                         | <a href="#">Last week's top 10 urls</a>                                    | Last Week  | Aggregate     | 20/07/2007 02:52 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |
| <a href="#">Daily URL Report</a>                          | <a href="#">Today's top 10 urls</a>                                        | Last Week  | Per Appliance | 20/08/2007 05:17 PM | <a href="#">Edit</a>   <a href="#">Delete</a> |  |

## NOTE

You can personalize your reports by adding your company's logo. To load your own logo navigate to Reports and click on an existing report of your choice. Once you are able to view the report you can choose to email or save it. Both these options will give you the choice of uploading a locally saved image as your logo before completing the email/save action. The Title image field is optional and supports jpg, gif, png and bmp image file formats only (max 1MB). The image will be scaled to fit in the PDF report. For best image quality, we recommend to upload image resolution of 1024 x 450 or similar.

## NOTE

All reports are based on your time zone settings. To configure your time zone settings navigate to **Setup > Timezone** settings.

## 0.16.8 Viewing the config log in SDP

Click **Config Log** on the top menu to enter the Config Log page. The log shows a list of actions including user login, configuration changes and firmware updates. You can choose to view all actions, completed actions or pending actions from the drop down menu. The config log is divided into the access containing the user login and configuration, containing the configuration changes and firmware updates. You can view both log types or select a specific one by checking the appropriate tick box.

Auto Refresh:   
 Disabled   
 Disabled   
 5 minutes   
**10 minutes**   
 15 minutes

You can select to refresh this page every 5, 10 or 15 minutes. If you don't wish this page to be automatically refreshed select "Disable" from the drop down menu.

Detailed action information

Action type specifies whether the action is related to a config or firmware update.

The status of the action can be completed or pending.

| Actions                  | Type   | Applied By                   | Time Entered        | Appliance   | Status    | Time Executed       |
|--------------------------|--------|------------------------------|---------------------|-------------|-----------|---------------------|
| no dns secondary         | Config | demo@exinda.com              | 10/12/2007 01:39 AM | Test-Client | Pending   | -                   |
| dns primary 192.168.0.1  | Config | demo@exinda.com              | 10/12/2007 01:39 AM | Test-Client | Pending   | -                   |
| region Australia/Sydney  | Config | roshan.hendahewwa@exinda.com | 07/12/2007 05:46 AM | Node-2      | Completed | 07/12/2007 05:50 AM |
| manual config            | Config | demo@exinda.com              | 05/12/2007 01:12 AM | Node-2      | Completed | 05/12/2007 01:15 AM |
| Exinda Firmware v4.90-03 | Update | demo@exinda.com              | 04/12/2007 11:38 PM | DEFAULT     | Completed | 04/12/2007 11:38 PM |
| Exinda Firmware v4.90-03 | Update | demo@exinda.com              | 04/12/2007 11:28 PM | DEFAULT     | Completed | 04/12/2007 11:28 PM |
| Exinda Firmware v4.90-03 | Update | demo@exinda.com              | 04/12/2007 11:12 PM | DEFAULT     | Completed | 04/12/2007 11:13 PM |
| Exinda Firmware v4.90-03 | Update | demo@exinda.com              | 04/12/2007 08:15 AM | DEFAULT     | Completed | 04/12/2007 08:18 AM |

If the status is complete, then the executed time is recorded. If the action is pending then this field will be blank.

Email address of the user who requested the action.

The time the changes were submitted. This time is base on your time zone settings.

The hostname of the appliance that the changes were Applied.

**NOTE**

The time/date are based on your time zone settings.To configure your time zone settings navigate to **Setup > Timezone settings.**

**0.16.9 Install SDP as a Virtual Appliance**

Exinda's Service Delivery Point (SDP) is a revolutionary platform for centrally managing Exinda appliances distributed throughout the corporate network.

Exinda's Service Delivery Point (SDP) is available as fully featured virtual appliances. There is no difference between the software that runs on bare metal hardware, and the software that runs on virtual appliances. As a virtual appliance, SDP further simplifies the task of installing, configuring, monitoring and reporting WAN optimization. Exinda provides full support for Virtual Appliances running on VMware vSphere (ESX and ESXi) version 5.1 and later.

**NOTE**

Particular attention should be paid to the sizing and licensing requirements of SDP on Virtual Appliances. See [Understanding virtual appliance resource requirements](#) and [License the SDP Virtual Appliance](#)

**License the SDP Virtual Appliance**

The license for SDP is generated by the Exinda Network Support Services team, based on the Host ID of the virtual appliance.

1. Obtain the host ID of the virtual appliance.

- a. In a browser, go to **https://<vm\_ip\_address>:10000**.
  - b. Login with the username **admin** and the password **exinda**.
  - c. In the menu, select **Others > Custom Commands**.
  - d. Click **Display Host ID**. The Execute Command page is displayed with the Host ID.
2. Copy the Host ID, and send an email with the information at [support@exinda.com](mailto:support@exinda.com).
  3. After receiving the license from Exinda, enter the license in SDP.
    - a. In a browser, go to **https://<vm\_ip\_address>:10000**.
    - b. Login with the username **admin** and the password **exinda**.
    - c. In the menu, select **Others > Custom Commands**.
    - d. Paste the license into the Activation Key field, and type the number of licenses.
    - e. Click **Update SDP Licenses**. The SDP virtual appliance is now licensed.

### Requirements for the ESXi server running the virtual appliance

For the virtual appliance to perform at a good level, the ESXi server that hosts the SDP virtual appliance must meet the following minimum requirements:

- » CPU—one quad-core CPU. The CPU must be VT Enabled and 64-bit.
- » RAM—8 GB
- » Storage—120GB. The underlying storage should be RAID-based, ideally RAID 10.

These are guidelines only and there are several factors that may require more or less resources in individual environments. For example, the quality, speed, and performance of the CPU, quality and performance of the NICs, and the available bandwidth.

### Install the SDP Virtual Appliance

The VMware vSphere client must be installed on a server that meets the minimum requirements specified in [Requirements for the ESXi server running the virtual appliance](#) (page 491).. The SDP software is included in the virtual appliance, and is available as soon as the virtual appliance is installed.

1. Download and install the VMware vSphere client.
2. Download the SDP Virtual Appliance package at: [http://updates.exinda.com/sdp/virtual/vmware/release/current/sdp\\_appliance.ovf](http://updates.exinda.com/sdp/virtual/vmware/release/current/sdp_appliance.ovf)
3. Open the VMware vSphere client.
4. Select **File > Deploy OVF Template**.
5. Copy the URL of the latest release of the SDP Virtual Appliance from Exinda.com, and paste it into the **Deploy from...** field. Click **Next**.
6. Confirm the OVF template details are correct, and click **Next**.
7. Review and accept the End User License Agreement (EULA). Click **Next**.
8. Specify a name for the virtual appliance. If prompted, choose the location to deploy the virtual appliance. Click **Next**.
9. Choose the format to store the virtual disks for the virtual appliance. Exinda recommends **Thick Provisioning**(the

default). By default, the Virtual Appliance is configured with a single 50GB disk.

10. Connect the network interfaces to the appropriate network.
  - a. Connect the Management interface to a network where you can manage the virtual appliance.
  - b. If you are configuring the virtual appliance for clustering, high availability, or out-of-path deployments, map the AUX interface to the appropriate network. This interface can be left disconnected if it is not required.
11. Click **Next**.
12. Review the deployment settings, and click **Finish**. The Virtual Appliance is deployed.
13. Start the Virtual Appliance.
  - a. Log into the VMware vSphere Client.
  - b. On the **Home** page, click **Inventory** and navigate to the virtual appliance.
  - c. Select the virtual appliance, and click **Power On**.

## Connect the Virtual Appliance to the Network

By default, the virtual appliance is configured to connect to the network through DHCP. Exinda recommends using the following instructions to configure a static IP address for your virtual appliance.

1. Login to the appliance using console window in ESXi with the following credentials:

Username: `admin`

Password: `exinda`

2. At the prompt, type `ifconfig`. The IP Address and interface of the SDP virtual appliance is displayed. Make note of these values. By default, the interface used by the virtual appliance is `eth0`.
3. In a browser, go to `https://<vm_ip_address>:10000`.
4. Login with the username `admin` and the password `exinda`.
5. In the menu, select **Network > Network Configuration > Network Interfaces**.
6. Click the virtual appliance interface name. For example, click **eth0**.
7. In the Boot Time Interface Parameters, select **From DHCP** or specify a static IP address for the virtual appliance.
8. Ensure the following options are selected:
  - Activate at boot?—**Yes**
  - IPv6 addresses—**IPv6 disabled**
  - MTU—**Default**
  - Hardware Address—**Default**
9. Click **Save**.

## Exinda SDP VA

### NOTE

The following CLI commands should be pasted into the Exinda CLI (configure terminal mode) or uploaded via the **System > Maintenance > Import Config** page on the Web UI, advanced mode.

To install the Exinda SDP Virtual Appliance, first assign a physical interface to the virtual infrastructure so that the virtual machine can have network connectivity.

The following command will add the physical interface specified, to a bridge that can later be attached to virtual machines. You should use a spare, unused interface on the Exinda appliance for this purpose (eth2 is usually a good choice). If the interface specified here is "eth2" for example, the bridge will be called "brvm2". You will need to use this bridge later on when configuring the virtual machine's network interfaces.

```
virt interface eth2
```

This command will fetch the Exinda SDP virtual disk image. The file is approximately 800MB, and once downloaded, will be uncompressed to about 2.7GB.

```
virt volume fetch url http://updates.exinda.com/vm/exinda/Exinda-SDP-x86\_64-0.img
```

After the previous command has completed, go ahead and paste the following commands into the CLI to create the Virtual Machine. This will create a VM called "Exinda-SDP".

```
virt vm Exinda-SDP
  virt vm Exinda-SDP arch x86_64
  virt vm Exinda-SDP boot auto-power last
  virt vm Exinda-SDP boot device order hd
  virt vm Exinda-SDP comment "Exinda SDP Server"
  virt vm Exinda-SDP feature acpi enable
  virt vm Exinda-SDP feature apic enable
no virt vm Exinda-SDP feature pae enable
  # Specify the bridge created above when setting the virtual interface.
  virt vm Exinda-SDP interface 1 bridge brvm2
  virt vm Exinda-SDP interface 1 model e1000
no virt vm Exinda-SDP interface 2
  virt vm Exinda-SDP memory 2048
  virt vm Exinda-SDP storage device bus virtio drive-number 1 source file Exinda-SDP-x86_64-0.img mode read-write
  virt vm Exinda-SDP vcpus count 2
```

Now that the VM has been created, navigate to the **System > Virtualization** page on the Web UI, advanced mode. Power on the VM from this screen.

By default, the Exinda SDP Virtual Appliance is not activated or licensed. Please contact Exinda Support for assistance.

#### NOTE

The Exinda SDP's ethernet port will be attached to the bridge created during the first step. This means the virtualization interface (e.g. eth2) on the Exinda appliance will need to be connected to the network in order to access the SDP Virtual Appliance.

#### NOTE

The virtual disk image used for this Exinda SDP Virtual Appliance is an automatically expanding image. When first installed, it will be about 2.7GB, however, it is capable of automatically growing up to 50GB as space is used. Ensure you allow enough free space on the Exinda appliance for the image to grow.

## 0.17 Introduction

Software-defined Wide Area Network (SD-WAN) refers to the concept of combining different WAN interfaces to create a unified WAN.

SD-WAN gives more flexibility and reliability due to automatic failover in case there is an issue with the primary internet link. It also increases the application performance as it maximizes all WAN links simultaneously.

Exinda SD-WAN are highly versatile devices that can provide true bandwidth aggregation in a variety of settings. These include standalone operation, point-to-point communication between two sites, as well as supporting robust data communication across a large global enterprise. In all cases, the Exinda SD-WAN devices enable high performance, high reliability data communications by making concurrent use of all available communication resources. By using the most cost effective communication resources available at a site location, the WAN performance and reliability can be significantly improved in a cost effective manner.

### 0.17.1 Features

- » Ethernet WAN ports are compatible with any broadband modem technology, such as ADSL, ADSL2, ADSL2+, VDSL, VDSL2, T1, Cable Modem, Satellite, 2G, 3G, 4G LTE etc. via 10/100 baseT, 1GigE and 10GigE Ethernet, and USB interfaces.
- » Web-based Exinda SD-WAN Management Interface allowing easy configuration and monitoring.
- » Bonded Internet service is provided through one or more Ethernet LAN ports.
- » Stateful Firewall.
- » Traffic Monitoring.
- » Traffic filtering and interface binding.
- » Dynamic Traffic shaping, prioritization, QoS shaping, and QoS reservations.
- » Seamless support of legacy networks without re-configuration of existing network.
- » Configurable DHCP server for auto-addressing (optionally enabled).
- » Support for DMZ and port forwarding and firewall functions (optionally enabled).
- » UPnP support (optionally enabled).
- » VRRP support (optionally enabled).
- » Parallel DNS queries on all active WAN interfaces for reduced DNS latency and increased reliability.
- » Domain name blocking and DNS redirection.
- » Session Continuity for automatic allocation of available WAN resources without interruption.
- » Syslog server and SMTP MIBv2 support.

### 0.17.2 Hardware series

The Exinda SD-WAN dedicated hardware series include the EXSD-300, EXSD-1000, and EXSD-10G.

The **EXSD-300** model is a rack mountable device with 4 WAN connections and 2 LAN connections that are available through Ethernet connector (RJ-45, 10/100/1000 Mbps) on the front panel. Two USB ports are available for wireless WAN connections. There is also a serial port connector on the front panel reserved for future use.

Additional models include **EXSD-1000** with 4 additional WAN ports model and Performance Pack, **EXSD-1000** with 8 additional WAN copper ports model, **EXSD-1000** with 8 additional WAN fiber optic ports model.

The **EXSD-10G** model is the device with two 10 Gbps fiber optic LAN and two 10 Gbps fiber optic WAN ports and ten 1 Gbps copper WAN ports.

The USB ports on any Exinda SD-WAN device can also be used for resetting the unit to factory default settings.

Each Ethernet connector has two LED indicators associated with it. One LED is lit for a logical connection to a device on the opposite end of the cable. The other LED is lit indicating data transfer. An LED on the front panel indicates the unit is powered up.

## Exinda SD-WAN hardware and software specifications

Below are Exinda SD-WAN hardware (standalone) and software specifications.

### Exinda SD-WAN EXSD-300

#### Exinda SD-WAN EXSD-300 Hardware Specifications

| Specification                              | Details                                                                           |
|--------------------------------------------|-----------------------------------------------------------------------------------|
| Mechanical dimensions                      | 438(W) x 292.1(D) x 44(H) mm, 17.25"(W) x 11.5"(D) x 1.73"(H)                     |
| Weight                                     | 10.4 lbs                                                                          |
| Input power requirement                    | 200 W full-range ATX                                                              |
| USB ports                                  | 2 (for cellular data)                                                             |
| LAN ports (10/100/1000 baseT, auto-sensed) | 4 RJ45 Ethernet WAN ports. 4 more or 8 more RJ45 Ethernet WAN ports add-on option |
| Certifications                             | FCC, CE, RoHS-5, ICES-03, UL, cUL                                                 |
| Operating temperature                      | 32 – 104°F, 0 – 40°C                                                              |
| Operating humidity range                   | 20 - 90% RH                                                                       |
| Storage temperature range                  | -4 – +158°F, -20 – +70°C                                                          |

#### Exinda SD-WAN EXSD-300 Software Specifications

| Specification                        | Details                                                                                                                      |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Max throughput                       | Base: 1 Gbps (load balancing), 300 Mbps (standalone) , 100 Mbps (peered). Higher performance available in EXSD-1000 models   |
| Max number of concurrent IP sessions | 150,000 (1,000,000 with EXSD-1000 and EXSD-10G models)                                                                       |
| Device management                    | SNMP, Remote Syslog, Email alerts                                                                                            |
| DHCP and DNS servers                 | DNS relay, Parallel DNS optimization, DHCP server                                                                            |
| DDNS                                 | Support for dynamic DNS for multiple interfaces                                                                              |
| WAN Configuration                    | Static IP, PPPoE, DHCP or Pass Through. Failover-only or aggregate modes for cellular and other WAN ports. Interface binding |
| Routing                              | NAT and IP forwarding, QoS and inbound/outbound VOIP quality management                                                      |

## Exinda SD-WAN EXSD-1000

### Exinda SD-WAN EXSD-1000 Hardware Specifications

| Specification                | Details                                                                                               |
|------------------------------|-------------------------------------------------------------------------------------------------------|
| Mechanical dimensions        | 438(W) x 292.1(D) x 44(H) mm, 17.25"(W) x 11.5"(D) x 1.73"(H)                                         |
| Weight                       | 10.4 lbs                                                                                              |
| Input power requirement      | 80 plus 250 W full-range ATX                                                                          |
| USB ports                    | 2 (for cellular data)                                                                                 |
| LAN ports (GbE, auto-sensed) | 2 GbE copper RJ45 Ethernet connector. One WAN port to LAN port fail-to-wire in case of a power outage |
| WAN ports (GbE, auto-sensed) | 8 GbE copper WAN ports (model E)                                                                      |
| Certifications               | FCC, CE, RoHS-2, UL                                                                                   |
| Operating temperature        | 32 – 104°F, 0 – 40°C                                                                                  |
| Operating humidity range     | 20 - 90% RH                                                                                           |
| Storage temperature range    | -4 – +158°F, -20 – +70°C                                                                              |

### Exinda SD-WAN EXSD-1000 Software Specifications

| Specification                        | Details                                                                                                                      |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Max throughput                       | 1 Gbps (load balancing), 940 Mbps (standalone) , 800 Mbps (peered)                                                           |
| Max number of concurrent IP sessions | 1,000,000                                                                                                                    |
| Device management                    | SNMP, Remote Syslog, Email alerts                                                                                            |
| DHCP and DNS servers                 | DNS relay, Parallel DNS optimization, DHCP server                                                                            |
| DDNS                                 | Support for dynamic DNS for multiple interfaces                                                                              |
| WAN Configuration                    | Static IP, PPPoE, DHCP or Pass Through. Failover-only or aggregate modes for cellular and other WAN ports. Interface binding |
| Routing                              | NAT and IP forwarding, QoS and inbound/outbound VOIP quality management                                                      |

## Exinda SD-WAN EXSD-10G

### Exinda SD-WAN EXSD-10G Hardware Specifications

| Specification           | Details                                                       |
|-------------------------|---------------------------------------------------------------|
| Mechanical dimensions   | 438(W) x 292.1(D) x 44(H) mm, 17.25"(W) x 11.5"(D) x 1.73"(H) |
| Weight                  | 10.4 lbs                                                      |
| Input power requirement | 80 plus 250 W full-range ATX                                  |
| USB ports               | 2 (for cellular data)                                         |

| Specification                | Details                                                    |
|------------------------------|------------------------------------------------------------|
| LAN ports (GbE, auto-sensed) | 2x 10 GbE fiber Ethernet connectors                        |
| WAN ports (GbE, auto-sensed) | 10x 1 GbE copper WAN ports, 2x 10 GbE fiber Ethernet ports |
| Certifications               | FCC, CE, RoHS-2, UL                                        |
| Operating temperature        | 32 – 104°F, 0 – 40°C                                       |
| Operating humidity range     | 20 - 90% RH                                                |
| Storage temperature range    | -4 – +158°F, -20 – +70°C                                   |
| Storage humidity range       | 5 – 95%, non-condensing                                    |

### Exinda SD-WAN EXSD-10G Software Specifications

| Specification                        | Details                                                                                                                      |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Max throughput                       | 10 Gbps (standalone) , 800 Mbps (peered)                                                                                     |
| Max number of concurrent IP sessions | 1,000,000                                                                                                                    |
| Device management                    | SNMP, Remote Syslog, Email alerts                                                                                            |
| DHCP and DNS servers                 | DNS relay, Parallel DNS optimization, DHCP server                                                                            |
| DDNS                                 | Support for dynamic DNS for multiple interfaces                                                                              |
| WAN Configuration                    | Static IP, PPPoE, DHCP or Pass Through. Failover-only or aggregate modes for cellular and other WAN ports. Interface binding |
| Routing                              | NAT and IP forwarding, QoS and inbound/outbound VOIP quality management                                                      |

## 0.17.3 Exinda SD-WAN Management Interface

### Accessing the Exinda SD-WAN Management Interface

The Exinda SD-WAN takes approximately 60 seconds to initialize after powering up. After that, you can access the Exinda SD-WAN Management Interface. Connect an Ethernet cable from a PC to any of the Exinda SD-WAN's LAN ports and go to `http://192.168.254.99/` in a browser.

The Exinda SD-WAN Management Interface is accessible on any computer with a web browser on the local network. With the local network configured, it is not necessary to plug the computer directly into the Exinda SD-WAN in order to access the Exinda SD-WAN Management Interface.

If the Exinda SD-WAN Management Interface cannot be accessed, make sure the management computer is configured to accept a "dynamic IP" address from a DHCP server which is in this case, the DHCP server of the Exinda SD-WAN. By default, the DHCP server of the Exinda SD-WAN is enabled, but it can be disabled through the Exinda SD-WAN Management Interface. In order to access the Exinda SD-WAN Management Interface when the DHCP server is disabled and no other active DHCP is accessible, a static IP address should be assigned to the management computer. It is recommended to use the static IP address `192.168.254.10` for the management computer in this case.

The Exinda SD-WAN Management Interface presents up the following tabs to access its functionality: Home, Advanced, VLL, Firewall, Status, and Admin.

### Exinda SD-WAN licensing

Exinda SD-WAN requires a specific license to be activated.

## Entering license key

Once you log in to the Exinda SD-WAN Management Interface, enter the Activation license on the Admin tab.

For more information, refer to [Exinda SD-WAN Management Interface](#) (page 497).

One license can only be used on a single appliance.

## Connect Exinda SD-WAN to broadband modems

Exinda SD-WAN requires connection of one or more modems, such as DSL, cable, or cellular broadband. Configure the modems for operation according to the ISP's instructions and then configure the WAN ports of the Exinda SD-WAN. In many cases the modems are set for DHCP by default, and no additional configuration is required. Some cable modems may require a power-cycle (turn off and on) to associate with a new MAC address after connecting to the Exinda SD-WAN.

Note that some older modem models may require a reset to associate with a new MAC address when connecting to the Exinda SD-WAN. This may further require the modems to be reconfigured with the parameters from the ISP.

The screenshot displays the Exinda SD-WAN Management Interface with the following sections:

- Navigation:** HOME, ADVANCED, VLL, FIREWALL, PERFORMANCE, STATUS, ADMIN.
- Management Interface - Configuration for [bbna](#)**
  - Image of the Exinda SD-WAN appliance.
  - Text: **LAN Interface : connects to your LAN router/switch/hub**  
**Interface 1 - 3 : connects to your WAN devices (e.g., DSL/Cable modem)**
- LAN Connection**

| Interface           | IP Address     | Type | Download/Upload | Peak         |
|---------------------|----------------|------|-----------------|--------------|
| <a href="#">LAN</a> | 192.168.254.99 | LAN  | 4/0 Kbps        | 61/1252 Kbps |

DHCP Server :  Enable  Disable    Start IP: 192.168.254.100    End IP: 192.168.254.254    [Apply](#)
- Wired WAN**

| Interface         | IP Address    | External IP    | Status    | Type | Download/Upload | Peak         |
|-------------------|---------------|----------------|-----------|------|-----------------|--------------|
| <a href="#">1</a> | 10.170.29.144 | 212.32.226.129 | Connected | DHCP | 1/0 Kbps        | 54/1250 Kbps |
| <a href="#">2</a> | 10.170.28.168 | 212.32.226.129 | Connected | DHCP | 1/0 Kbps        | 10/2 Kbps    |
| <a href="#">3</a> | 10.170.29.143 | 212.32.226.129 | Connected | DHCP | 1/0 Kbps        | 11/2 Kbps    |
- Remote LAN (VLL)** [Add Remote LAN](#)

| Connection          | Remote Network   | Status              | Notes | Download/Upload | Peak     |
|---------------------|------------------|---------------------|-------|-----------------|----------|
| <a href="#">Con</a> | 192.168.253.0/24 | Client Initializing | None  | 0/0 Kbps        | 0/0 Kbps |
- Cellular WAN**

| Interface                     | IP Address | Signal | Status | RTT | Carrier/Type | Download/Upload | Peak |
|-------------------------------|------------|--------|--------|-----|--------------|-----------------|------|
| No cellular devices connected |            |        |        |     |              |                 |      |

Screenshot 213: Exinda SD-WAN Management Interface

In order to connect each wireline modem, use an Ethernet cable to connect the “LAN” or “Ethernet” port of the modem and connect the other end of the cable to any wired WAN port of the Exinda SD-WAN. One of the LEDs on the connector on the Exinda SD-WAN lights up after both the modem and the Exinda SD-WAN are powered on and the modem is properly connected to the Exinda SD-WAN.

If the modems do not use DHCP, the WAN connection details need to be configured for each non-DHCP modem in the Exinda SD-WAN Management Interface, as detailed below.

## Configuration of wired WAN ports

For each wired WAN port there is a corresponding status row on the Home tab. Each port is labeled with an index from 1 (WAN1) to 4 (WAN4) (or 2, 8 or 12 depending on the model) which is visible on the front panel of the Exinda SD-WAN, as well as in the left column of the tab of the Exinda SD-WAN Management Interface.

Each active wired WAN port on the Exinda SD-WAN which is connected to a modem needs to be configured with an IP address.

To configure a wired WAN port, click on the corresponding port index in the first column of the tab. If the wired WAN port is not in use, select **Disable** and click **OK**. Otherwise, select **Enable**, then specify the configuration mode for the WAN port, either DHCP, Static, PPPoE, or Pass Through, as appropriate.

After configuring each wired WAN port, click **OK** to save the corresponding WAN port settings.

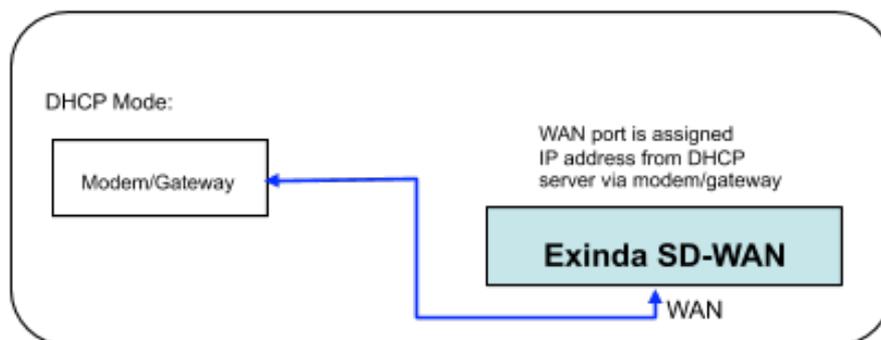
Since there is probably more than one service account being used with the Exinda SD-WAN, make sure the information entered corresponds to the correct modem and port index as determined by how the modems are connected to the Exinda SD-WAN with the Ethernet cables.

Once a WAN port is configured, the **IP Address** column for the corresponding row is filled in automatically. The “External IP Address” is the IP address that can be used to reach the corresponding WAN port from the Internet. This address usually differs from the IP address of the WAN port when a dynamic IP address is assigned via DHCP. The **Status** column corresponding to the WAN port is updated, as appropriate.

For each active WAN port, the table menu shows the rate information on each interface for monitoring the status of each Internet connection in real time.

### DHCP mode

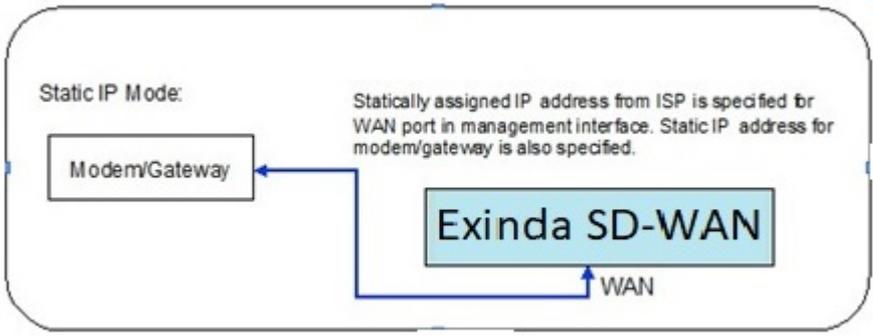
This is the default configuration mode for each wired WAN port on the Exinda SD-WAN. Typically a DSL or cable modem has a DHCP server which is capable of assigning an IP address, called a “dynamic IP address”. In this case, the DSL or cable modem automatically configures the attached wired WAN port on the Exinda SD-WAN and no manual configuration is necessary. Optionally, alternate DNS servers may be specified that override those that are provided by the modem.



Screenshot 214: The DHCP Mode

### Static IP mode

If a static IP address is assigned to a service account from the ISP, then the Static IP mode should be selected for the corresponding Exinda SD-WAN wired WAN port. The assigned static IP address needs to be entered in the User Configuration tab for the corresponding Exinda SD-WAN wired WAN port. If the ISP has assigned several static IP addresses for the same service account, one of these static IP addresses needs to be assigned to the corresponding wired WAN port on the Exinda SD-WAN.



Screenshot 215: The Static IP mode

Screenshot 216: An example of Static IP configuration for WAN interface

The **Gateway IP** address, the **Mask** of the network, and **Primary DNS** for the service account also need to be filled in; these parameters should be provided by the ISP. You can check the **Inbound Subnet Aliasing** check box if required. A field for optionally entering an IP address for a secondary DNS server is also provided.

Finally, one or more IP subnets may be specified in **Additional Subnets**. The additional subnets should be specified in CIDR notation (for example, 172.16.1.0/24) with comma delimiters. Normally, this field should be left blank, but if one or more additional subnets are specified, the designated WAN interface which is in Static IP mode is prepared to send and receive traffic to and from the specified subnet.

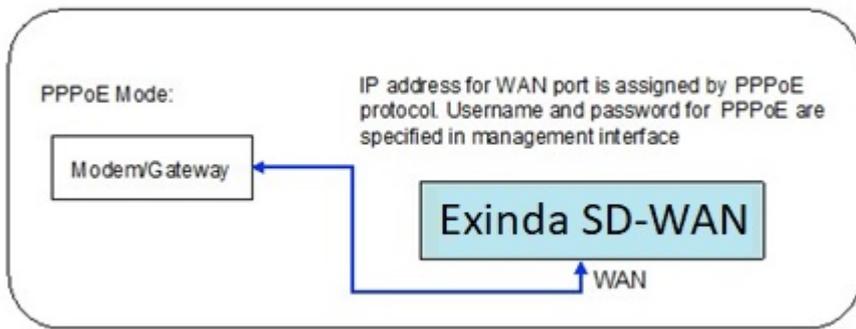
**Inbound subnet aliasing**

For each WAN port configured in the Static IP mode, there is an option to activate the Inbound Subnet Aliasing. If this option is enabled, the Exinda SD-WAN device acts as a proxy in the ARP protocol for all IP addresses in the IP address range that are specified by the given IP address and subnet mask. This may be necessary in order for inbound traffic from the Internet to reach all of the devices in the specified subnet.

**PPPoE**

Some of the older DSL modems use the PPPoE protocol to configure IP addresses. If that is the case, the PPPoE mode should be selected for the corresponding Exinda SD-WAN wired WAN port. The username and password given by the

ISP for the corresponding service account should be entered for the corresponding Exinda SD-WAN wired WAN port.



Screenshot 217: The PPPoE mode

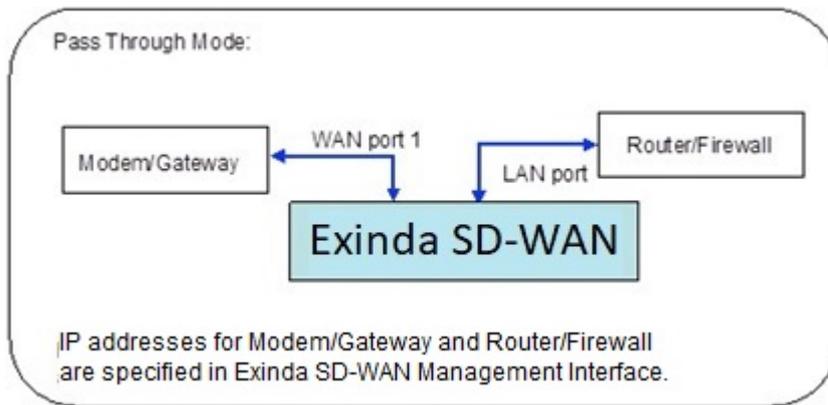
Configure WAN Interface 2  
MAC Address: 08:35:71:02:93:cf  
 Disable  
 Enable PPPoE  
User Name: abc  
Password: .....  
Service (Optional):  
Alternate DNS 1 (Optional): 8.8.8.8  
Alternate DNS 2 (Optional):  
 Standby  
Cancel OK

Screenshot 218: The PPPoE mode menu

If there is uncertainty as to what type of IP address configuration schemes are used by the modems, contact your ISP.

### Pass Through mode

When installing the Exinda SD-WAN in an existing network with a single wired WAN backhaul connection, reconfiguration of the legacy network is not required for an Exinda SD-WAN that is configured in the Pass Through mode. This mode is available on the WAN ports. This can be useful if the legacy network is relatively complex, and the network administrator does not desire to make any configuration changes to the existing network. For example, the network administrator may wish to retain all of the existing configuration parameters of the firewall device in the legacy network.



Screenshot 219: The Pass Through mode

The screenshot shows a configuration window titled "Configure WAN Interface 2". The MAC Address is "08:35:71:02:93:cf". There are two radio buttons: "Disable" (unselected) and "Enable" (selected). Next to "Enable" is a dropdown menu showing "Pass Through". Below this are several input fields: "Firewall IP" (10.10.5.2), "Gateway Address" (10.10.5.1), "Mask" (255.255.255.0), "Primary DNS" (10.10.5.1), "Secondary DNS (Optional)" (empty), "MTU" (1500), and "Additional Subnets (Optional)" (empty). At the bottom are "Cancel" and "OK" buttons.

Screenshot 220: The Pass Through mode menu

In legacy networks, a gateway or modem device that provides the WAN backhaul connection is connected directly to a router/firewall device. After installation of the Exinda SD-WAN in the Pass Through mode, the Exinda SD-WAN is inserted in-line between the gateway router/modem and the router/firewall device with the gateway/modem connected to one of the WAN ports and the router/firewall connected to one of the Exinda SD-WAN LAN ports, as indicated in the illustration above. In order to configure the Pass Through mode, the Exinda SD-WAN Management Interface prompts the user to enter the IP address of the gateway/modem to be connected to that WAN port, the IP address of the firewall, the netmask of the network behind the firewall device, and the IP address of the primary DNS server. An entry for the IP address of a secondary DNS server is optional.

Note that in the Pass Through mode, WAN port 1 is not assigned an IP address. Also, it is important to note that by configuring WAN port 1 in the Pass Through mode, the Exinda SD-WAN disables its DHCP server\*. The values entered on the Pass Through mode settings are exactly identical to the values for the firewall.

\* As a result, when the Exinda SD-WAN is configured in Pass Through mode, it may be necessary to manually assign a static IP address to the PC which is used to access the Exinda SD-WAN Management Interface, if there is no active DHCP server to assign it a dynamic IP address.

### Standby mode

Each enabled WAN port can be put in the Standby mode if desired. In this mode, a WAN port is not normally used unless no other WAN connectivity is available. This mode should be used if it is desired to use the WAN resource only as a backup in case of failure of the primary WAN connections. In order to put a WAN port in the Standby mode, select the displayed check box.

## Using a Cellular Mobile Data Card

The Exinda SD-WAN has two USB ports which may be used as WAN interfaces for connecting to a cellular data modem. Please contact [Support](#) for the list of supported air cards.

In order to use your data card, first configure it with your personal computer using the instructions from your mobile wireless service provider. Once the card is configured, you can use it with your Exinda SD-WAN. To do this, plug your wireless data card into the Exinda SD-WAN after it has booted up. After a few seconds, the Exinda SD-WAN Management Interface includes the status for Cellular WAN interface or interfaces in the table displayed at the bottom of the Home tab.

In order to configure the cellular WAN interface, on the home page of the Exinda SD-WAN Management Interface, click the mouse on the interface index on the left side of the Cellular WAN connection table. A menu is displayed for configuring the corresponding interface as CDMA or HSPA/GSM and for optionally setting the interface into standby mode. The **Carrier** and, depending on the mode, **APN** may need to be entered. The **Username** and **Password** may also be required for configuring the interface.

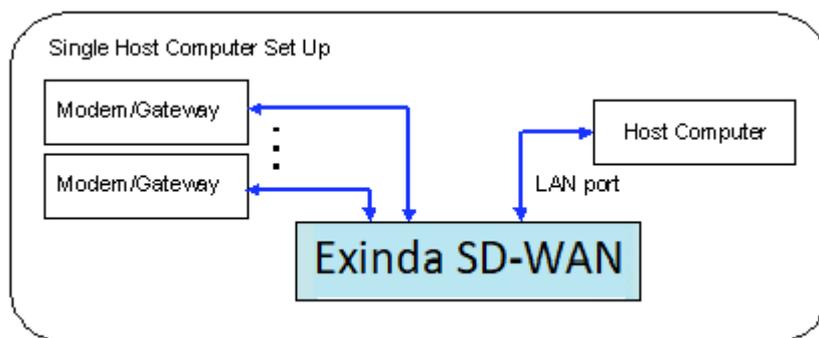
When the cellular WAN interface is in the Standby mode, it is only used when all other WAN interfaces are in the inactive state. This is useful if the cellular WAN interface is used for a service provider that charges a fee that depends on data usage. When at least one of the other interfaces enters the active state, the cellular WAN interface is not used while it is in the Standby mode. When the cellular WAN interface is not in the Standby mode, the Exinda SD-WAN uses the cellular WAN interface as much as possible in order to improve performance.

## Connect PC and local network to Exinda SD-WAN

After setting up the modems that are connected to the Exinda SD-WAN, any of the LAN ports on the Exinda SD-WAN accesses bonded Internet links. There are several ways the Exinda SD-WAN LAN ports can be used.

### As a regular Ethernet connection

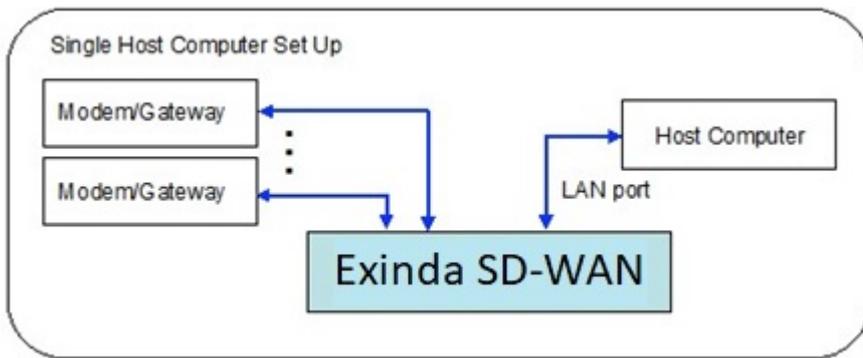
Connect an Ethernet cable to any of the LAN ports of the Exinda SD-WAN and attach the other end of the Ethernet cable to a PC or laptop which is provided bonded Internet access. Other PC or laptops can plug into the other LAN ports of the Exinda SD-WAN.



Screenshot 221: Exinda SD-WAN LAN ports used as a regular Ethernet connection

### By using an Ethernet Switch

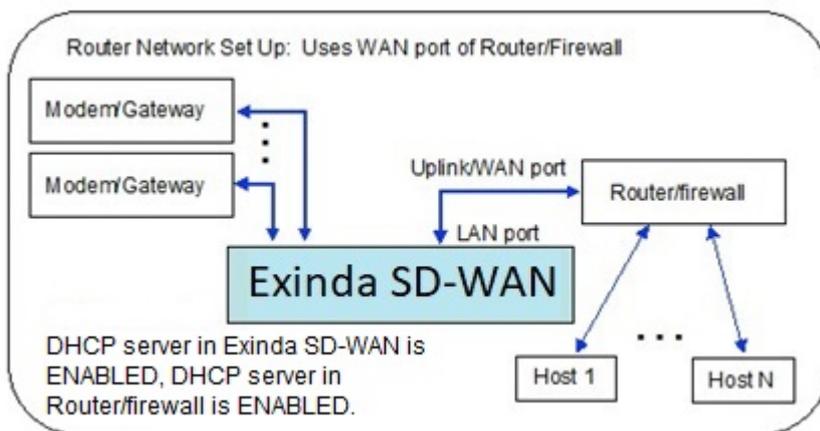
If the number of computer devices (such as PCs and laptops) to be supported is greater than the number of LAN ports on the Exinda SD-WAN (two), an Ethernet switch can be connected to one of the LAN ports of the Exinda SD-WAN. Any port of the Ethernet switch may also be used to provide bonded Internet access to a local device.



Screenshot 222: Exinda SD-WAN LAN ports used with an Ethernet switch

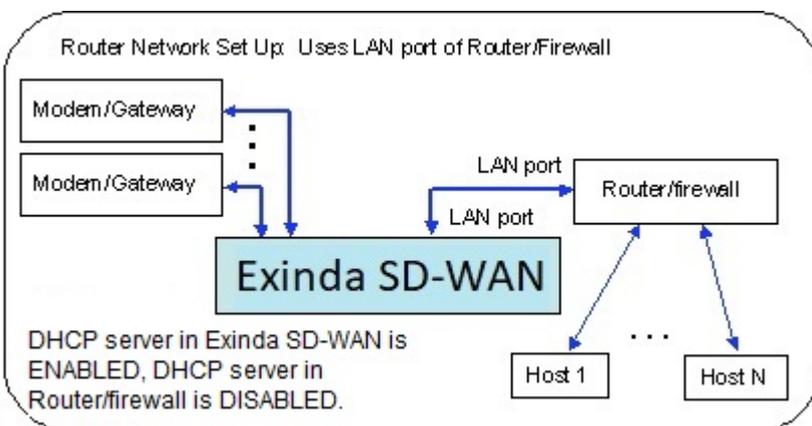
### With a Router

Connect the WAN/Internet port of the router to one of the Exinda SD-WAN LAN ports. The DHCP server in the Exinda SD-WAN should be enabled to assign the router/firewall an IP address. The DHCP server in the router/firewall should also be enabled for the host computer devices which are attached to the router/firewall through the LAN port of the router/firewall. It may be necessary to configure the DHCP server of the router/firewall to set the DNS server address for its clients to the IP address of the Exinda SD-WAN (by default this is set to 192 . 168 . 254 . 99) rather than the IP address of an external DNS server.



Screenshot 223: Exinda SD-WAN LAN ports used with a router WAN port

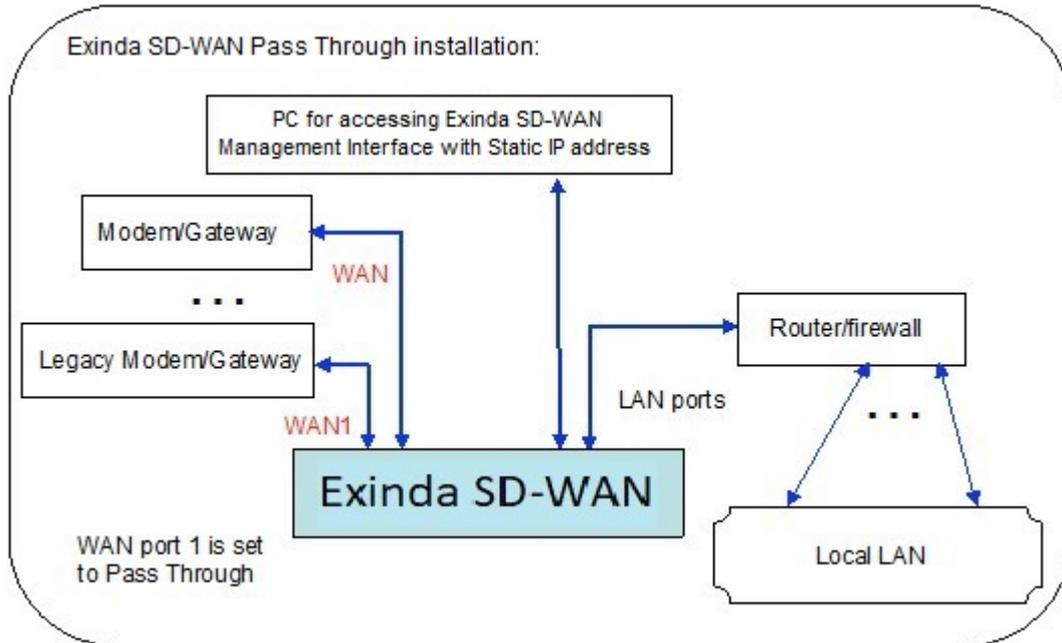
Alternatively, one of the Exinda SD-WAN LAN ports can be connected to a LAN port on the router. In this case, the DHCP server needs to be disabled on the router, so that it does not conflict with that of the Exinda SD-WAN. With this configuration, UPnP support is available.



Screenshot 224: Exinda SD-WAN LAN ports used with a router LAN port

## Configuring the Pass Through mode

The Exinda SD-WAN device is inserted in-line between the legacy modem/gateway and the router/firewall. The legacy modem/gateway is connected to WAN port 1 of the Exinda SD-WAN and the router/firewall is connected to the LAN port of the Exinda SD-WAN. Additional modem/gateways for the network are connected to other WAN ports of the Exinda SD-WAN. The following illustration shows the network above after installation of the Exinda SD-WAN.



Screenshot 225: Exinda SD-WAN Pass Through configuration

The Exinda SD-WAN can first be configured by plugging a PC into one of the LAN ports of the Exinda SD-WAN, as shown above, and accessing the Exinda SD-WAN Management Interface. After the Exinda SD-WAN is appropriately configured, the Exinda SD-WAN Management Interface can be accessed by any PC on the Local LAN.

Note that when the Exinda SD-WAN is configured in the Pass Through mode, the DHCP server within the Exinda SD-WAN is disabled. Therefore, it is generally necessary to manually assign a static IP address to the PC through which the Exinda SD-WAN Management Interface is accessed.

## Advanced features

Here you can find the advanced features of the Exinda SD-WAN.

### IP address for Exinda SD-WAN LAN

By default, the LAN interface of the Exinda SD-WAN is assigned the IP address 192 . 168 . 254 . 99. This is the IP address through which the Exinda SD-WAN Management Interface is accessed.

It is possible to change the IP address of the Exinda SD-WAN LAN interface. This should be done with extreme care, as changing this affects the manner in which the Exinda SD-WAN Management Interface is accessed. In particular, it is important to remember the IP address that is assigned to the Exinda SD-WAN LAN. It is not possible to access the Exinda SD-WAN Management Interface without this information.

To set the Exinda SD-WAN LAN IP address, click the Home tab of the Exinda SD-WAN Management Interface. In the LAN configuration table, in the Interface column, click **LAN**. A prompt appears requesting for an IP address. If the value is changed for this IP address, it is imperative that the new value be remembered. The network mask for the Exinda SD-WAN LAN can also be set under **Mask**. By default, the network mask is set to 255 . 255 . 255 . 0.

## DHCP server

The DHCP server on the Exinda SD-WAN can be enabled or disabled via the radio buttons on the Home tab in the LAN configuration table. The “starting IP address” for the DHCP server can also be configured here. When the DHCP server inside the Exinda SD-WAN is enabled, it starts assigning dynamic IP addresses starting from the configured address. This may be useful in some cases to avoid conflicts with manually assigned static IP addresses. The default starting IP address is 192.168.254.100. An ending IP address can also be specified. The default ending IP address is 192.168.254.254.

The DHCP server is enabled by default. It is automatically disabled when the Pass Through mode is enabled on any of the WAN ports.

The specified IP address range for DHCP is also the range of local IP addresses that can be reached by VPN clients that are terminated on the Exinda SD-WAN.

The DHCP lease time and maximum number of leases can be set from the “DHCP Server Additional Settings” section in the Advanced tab. The gateway IP and WINS Server IP may also be set there. The name of the server IP for internal domains may be specified as well.

## Support for PPTP

The Exinda SD-WAN is able to terminate VPN connections by using the PPTP protocol. By default, this feature is disabled.

To configure this feature, select the Advanced tab within the Exinda SD-WAN Management Interface and click on the entry for PPTP in the table (the default entry is **Disabled**). A pop-up window appears, which allows this feature to be configured. To enable the feature, select **Enable** and enter a name and password pair, and click **Apply**.

After the feature has been enabled, PPTP clients may access the LAN by using the same address range as DHCP with the specified name and password.

## Support for UPnP

By default, the Exinda SD-WAN is configured to support UPnP which enables auto-configuration of IP addressing features for certain applications.

Normally, support for UPnP should be disabled when running the Exinda SD-WAN in the Pass Through mode since this functionality would already be provided by a device in the legacy network.

## Inbound port forwarding

Inbound port forwarding allows external computers on the internet to access devices in the local network through the same IP address. Normally this feature is used with static IP address assignment, but it can be used with any of the IP address configuration modes for the Exinda SD-WAN WAN ports. Inbound port forwarding is not recommended for the Pass Through mode of operation on WAN interface 1, since this function, if required, would already be handled by the legacy network that was kept intact without any changes during the installation.

By default, all inbound Internet traffic is blocked, except for ICMP traffic. Selected inbound internet traffic can be allowed or blocked by configuration of one or more forwarding rules. To configure a rule for forwarding inbound Internet traffic, select the Firewall tab, and click **Add Firewall Rule**. Select **Inbound** as a **Direction**.

In order to configure a rule for forwarding inbound Internet traffic, an **Interface** must also be specified on the pop-up menu. You can specify **ALL WAN** which means that traffic from all WAN interfaces, wired or cellular is the target of the rule. Alternatively, a forwarding rule for a single WAN interface can be specified in the drop-down menu. In order for cellular WAN interfaces to become available, the corresponding cellular WAN interface should be configured first. In addition, if any remote LAN connections have been configured (with the optional VLL feature), they also appear in the menu as possible targets of the forwarding rule.

Traffic from a specific protocol or from all protocols can be specified as the target of the forwarding rule by selecting under **Protocol/Action**. **TCP**, **UDP**, **ICMP**, **IPSec**, and **PPTP** are available, as is **Any** which is interpreted as all protocols. A specific protocol other than the above can also be specified by selecting **Other** and entering the corresponding protocol number.

Adding an inbound forwarding rule causes the targeted traffic to be forwarded to the **Local IP** address designated in the pop-up menu, unless the **Discard** option is selected in which case the targeted traffic is discarded.

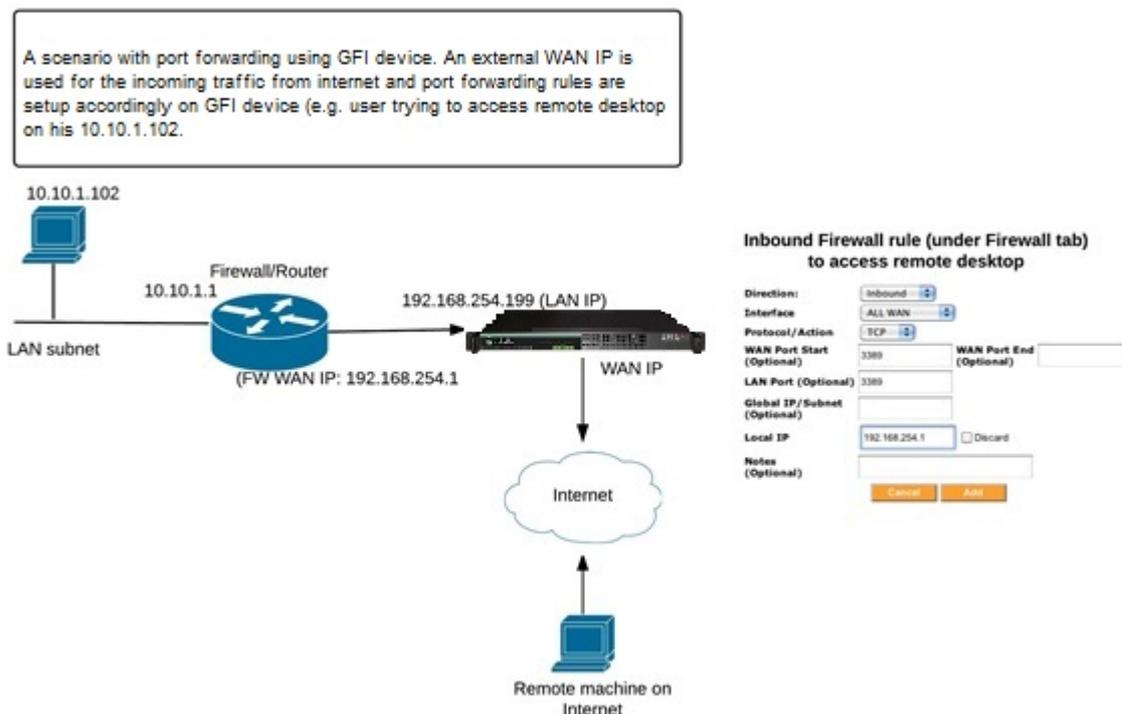
Specifying a WAN port number causes all traffic destined to the specified port to be forwarded to the specified Local IP address. More generally, a range of WAN ports can also be specified by designating a second WAN port, which is the end of the port range. It is delivered to the same destination port unless a LAN port is specified, in which case it is delivered to the specified LAN port.

A **Global IP** address may also be optionally specified, which signifies that only that traffic which is destined to the specified Global IP address should be forwarded to the specified Local IP address.

A text string can also be entered in **Notes**, to provide a mechanism for the operator to remember the reason or context of the forwarding rule.

After specifying all desired conditions in a forwarding rule for inbound traffic, click **Add** to load the forwarding rule into the Exinda SD-WAN. Multiple forwarding rules can be added to provide more flexible forwarding of traffic. If a forwarding rule that was added previously needs to be deleted, click **Delete** next to the listed forwarding rule. If a forwarding rule needs to be edited, the rule can first be deleted and the modified rule can then be entered as before. The Exinda SD-WAN processes arriving traffic by sequential application of the specified forwarding rules. This means if there are conflicts in the forwarding rules, the forwarding rule nearest the top of the list takes precedence. By default, the list of forwarding rules is such that rules entered earlier is lower in the list, so the most recently added rule always take precedence.

Here is an example scenario for a user accessing his desktop remotely:



Screenshot 226: Example scenario for a user accessing his desktop remotely

## Outbound port blocking and NAT

Outbound port blocking allows a network operator to block all outgoing traffic from a specified source port, protocol, and optionally a specified source IP address or subnet.

By default, all outbound internet traffic is allowed. To configure a rule for blocking outbound Internet traffic, click the **Firewall** tab to display the menu and click **Add Firewall Rule**. Select **Outbound** as a **Direction**. Use the menu under **Protocol/Action** to select which protocol should be denied outbound access.

Alternatively, you can specify an arbitrary protocol by selecting **Other** and entering the corresponding protocol number. You can also select **Deny Any** to deny all outbound traffic based on protocol. Finally, you can select the **No NAT** option if you wish to disable the Network Address Translation (NAT) function on the specified outbound traffic. By default, NAT is applied to all outgoing traffic, meaning that the source IP address of egress packets that originate on the Exinda SD-WAN LAN is replaced with the IP address of the corresponding WAN interface, and selecting the No NAT option disables the NAT function for the targeted traffic.

Optionally, a Source Port index for the targeted traffic can be specified. Thus, only packets with a Source Port value matching the entered value is targeted. Also optionally, a source IP address can be specified so that only traffic with a source IP address matching the designated address is targeted. Generally, an IP subnet can also be specified in CIDR notation, for example, 192 . 168 . 1 . 0 / 24. In this case, traffic originating from the specified subnet is targeted. The targeted traffic by definition satisfies all the matching conditions, so that effectively the conditions are logically concatenated to specify the targeted traffic.

A text string can also be entered in **Notes**, to provide a mechanism for the operator to remember the reason or context of the outbound rule.

Multiple outbound blocking rules can be added to provide more flexible blocking of traffic. If a traffic blocking rule that was added previously needs to be deleted, click **Delete** next to the listed blocking rule. If a blocking rule needs to be edited, the rule can first be deleted and the modified rule can then be entered as before. The Exinda SD-WAN processes outgoing traffic by sequential application of the outbound rules specified. This means if there are conflicts in the outbound rules, the outbound rule nearest the top of the list takes precedence. By default, the list of outbound rules is such that rules entered earlier is lower in the list, so the most recently added rule always takes precedence.

Normally outbound blocking rules should not be specified when the Exinda SD-WAN is configured in Pass Through mode since such blocking would already be done by the legacy network.

### Advanced firewall rules for filtering

Firewall Rules 

| Field                 | Value  |
|-----------------------|--------|
| Type                  | Drop ▼ |
| Priority*             | 0      |
| Incoming Device Type  | - ▼    |
| Incoming Device Index |        |
| Outgoing Device Type  | - ▼    |
| Outgoing Device Index |        |
| Source                |        |
| Destination           |        |
| DSCP                  | - ▼    |
| Protocol              | - ▼    |
| Source Ports          |        |
| Destination Ports     |        |
| Notes                 |        |

Apply

Screenshot 227: The Firewall Rules configuration screen

The advanced firewall rules allow characterizing traffic based on the Incoming Device Type and the Outgoing Device Type, or both. The “Type” field is the action that the rule takes on the traffic that is matched. The type can be either

ACCEPT, DROP or REJECT. ACCEPT takes the matched traffic, DROP drops the traffic and REJECT would send an ICMP unreachable packet to the sender and then DROP the packet.

The traffic can be filtered based on the Source IP, destination IP, source ports or destination ports. An IP subnet also can be specified with a CIDR notation, e.g. 192.168.1.0/24.

Multiple firewall rules can be added to achieve a combination by means of priority. The priority is in ascending order, so the rule with a lower priority is processed first. So to accept a particular IP address and drop all other traffic, we have the ACCEPT rule with a lower priority and the DROP rule with a higher priority.

## DMZ support

Support for a DMZ (demilitarized zone) can be enabled through the Firewall tab within the Exinda SD-WAN Management Interface.

The DMZ feature of the Exinda SD-WAN allows the user to specify the IP address of a local computer that has been designated as the DMZ. The specified IP address is called the "DMZ IP address". When the Exinda SD-WAN receives a packet from the internet over any of the WAN ports then the Exinda SD-WAN forwards the packet to the specified DMZ address if none of the port forwarding rules apply to the given packet.

If no port forwarding rules are specified and a DMZ IP address is specified, then all packets received on a WAN port is forwarded to the local computer with the specified DMZ IP address. It should be noted that if one or more port forwarding rules are specified that they take precedence over the forwarding of traffic to a DMZ, if the DMZ feature is enabled. Moreover, it should be noted that any port forwarding that occurs as a result of the UPnP feature being enabled takes precedence over any port forwarding rules.

Setting up a DMZ is useful if there are several computers in the local network, but only one of them (that has been designated as the DMZ) is desired to be directly accessible via the Internet. Normally, support for DMZ should be disabled when running the Exinda SD-WAN in the Pass Through mode since a DMZ would already be implemented in the legacy network.

The DMZ feature can be configured by clicking on the status **Enabled** or **Disabled** for DMZ in the configuration table and selecting the appropriate radio button. If **Enable** is selected, then the IP address for the DMZ must be specified. When you click **Apply**, the DMZ feature is configured accordingly.

## Dynamic DNS

The Exinda SD-WAN supports inbound load balancing and inbound link fail-over via Dynamic DNS (DDNS). This feature allows management of inbound traffic (traffic which is originally initiated from the Internet). This feature is configured via the Advanced tab. It is disabled by default and three modes are supported currently through the Dynamic DNS service at afraid.org, dyndns.org, and changeip.com. To properly configure this feature, a service account should first be set up with afraid.org, dyndns.org, or changeip.com. Details of the service are available from these service providers. Once the service account has been set up the corresponding mode can be selected by clicking on the status entry ("Disabled" or "Enabled") for "Dynamic DNS" in the configuration table. The various options is displayed, and the Exinda SD-WAN Management Interface then prompts for a Username, Password, and Hostname, which should be provided by the Dynamic DNS provider. By clicking Apply, the Dynamic DNS feature is configured accordingly.

For afraid.org's round-robin DDNS to work correctly, it requires unlinking of entries for the same domain (one entry for each Exinda SD-WAN WAN connection is required). For more information, see the [afraid.org](http://freedns.afraid.org/) FAQ at <http://freedns.afraid.org/>.

DDNS is a service that is hosted by companies with no connection to GFI. GFI makes no warranty as to the reliability of these services. For mission critical applications, it is strongly advised that the reliability of the service be used considered before deployment.

Support for dynamic DNS on multiple interfaces is also available.

Multiple DYNDNS setup 

| Field            | Value                                                                                                      |
|------------------|------------------------------------------------------------------------------------------------------------|
| Dyndns Provider* | dyndns.org ▼                                                                                               |
| Device Type*     | wan ▼                                                                                                      |
| Device Index     | <input type="text"/>      |
| DYNDNS Username* | <input type="text"/>                                                                                       |
| DYNDNS Password* | <input type="password"/>  |
| DYNDNS Hostname* | <input type="text"/>                                                                                       |

**Apply**

Screenshot 228: Multiple DYNDNS Setup screen

### Load balancing weights

Weights may be set per interface for load balancing. These weights adjust the proportion of traffic that is handled by each WAN interface.

To configure these, click the **Advanced** tab and then click **Add** next to **Load Balancing Weights**.

For setting the weights per interface, it is recommended that you use the rounded value of the kbit/s rate of the respective WAN interfaces.

### Connectivity check settings

The settings for connectivity check may be adjusted on the **Advanced** tab of the Exinda SD-WAN Management Interface. These settings determine how the Exinda SD-WAN discovers and monitors the connectivity of the individual WAN interfaces.

The Exinda SD-WAN attempts to reach external servers to determine if a WAN interface is successfully connected. The connectivity check proceeds in three phases:

1. Resolving connectivity check addresses
2. Pinging connectivity check addresses
3. Simultaneously pinging and attempting HTTP connections to connectivity check addresses

The maximum duration (in seconds) of each phase may be adjusted by clicking **Settings** in the **Index** column of the **Connectivity Check Settings** configuration table. If desired, to immediately check HTTP in addition to ICMP, the time for the ping phase may be reduced to zero. Be aware, though, of the bandwidth overhead of retrieving HTTP responses on very limited lines.

The Exinda SD-WAN internally manages a list of active server IPs that is used for these checks, which are automatically updated as long as the Exinda SD-WAN maintains a connection to the internet. Additional servers may be added manually by clicking **Add** next to **Additional Remote Servers for Connectivity Check** on the **Advanced** tab and entering the server's IP address. These servers should be responsive to ICMP or HTTP requests, or both, and be available with high reliability.

### List of DHCP leases

A list of all dynamic IP addresses leased out by the Exinda SD-WAN DHCP server can be viewed on the **Status** tab. The MAC address, IP address, and possibly the machine name are listed for each IP address lease.

### Remote syslog

A log of Exinda SD-WAN system events appears in the **Status** tab. The IP address of a Remote Syslog Server can be specified which is to be the recipient of log messages generated by the Exinda SD-WAN. After the IP address of the

Remote Syslog Server has been specified, click Apply to configure this feature.

## Domain Name Blocking / DNS Redirection

| Field                     | Value                |
|---------------------------|----------------------|
| Domain Name*              | <input type="text"/> |
| IP (leave blank to block) | <input type="text"/> |

Apply

Screenshot 229: The Domain Name Blocking / DNS Redirection Menu

Domain name blocking / DNS redirection is used to block a site or used to redirect to a particular IP for a site. For the DNS redirection to work, the Exinda SD-WAN device should be configured the same as the DNS server in the local network.

## VLAN setup

| Field         | Value                |
|---------------|----------------------|
| Device Type*  | wan                  |
| Device Index* | <input type="text"/> |
| VLAN Tag*     | <input type="text"/> |
| Bridge Index  | <input type="text"/> |

Apply

Screenshot 230: The VLAN Setup menu

The VLAN setup can be used for creating 802.11q format frames with a particular tag. The "Device Type" would be WAN, LAN, or Cellular, followed by the "Device Index" which would be the index number. For WAN 1, you would enter '1' here. The "VLAN Tag" field contains the tag number, and "Bridge Index" contains the bridge index.

## Netflow settings

| Field                | Value                |
|----------------------|----------------------|
| Netflow Server IP*   | <input type="text"/> |
| Netflow Server Port* | <input type="text"/> |
| Netflow Version      | <input type="text"/> |

Apply

Screenshot 231: The Netflow Settings menu

The netflow settings are used to configure sending netflow data to a netflow server on the local area network of the Exinda device. The required parameters are the **Netflow Server IP** and the **Netflow Server Port**. The Netflow Version supports versions 1, 5, and 9.

## Manual network routes

Manual routes are used in scenarios where you have multiple subnets on the LAN segment, or you want to specify certain subnet traffic to use a specific Remote LAN, or WAN, link. This feature can be configured by going to the Advanced tab on the Exinda SD-WAN Management Interface and clicking Add next to the Manual Network Routes

section:

**Manual Network Routes** ?

| Field                | Value                                                                                                  |
|----------------------|--------------------------------------------------------------------------------------------------------|
| Type                 | Simple ▾                                                                                               |
| Device Type*         | wan ▾                                                                                                  |
| Device Index*        | <input type="text"/>  |
| Notes                | <input type="text"/>                                                                                   |
| Use Default Gateway* | True ▾                                                                                                 |
| Gateway              | <input type="text"/>                                                                                   |
| Destination          | <input type="text"/>                                                                                   |
| MTU                  | <input type="text"/>                                                                                   |
| Metric*              | 0                                                                                                      |

**Apply**

Screenshot 232: Manual Network Routes, Simple

This generates a pop-up window where the information specifying a manual route rule can be entered. The **Type** field specifies **Simple Route** or **Advanced Route**. **Device Type** identifies the interface/tunnel type for the traffic we are setting the route for. As an example, WAN or LAN subnets or VLL if the traffic is destined for a VLL tunnel. The **Destination** field specifies the destination IP address subnet to which the packets are routed to. This can be accomplished either through a device route or through a VLL route, based on the option chosen previously.

Choosing the **Advanced Route** as the **Type** exposes additional configuration options to identify the traffic including **Source IP/Subnet**, **DSCP**, **Protocol**, **Source Ports**, and **Destination Ports**. You can also modify the **Priority** and the Maximum Transmission Unit (**MTU**) size of the traffic.

The failover mode is used to configure the failover action in case of primary interface failure. It has 2 modes: "sticky mode" and "available mode". In case of a primary interface failure, the sticky mode transfers the traffic on to the secondary interface, and when the primary line comes back up, transfers the traffic back to the primary interface. With the "available mode", after the primary interface recovers, the traffic that failed- over earlier continues to use the failover interface. The newer sessions is routed over the primary interface as usual.

Manual Network Routes ? Remove

| Field                | Value      |
|----------------------|------------|
| Type                 | Advanced ▼ |
| Device Type*         | wan ▼      |
| Device Index*        | 1          |
| Notes                |            |
| Use Default Gateway* | True ▼     |
| Gateway              |            |
| Destination          |            |
| MTU                  |            |
| Priority*            | 0          |
| Source               | 1.2.3.4    |
| DSCP                 | - ▼        |
| Protocol             | - ▼        |
| Source Ports         |            |
| Destination Ports    |            |
| Failover Mode        | - ▼        |

Apply

Screenshot 233: Manual Network Routes, Advanced

## Interface groups

An "interface group" is an abstraction of a set of real interfaces that the user can manage with finer control. The user can define what interfaces (WAN, Cellular, VLL) are to be associated with a particular interface group, what percentage of each interface is to be used and if the group requires failover interfaces. The user can also define what type of traffic would traverse this bundled virtual interface.

## Components of interface groups

### Interface groups

This block gives the user the functionality to create a group which could then have filters and the outgoing interfaces associated with it.

Interface Groups ? Remove

| Field | Value   |
|-------|---------|
| Type  | Hash ▼  |
| Notes | Group 1 |

Apply

Screenshot 234: Adding interface groups

## Interface group interfaces

Once an interface group is created, the group's specific interfaces can be configured and added to the interface group. For instance the user can add WAN 1 and Cellular WAN 2 to a particular interface group.



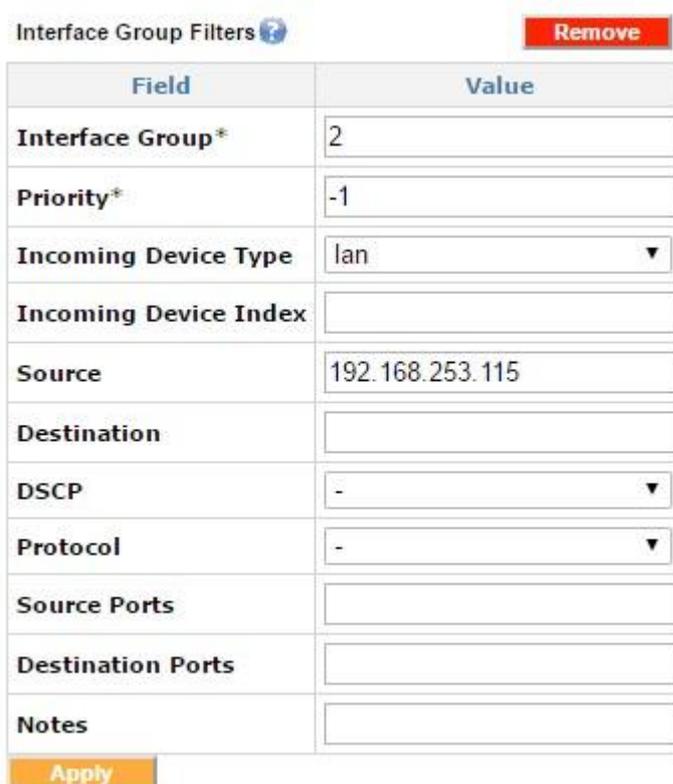
| Index | Interface Group | Device Type | Device Index | Weight |
|-------|-----------------|-------------|--------------|--------|
| 1     | 1               | wan         | 1            | 1.0    |
| 2     | 1               | cellular    | 2            | 1.0    |

Screenshot 235: Interface group interfaces

In this example only WAN1 and Cellular WAN2 is included in this group.

## Interface group filter

The interface group filter is used to filter the desired traffic types on a particular Interface Group.



| Field                 | Value           |
|-----------------------|-----------------|
| Interface Group*      | 2               |
| Priority*             | -1              |
| Incoming Device Type  | lan             |
| Incoming Device Index |                 |
| Source                | 192.168.253.115 |
| Destination           |                 |
| DSCP                  | -               |
| Protocol              | -               |
| Source Ports          |                 |
| Destination Ports     |                 |
| Notes                 |                 |

Screenshot 236: Adding interface group filters

This sample filter directs all traffic from 192 . 1 6 8 . 2 5 3 . 1 1 5 onto Interface Group 2.

The default Interface Group with index 1 should not be modified. Any additional desired Interface Groups should be added as new Interface Groups.

| Interface Groups <span>?</span> |      |       | Add |  |  |  |  |  |  |
|---------------------------------|------|-------|-----|--|--|--|--|--|--|
| Index                           | Type | Notes |     |  |  |  |  |  |  |
| <u>1</u>                        | Hash |       |     |  |  |  |  |  |  |

| Interface Group Interfaces <span>?</span> |                 |             |              |        | Add |  |  |  |  |
|-------------------------------------------|-----------------|-------------|--------------|--------|-----|--|--|--|--|
| Index                                     | Interface Group | Device Type | Device Index | Weight |     |  |  |  |  |
| <u>1</u>                                  | 1               | wan         | -            | 0.0    |     |  |  |  |  |
| <u>2</u>                                  | 1               | cellular    | -            | 0.0    |     |  |  |  |  |

| Interface Group Filters <span>?</span> |                 |                      |                       |        |             |      |          |              |                   | Add |
|----------------------------------------|-----------------|----------------------|-----------------------|--------|-------------|------|----------|--------------|-------------------|-----|
| Index                                  | Interface Group | Incoming Device Type | Incoming Device Index | Source | Destination | DSCP | Protocol | Source Ports | Destination Ports |     |
| <u>1</u>                               | 1               | lan                  | -                     | -      | -           | -    | -        | -            | -                 |     |

Screenshot 237: Interface Groups

The blank **Device Index** in this example indicates that all device type indexes should be selected.

The distribution of traffic within this interface group can be customized by using **Weights**. The “weight” indicates the proportion of traffic desired for a particular interface. A weight of 0 . 0 for all interfaces selected indicates an equalized distribution of the traffic. This would match all the incoming traffic on the LAN. These weights can be customized to load balance by changing the weights as desired.

### Enable Capture All with group interfaces

| Interface Groups <span>?</span> |      |       | Add |  |  |  |  |  |  |
|---------------------------------|------|-------|-----|--|--|--|--|--|--|
| Index                           | Type | Notes |     |  |  |  |  |  |  |
| <u>1</u>                        | Hash |       |     |  |  |  |  |  |  |

| Interface Group Interfaces <span>?</span> |                 |             |              |        | Add |  |  |  |  |
|-------------------------------------------|-----------------|-------------|--------------|--------|-----|--|--|--|--|
| Index                                     | Interface Group | Device Type | Device Index | Weight |     |  |  |  |  |
| <u>1</u>                                  | 1               | wan         | -            | 0.0    |     |  |  |  |  |
| <u>2</u>                                  | 1               | cellular    | -            | 0.0    |     |  |  |  |  |

| Interface Group Filters <span>?</span> |                 |                      |                       |        |             |      |          |              |                   | Add |
|----------------------------------------|-----------------|----------------------|-----------------------|--------|-------------|------|----------|--------------|-------------------|-----|
| Index                                  | Interface Group | Incoming Device Type | Incoming Device Index | Source | Destination | DSCP | Protocol | Source Ports | Destination Ports |     |
| <u>1</u>                               | 1               | lan                  | -                     | -      | -           | -    | -        | -            | -                 |     |

Screenshot 238: Enabling Capture All with interface groups

In this example, the weight of 1 . 0 for the VLL tunnel ensures all the traffic goes over the VLL interface. The other interfaces have weights of 0 . 0, which makes them a failover interface and only activate them in case the primary interface (in this example, the VLL interface) goes down.

## Load balancing by using the interface groups

| Interface Groups <span>?</span> |  |  |      |  |  |  |  |  |  | Add   |
|---------------------------------|--|--|------|--|--|--|--|--|--|-------|
| Index                           |  |  | Type |  |  |  |  |  |  | Notes |
| <u>1</u>                        |  |  | Hash |  |  |  |  |  |  |       |

| Interface Group Interfaces <span>?</span> |                 |             |              |        | Add |
|-------------------------------------------|-----------------|-------------|--------------|--------|-----|
| Index                                     | Interface Group | Device Type | Device Index | Weight |     |
| <u>1</u>                                  | 1               | wan         | 1            | 0.7    |     |
| <u>2</u>                                  | 1               | wan         | 2            | 0.3    |     |

| Interface Group Filters <span>?</span> |                 |                      |                       |        |             |      |          |              |                   | Add |
|----------------------------------------|-----------------|----------------------|-----------------------|--------|-------------|------|----------|--------------|-------------------|-----|
| Index                                  | Interface Group | Incoming Device Type | Incoming Device Index | Source | Destination | DSCP | Protocol | Source Ports | Destination Ports |     |
| <u>1</u>                               | 1               | lan                  | -                     | -      | -           | -    | -        | -            | -                 |     |

Screenshot 239: Load balancing with interface groups

Load balancing weights can be defined for each interface contained in the interface group interfaces. These weights indicate the proportion of traffic desired for the given interface.

In this example, the weights are set to 0.7 for WAN1 and 0.3 for WAN2, which results in 70% of the traffic being directed to WAN1 and 30% of the traffic being directed to WAN2.

## Random group

| Interface Groups <span>?</span> |  |  |        |  |  |  |  |  |  | Add   |
|---------------------------------|--|--|--------|--|--|--|--|--|--|-------|
| Index                           |  |  | Type   |  |  |  |  |  |  | Notes |
| <u>1</u>                        |  |  | Random |  |  |  |  |  |  |       |

| Interface Group Interfaces <span>?</span> |                 |             |              |        | Add |
|-------------------------------------------|-----------------|-------------|--------------|--------|-----|
| Index                                     | Interface Group | Device Type | Device Index | Weight |     |
| <u>2</u>                                  | 1               | wan         | -            | 1.0    |     |

| Interface Group Filters <span>?</span> |                 |                      |                       |        |             |      |          |              |                   | Add |
|----------------------------------------|-----------------|----------------------|-----------------------|--------|-------------|------|----------|--------------|-------------------|-----|
| Index                                  | Interface Group | Incoming Device Type | Incoming Device Index | Source | Destination | DSCP | Protocol | Source Ports | Destination Ports |     |
| <u>1</u>                               | 1               | lan                  | -                     | -      | -           | -    | -        | -            | -                 |     |

Screenshot 240: Random interface group

In this example, the traffic matching the filter is randomly distributed across all interfaces that are included in the interface group interfaces for that group. Note that since the **Device Index** is left blank, then all WAN interfaces are included in this group. This load-balancing algorithm is essentially a round-robin algorithm.

Advanced routes take precedence over the interface group filters, and therefore advanced routes can be used to create exceptions to the interface group filters.

## VLAN tagging

**VLAN setup** ⓘ

| Field         | Value                |
|---------------|----------------------|
| Device Type*  | wan ▼                |
| Device Index* | <input type="text"/> |
| VLAN Tag*     | <input type="text"/> |
| Bridge Index  | <input type="text"/> |

**Apply**

Screenshot 241: VLAN tagging

The VLAN tags can be set for both WAN and LAN interfaces. The **Device Type** can be WAN or LAN, the **Device Index** here is the index number. The **VLAN Tag** field contains the tag number.

## Bridge alias

**Bridge Aliases** ⓘ

| Field  | Value                |
|--------|----------------------|
| Subnet | <input type="text"/> |

**Apply**

Screenshot 242: A Bridge alias

The Bridge alias is used to add a subnet alias to the LAN interface. This feature can be useful if there is more than one subnet in the local network.

## QoS (Quality of Service)

The QoS feature can be configured through the **Advanced** tab on the Exinda SD-WAN Management Interface. This feature is used to guarantee a certain bandwidth for an application or to prioritize certain traffic, such as Citrix or SIP.

## Quality of Service shapers

In order to use the QoS feature, a user must first create the WAN Shaper rule (**Quality of Service Shaper**) on the **Advanced** tab. A rate that is 85% of the rate provided by the ISP should be used as the shaper rule both for uplink (egress) and downlink (ingress).

Quality of Service Shapers

| Field                                   | Value      |
|-----------------------------------------|------------|
| Type                                    | QoS Shaper |
| Device Type*                            | wan        |
| Device Index*                           | 1          |
| Direction*                              | egress     |
| Active                                  |            |
| Rate (Kbit)*                            | 600        |
| Burst (bytes)                           |            |
| Unclassified Rate (Kbit or percentage)* | 10         |
| Unclassified Buffer (packets)*          | 5          |
| Unclassified Burst (bytes)              |            |

Apply

WAN 1  
uplink  
uplink rate

Screenshot 243: Quality of Service shaper example

In this example, the ISP uplink rate is 705 kbps, so the egress shaper for WAN1 is set to 85% of the available bandwidth, or 600 Kbps.

### Quality of Service reservations

After creating the shaper rule for both ingress and egress, the application traffic reservation and prioritization can be created using **Quality of Service Reservations**. The following example illustrates how to reserve 200 Kbps uplink for IPSec traffic:

|                                       |                 |
|---------------------------------------|-----------------|
| Type                                  | QoS Reservation |
| Device Type*                          | wan             |
| Device Index*                         | 1               |
| Active                                |                 |
| Notes                                 |                 |
| Guaranteed Rate (Kbit or percentage)* | 200             |
| Max Rate (Kbit or percentage)         |                 |
| Buffer Limit (packets)*               | 5               |
| Burst (bytes)                         |                 |
| Priority*                             | Normal Priority |
| Direction*                            | egress          |
| Source IP/netmask                     |                 |
| Destination IP/netmask                |                 |
| Protocol                              | esp             |
| Source Port                           |                 |
| Destination Port                      |                 |

Apply

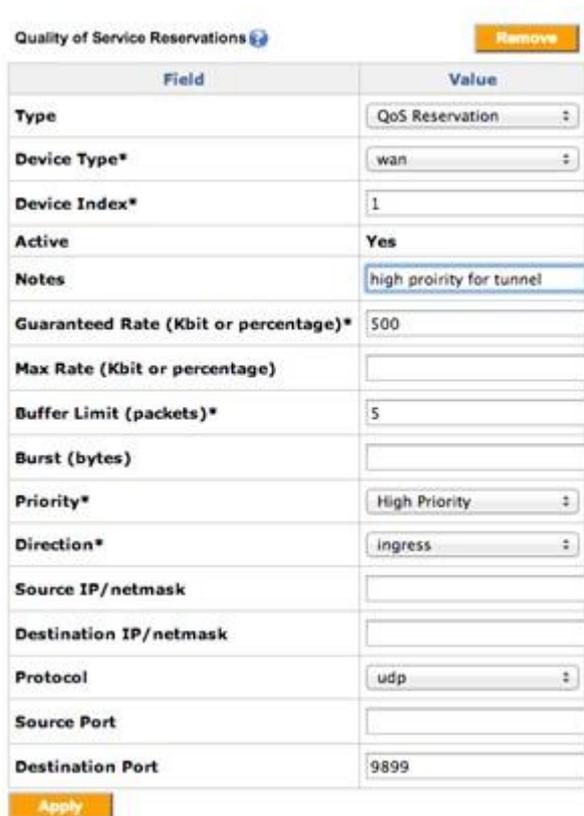
WAN1  
Reserved rate  
Outgoing traffic  
IPSec traffic

Screenshot 244: Quality of Service reservation example

For prioritizing real time traffic, the **Priority** field can be changed from **Normal Priority** to **High Priority**. Advanced application specific QoS can be achieved with Exinda Network Orchestrator. For more information, see [www.gfi.com/exinda](http://www.gfi.com/exinda).

## Prioritization of VLL traffic (example)

1. On the **Advanced** tab, add a new QoS reservation rule with the destination port as 9899 for the UDP protocol .
2. Create another rule using port 9899 in case your remote LAN has both client and server configured on the same GFI device:



| Field                                 | Value                    |
|---------------------------------------|--------------------------|
| Type                                  | QoS Reservation          |
| Device Type*                          | wan                      |
| Device Index*                         | 1                        |
| Active                                | Yes                      |
| Notes                                 | high priority for tunnel |
| Guaranteed Rate (Kbit or percentage)* | 500                      |
| Max Rate (Kbit or percentage)         |                          |
| Buffer Limit (packets)*               | 5                        |
| Burst (bytes)                         |                          |
| Priority*                             | High Priority            |
| Direction*                            | ingress                  |
| Source IP/netmask                     |                          |
| Destination IP/netmask                |                          |
| Protocol                              | udp                      |
| Source Port                           |                          |
| Destination Port                      | 9899                     |

Screenshot 245: Example of prioritization of VLL traffic

## Virtual leased line

The Exinda SD-WAN can be configured to peer/mate with another Exinda SD-WAN device that has a server license to create a data communication tunnel between the pair of Exinda SD-WAN units. By virtue of combining the capacity of multiple resources, the capacity of the tunnel created can be close to the sum of the capacities of the individual resources. Moreover, the reliability of the communication tunnel is increased since the communication tunnel remains operational as long as there is at least one active communication path between the Exinda SD-WAN devices.

It is also possible for each Exinda SD-WAN to peer with multiple Exinda SD-WAN devices. A bi-directional tunnel can be set up between each pair of Exinda SD-WAN devices as long as one Exinda SD-WAN has the server license activated. For each bidirectional tunnel, one Exinda SD-WAN device is designated as a “server” and the other is designated as a “client”. For example, one Exinda SD-WAN device might act as a server for several other Exinda SD-WAN devices, each of which acts as a client.

A tunnel can be configured in two possible modes, “default” and “proxy”. For each end of the tunnel, there is an associated LAN address. In the default mode, traffic that appears at one Exinda SD-WAN that matches the associated LAN address of the remote Exinda SD-WAN is forwarded to the remote Exinda SD-WAN through the tunnel. In the proxy mode, all traffic that appears at the Exinda SD-WAN on the client side that does not match the local LAN address is sent to the remote Exinda SD-WAN device over the tunnel. Also, in proxy mode, traffic appearing at the Exinda SD-WAN on the server side that matches the LAN address of the Exinda SD-WAN on the client side is forwarded to the Exinda SD-WAN device on the client side over the tunnel. LAN addresses are specified by an IP address and associated network mask.

The VLL feature is configured on the **Home** tab of the Exinda SD-WAN Management Interface, in the **Remote LAN** table. It looks similar to the following:

| Connection                | Remote Network   | Status                                | Download/Upload | Peak     |
|---------------------------|------------------|---------------------------------------|-----------------|----------|
| <a href="#">server218</a> | 192.168.250.0/24 | <a href="#">Server Connections: 2</a> | 0/0 Kbps        | 0/0 Kbps |

Screenshot 246: Remote LAN example

To create a tunnel, two Exinda SD-WAN devices are required. A tunnel is configured by creating an instance of a server connection at one Exinda SD-WAN and creating a matching instance of a client connection at the other Exinda SD-WAN. For security purposes, each bi-directional tunnel has an associated connection name and password, which must be specified identically at each end of the tunnel when the tunnel is configured.

By default, packets that are sent over the tunnel are encapsulated in UDP packets in order to provide robustness against certain types of packet filtering that may be commonly present in ISP access networks.

To create an instance of a server connection or a client connection at a Exinda SD-WAN device:

1. Go to the **Home** tab of the Exinda SD-WAN Management Interface for that Exinda SD-WAN device.
2. Click **Add Remote LAN** at the top of the Remote LAN table.
3. Set the **Connection Mode** as **Client** or **Server** in the menu at the top of the pop-up.

To configure the device in the **Server** mode:

1. Select **Server** for the **Connection Mode**.
2. Enter a text string for the **Connection Name** field, and an associated password in the **Password** field. To provide security, these fields must exactly match the corresponding fields for the associated connection at the other Exinda SD-WAN device in order for the tunnel to become operational.
3. Specify the network address of the LAN associated with the remote Exinda SD-WAN device in the **Remote LAN Network** and the **Remote LAN Netmask** fields.
4. After the parameters for the connection are entered, click **Add** to add the corresponding connection.

After a connection is added, the parameters can be edited by clicking on the corresponding entry in the **Remote LAN** table, modifying the entries in the pop-up menu, and clicking **Edit**. You can also remove a connection by deleting it.

To configure the device in the **Client** mode:

1. Select **Client** for the **Connection Mode**.
2. Specify which WAN ports on the Exinda SD-WAN are used for the tunnel by selecting or clearing each listed interface as appropriate, or selecting **ALL** to use all available interfaces for the tunnel.
3. The specification of which WAN ports is used on the remote side of the tunnel (that is, the server side) is done by providing a list of IP addresses (separated by commas) of these WAN ports in the **Remote IP/Name** field. Alternatively, instead of a list of IP address, a DNS domain name can be entered. The domain name should resolve to the list of IP addresses associated with the server that is used for the tunnel.
4. Specify the network address of the remote LAN in the **Remote LAN Network** and the **Remote LAN Netmask** fields.
5. The **Connection Name** and **Password** fields should contain text strings which exactly match the associated connection entries at the remote Exinda SD-WAN.
6. After the parameters for the connection are entered, click **Add** to add the corresponding connection.

The order in which the server and client are configured is arbitrary. Currently, the Exinda SD-WAN does not support overlapping LAN addresses on each Exinda SD-WAN; the network addresses for the Exinda SD-WAN devices at each end of the VLL tunnel must be distinct and non-overlapping. Once both the client and server devices are configured, the **Status** field of the **Remote LAN** table is appropriately updated on the Exinda SD-WAN Management Interface for each device. The entries in the **Status** column can be clicked on, and a pop-up menu appears that provides additional information regarding the state of the tunnel.

## VLL Interface Socket buffers

| Field                  | Value                              |
|------------------------|------------------------------------|
| Connection Name        | <input type="text"/>               |
| Device Type            | - <input type="button" value="v"/> |
| Device Index           | <input type="text"/>               |
| Destination IP         | <input type="text"/>               |
| Local Receive Buffer*  | <input type="text"/>               |
| Local Send Buffer*     | <input type="text"/>               |
| Remote Receive Buffer* | <input type="text"/>               |
| Remote Send Buffer*    | <input type="text"/>               |

Apply

Screenshot 247: VLL Interfaces Sockets buffers

The VLL interface socket buffers are required in scenarios where the lines bonded by the VLL have huge difference in the latencies (for example, combining a DSL and a satellite line). In these cases, fine-tuning the VLL socket buffers leads to an optimized performance.

## Interface binding support

The Exinda SD-WAN supports binding of specific outbound traffic to a specific WAN interface or VLL data tunnel (with the optional VLL feature). This is sometimes useful for some applications or users that require certain traffic to be routed over the same WAN interface or over a certain tunnel.

To configure the interface binding feature, go to the **Advanced** tab. To configure an interface binding rule, click **Add Interface Binding Rule**. This generates a pop-up window where the information specifying an interface binding rule can be specified. The rule action is specified by an Interface which can be either WIRED WAN 1, WIRED WAN 2, WIRED WAN 3, WIRED WAN 4, or possibly Cellular WAN 1 if a cellular broadband data modem is plugged into the Exinda SD-WAN. If a remote LAN connection has been defined (with the optional VLL feature), that can also be specified as an interface.

The specified traffic is preferentially routed over the WAN interface with the index that is selected for the rule, or to the corresponding remote LAN connection.

The traffic specification for the rule consists of a protocol designation, a port number, and an IP address or IP subnet. Only packets which match the corresponding protocol designation, port number, and IP address or IP subnet is affected by the rule (that is, the “and” of these three conditions). A match for the port number occurs when either the **Source Port** field or the **Destination Port** field matches with the port numbers specified. A port number specification can be a single integer, but it can also be a range. For example, 56–59 would be the range that includes 56, 57, 58, and 59. Values and ranges can also be separated by commas. For example, 56, 59 corresponds to the set containing the values 56 and 59. If the **Port Number** field is left blank, there is no match condition applied to the **Source Port** or **Destination Port** fields, that is, all ports match.

A match for the IP address occurs when either the source IP address of the packet or the destination IP address of the packet matches the given IP address. If an IP subnet is given, a match occurs when the source IP address or the destination IP address of the packet is within the IP subnet. If the **IP address / subnet** field is left blank then there is no match condition applied to the IP address field, that is, all IP address / subnet values matches.

The protocol designation can be **ANY**, **TCP**, **UDP**, **IPSEC**, **PPTP**, or **Other**. The **ANY** choice means any protocol matches. With the **Other** protocol designation, the corresponding protocol number is also specified.

All of the match conditions, if any, for the port number, the IP address or subnet, and the protocol designation must be met in order to trigger the port forwarding rule. When the port forwarding rule is triggered, the traffic is routed in the manner specified, for example, **WIRED WAN 1**, **WIRED WAN 2**, **WIRED WAN 3**, **WIRED WAN 4**, **Cellular WAN 1** if a cellular broadband data modem is plugged into the Exinda SD-WAN, and **Remote LAN 1** if a VLL connection has been configured (with the optional VLL feature).

There are two port numbers that need special consideration for interface binding. Specifically, traffic on port 80 (**HTTP**) and port 443 (**SSL**) should be selected for interface binding only with extreme care. Nominally, traffic for these ports should be handled by the same WAN interface.

An optional **Note** field is provided for convenience and typically is used to document what the interface binding rule is used for (for example, for a particular application or user). Once all the information for an interface binding rule has been entered, click **Add** to enable the rule.

Multiple interface binding rules can be added in order to provide more flexible binding of traffic. If an interface binding rule that was added previously needs to be deleted, click **Delete** next to the listed binding rule. If an interface binding rule needs to be edited, the rule can first be deleted and the modified rule can then be entered as before. The Exinda SD-WAN processes outgoing traffic by sequential application of the interface binding rules specified. This means if there are conflicts in the interface binding rules, the rule nearest the top of the list takes precedence. By default, the list of rules is such that rules entered earlier is lower in the list.

## VLL Cloud Relay

Each Exinda SD-WAN subscription can have up to 1 VLL connected to our Cloud Relay service. The Cloud Relay Service connects the Exinda SD-WAN device to another GFI hosted Exinda SD-WAN device, hosted on Amazon AWS. The Cloud Relay service utilizes the VLL feature explained above. To set it up, you should contact [Support](#).

## Performance tab

The **Performance** tab displays graphical performance metrics. A history of your traffic usage based on protocol, including total, network, HTTP, TCP, UDP, other, or interface, such as LAN or WAN1, are presented with multicolor graphs with a time scale from seconds, minutes, hours, days and months. There is also a check-box for live, realtime performance.

A typical example is shown below:

Management Interface - Status Information for Exinda SD-WAN

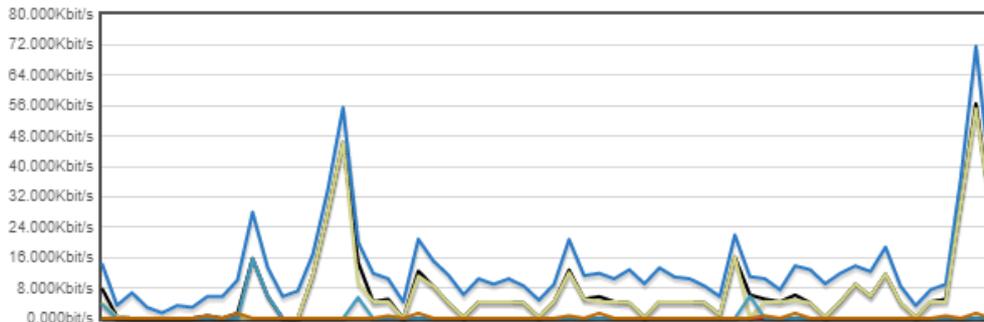


Live Last minute ▾

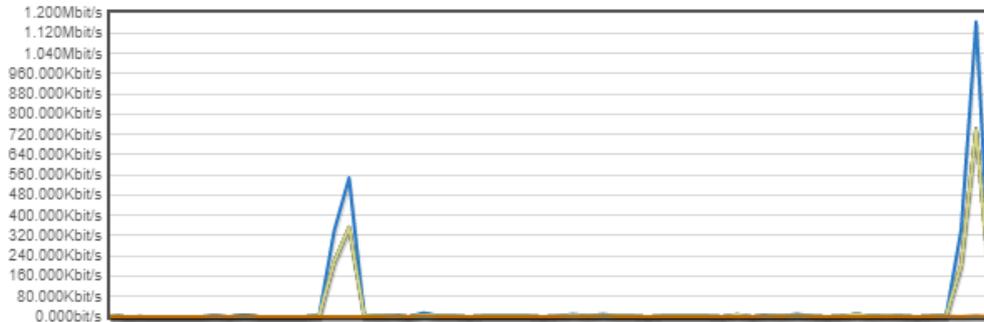
| Include                             | Name         | Total               | Incoming            | Outgoing            |
|-------------------------------------|--------------|---------------------|---------------------|---------------------|
| <input type="checkbox"/> all/none   | <b>Total</b> | <b>95.008Kbit/s</b> | <b>27.384Kbit/s</b> | <b>67.624Kbit/s</b> |
| <input type="checkbox"/>            | LAN          | 0.000bit/s          | 0.000bit/s          | 0.000bit/s          |
| <input checked="" type="checkbox"/> | WAN 1        | 95.008Kbit/s        | 27.384Kbit/s        | 67.624Kbit/s        |
| <input checked="" type="checkbox"/> | WAN 2        | 0.000bit/s          | 0.000bit/s          | 0.000bit/s          |
| <input checked="" type="checkbox"/> | WAN 3        | 0.000bit/s          | 0.000bit/s          | 0.000bit/s          |

Protocols Interfaces

Incoming



Outgoing

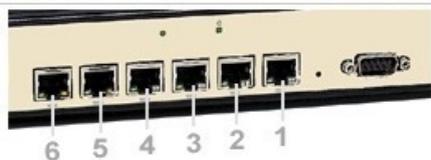


Screenshot 248: The Performance tab

### Status tab

The **Status** tab displays auxiliary information about the Exinda SD-WAN, including the software revision number, as well as links to a menu to upgrade the software version.

A log of system events is displayed. The Remote Syslog feature can be configured, as described in For more information, refer to [Remote syslog](#) (page 510).. A list of DHCP leases is also displayed.

**Management Interface - Status Information for Exinda SD-WAN**


| <b>About</b>               | Refer to your user guide for support, contact and warranty information.                                                                       |                                  |           |         |  |  |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|-----------|---------|--|--|
| <b>Model</b>               |                                                                                                                                               |                                  |           |         |  |  |
| <b>Software Revision</b>   | Exinda-SDWAN-1.17.1456881548                                                                                                                  | <a href="#">Firmware Upgrade</a> |           |         |  |  |
| <b>Software Build Date</b> | 2016-03-01 18:10:56-08:00                                                                                                                     |                                  |           |         |  |  |
| <b>DHCP Leases</b>         | <input type="text"/>                                                                                                                          |                                  |           |         |  |  |
| <b>Timezone</b>            | <a href="#">America/Los Angeles, 2019-05-21 23:51:38</a><br>NTP synchronized: <b>False</b>                                                    |                                  |           |         |  |  |
| <b>Logs</b>                | <table border="1"> <thead> <tr> <th>Date/Time</th> <th>Message</th> </tr> </thead> <tbody> <tr> <td> </td> <td> </td> </tr> </tbody> </table> |                                  | Date/Time | Message |  |  |
| Date/Time                  | Message                                                                                                                                       |                                  |           |         |  |  |
|                            |                                                                                                                                               |                                  |           |         |  |  |
| Remote Syslog Server:      | <input type="text"/>                                                                                                                          |                                  |           |         |  |  |
|                            | <input type="button" value="Apply"/>                                                                                                          |                                  |           |         |  |  |

Screenshot 249: The Status tab

## Administrative features

In this section, the administrative features of the Exinda SD-WAN are described.

### Web interface password

The Exinda SD-WAN Management Interface can be configured so that it can be accessed only with the entry of a username and password. With the default factory settings, entry of a password is not required.

The Exinda SD-WAN Management Interface password can be configured on the **Admin** tab so that it would require a password to access it.

The password should be changed with extreme care — if it is forgotten or mistyped it is not possible to access the Exinda SD-WAN Management Interface.

To configure the feature:

1. Go to the **Admin** tab.
2. Click the status under **Web Interface Password** and set it to **Enabled**. The default username is `admin`.
3. Set the password by entering it in the **Web Interface Password** field. Type in the password twice to ensure it is accurately recorded, then click **Apply**.

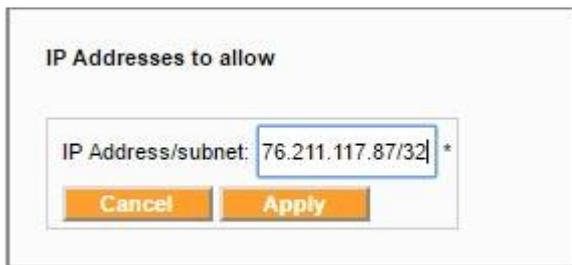
Configuration of this feature causes the Exinda SD-WAN to restart.

### Allow web IPs

This feature allows restricting the access to the Exinda SD-WAN Management Interface to only the IPs provided in this field. For example, if the administrator for the device wants to access the device remotely from the IP

76 . 211 . 117 . 87, the **Web Allow IPs** field can be configured accordingly, which would restrict the access to the

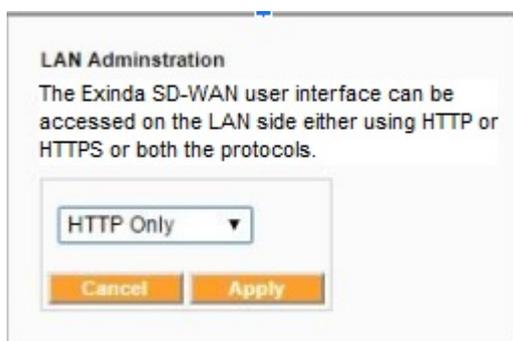
Exinda SD-WAN Management Interface to only this particular IP.



Screenshot 250: Allowing web IPs

### LAN user interface

The LAN user interface defines the protocol used to access the Exinda SD-WAN Management Interface from the LAN side. By default this would be set to HTTP.



Screenshot 251: LAN user interface

### WAN HTTP Port (Remote web interface port)

The Exinda SD-WAN Management Interface can be configured so that it can be accessed externally from the internet via the Exinda SD-WAN WAN ports. To enable this feature, go to the **Admin** tab and click on the status indicator of the **WAN HTTP Port** entry in the displayed status table. The pop-up window allows configuration of the feature. If the feature is enabled, a port number must be specified.

If the feature is enabled, the Exinda SD-WAN Management Interface can be remotely accessed from the internet through any of the public IP addresses for the Exinda SD-WAN. These public IP addresses can be obtained from the **External IP** column in the WAN status table on the **Home** tab of the Exinda SD-WAN Management Interface. The URL through which the Exinda SD-WAN Management Interface can be accessed is given by the external IP address appended with the assigned port number. For example, if one of the external IP addresses for the Exinda SD-WAN is 76 . 211 . 117 . 87 and the Remote Web Interface is enabled through port 8080, then the Exinda SD-WAN Management Interface can be accessed through the URL `http://76.211.117.87:8080`.

If Pass Through mode is enabled on WAN interface 1, the Exinda SD-WAN Management Interface can be accessed externally from the internet through the IP address of the router/firewall that is connected to the LAN port of the Exinda SD-WAN. For example, if the IP address of this router is 76 . 37 . 181 . 2 and the WAN HTTP Port is accessed through port 8080, then the Exinda SD-WAN Management Interface can be externally accessed through the URL `http://76.37.181.2:8080`. Note that in this example from a host on the LAN of the Exinda SD-WAN, packets addressed to 76 . 37 . 181 . 2 are forwarded to the router/firewall and not the Exinda SD-WAN. To access the Exinda SD-WAN Management Interface from the Exinda SD-WAN LAN, the Exinda SD-WAN LAN address should be used (for example, the default Exinda SD-WAN LAN address 192 . 168 . 254 . 99).

If the WAN HTTP Port is enabled, then it is highly recommended that a web interface password be configured to prevent unauthorized access to the Exinda SD-WAN from the public Internet.

## WAN HTTPS Port

The WAN HTTPS Port 8081 allows remote access of the Exinda SD-WAN Management Interface by using HTTPS. The URL through which the Exinda SD-WAN Management Interface can be accessed is given by the external IP address appended with the assigned port number. For example, if one of the external IP addresses for the Exinda SD-WAN is 76.211.117.87 and the Remote Web Interface is enabled through port 8081, then the Exinda SD-WAN Management Interface can be accessed through the URL `http://76.211.117.87:8081`.

If Pass Through mode is enabled on WAN interface 1, the Exinda SD-WAN Management Interface can be accessed externally from the Internet through the IP address of the router/firewall connected to the LAN port of the Exinda SD-WAN. For example, if the IP address of this router is 76.37.181.2 and the WAN HTTP Port is accessed through port 8081, then the Exinda SD-WAN Management Interface can be externally accessed through the URL `http://76.37.181.2:8081`. Note that in this example from a host on the LAN of the Exinda SD-WAN, packets addressed to 76.37.181.2 is forwarded to the router/firewall and not the Exinda SD-WAN. To access the Exinda SD-WAN Management Interface from the Exinda SD-WAN LAN, the Exinda SD-WAN LAN address should be used (e.g. the default Exinda SD-WAN LAN address 192.168.254.99).

## Maintenance mode

The Exinda SD-WAN has a maintenance mode for remote troubleshooting by GFI Support. This allows external access to the Exinda SD-WAN by GFI Support over the internet. Normally this mode should be disabled, and it is disabled in the default factory settings.

To configure the feature:

1. Go to the **Admin** tab.
2. Click the status for **Maintenance Mode**, then set it to **Disabled** or **Enabled**. This generates a pop-up window to prompt the user to configure the feature.
3. Click **Apply** to make the changes.

## VRRP configuration

The Exinda SD-WAN has an option to configure VRRP (Virtual Router Redundancy Protocol), which is designed to increase the availability of the default gateway-servicing hosts on the same subnet. By default this feature is disabled but can be enabled on the Admin tab.

The screenshot shows the 'exinda SD-WAN' administration interface. At the top, there are navigation tabs: HOME, ADVANCED, VLL, FIREWALL, PERFORMANCE, STATUS, and ADMIN. The 'ADMIN' tab is selected. Below the tabs, the page title is 'Management Interface - Administration for Exinda SD-WAN'. On the left side, there is a section for 'Backup/Restore Settings' with buttons for 'Download' and 'Send', and a 'Browse...' button. The main content area is a table of configuration settings. The 'VRRP Config' setting is highlighted with a red box and is currently set to 'Disabled'. Other settings include 'Web Interface Password' (Disabled), 'Web Allow IPs' (0.0.0.0/0), 'LAN User Interface' (HTTP only), 'WAN HTTP Port' (Disabled), 'WAN HTTPS Port' (8081), 'VOM Config' (Disabled), 'Maintenance Mode' (Enabled), 'Force VLL Reroutes' (Disabled), 'Mail Notification' (Disabled), and 'SNMP Config' (Disabled).

| Setting                | Status   | Value          | Setting   | Status             |           |
|------------------------|----------|----------------|-----------|--------------------|-----------|
| Web Interface Password | Disabled | Web Allow IPs  | 0.0.0.0/0 | LAN User Interface | HTTP only |
| WAN HTTP Port          | Disabled | WAN HTTPS Port | 8081      | VOM Config         | Disabled  |
| Maintenance Mode       | Enabled  | VRRP Config    | Disabled  | Force VLL Reroutes | Disabled  |
| Mail Notification      | Disabled | SNMP Config    | Disabled  | Ignore missing DNS | Disabled  |

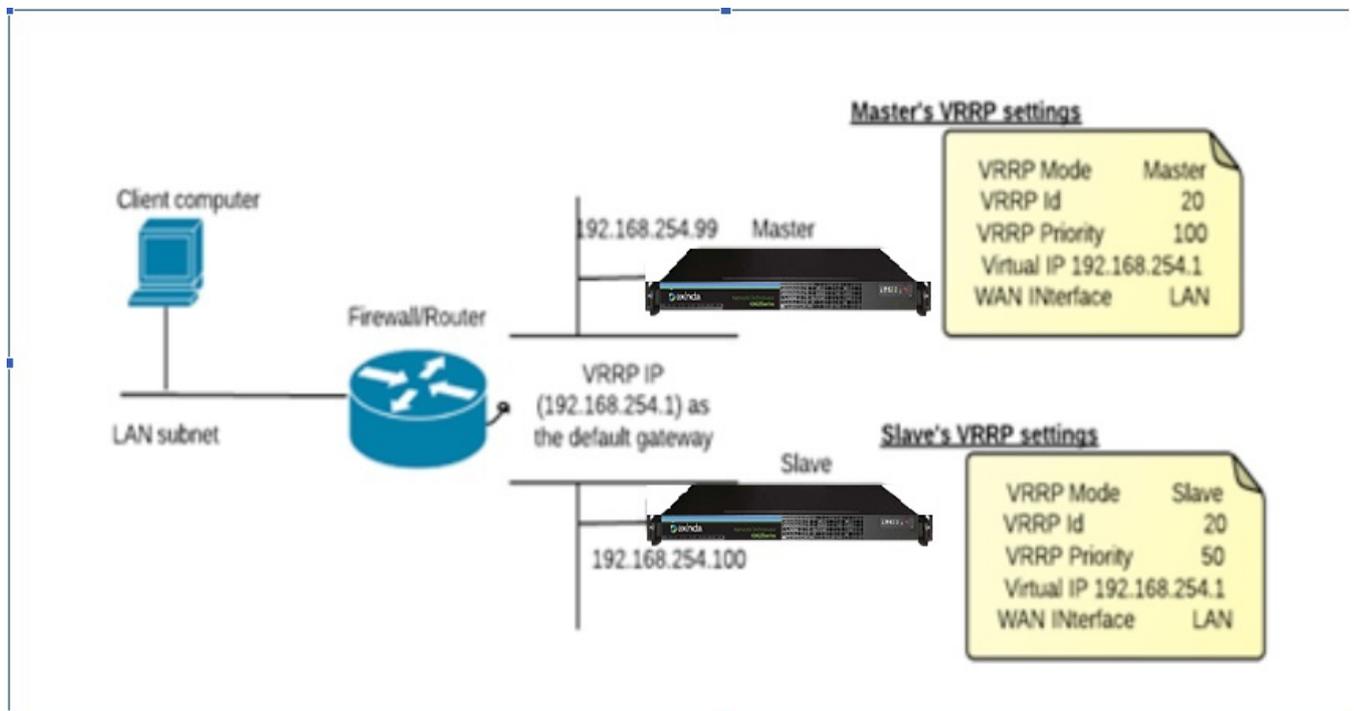
Screenshot 252: VRRP configuration

VRRP mode can be configured to be either as master or slave depending on which physical router is doing the actual routing. In case of failure of the master router, the slave router which was configured as a virtual router automatically

replaces it. **VRRP ID** is the Virtual Router Identifier (VRID), which uniquely identifies each virtual router in the subnet. This is a configurable item in the range 1-255 (decimal) and has no default value.

**VRRP Priority** is an 8-bit unsigned integer field with higher value indicating higher priority. The master should be given the highest priority. Care should be taken in configuring this field as the master should always be given highest priority compared to the slaves or this could cause instability in the network. **Virtual IP** is the IP address of the Exinda SD-WAN device depending on which interface is connected to the subnet.

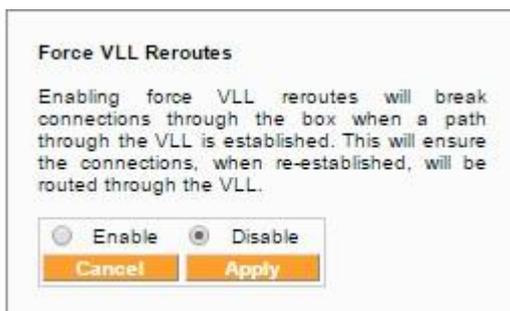
Here is an example setup:



Screenshot 253: Example of VRRP setup

### Force VLL reroutes

Enabling force VLL reroutes break connections through the box when a path through the VLL is established. This ensures the connections, when re-established, is routed through the VLL.



Screenshot 254: Force VLL reroutes configuration

### Mail notification (E-mail alarms)

The Exinda SD-WAN supports sending of email alarms to a specified address when a serious event occurs. A "serious" event is defined by a WAN interface going down or coming back up without manually enabling or disabling the WAN interface.

To configure email alarms:

1. Go to the **Admin** tab.
2. Click the status for **Mail Notification**, which is either **Disabled** or **Enabled**.
3. Enter the email address and the IP address of the SMTP server in the pop-up window to receive the alarms. Typically, this is the SMTP server on the Exinda SD-WAN LAN.

In order to test the configuration, a WAN interface can be manually brought down and up again by disconnecting the cable from an active WAN interface which triggers an email alarm to be sent to the designated address and server.

### SNMP configuration

The Exinda SD-WAN has a built in SNMP (Simple Network Management Protocol) agent running on the device which can be accessed by any SNMP browser or SNMP client application. The Exinda SD-WAN supports MIB2 (RFC 1213) management information base (MIB). You can enable the SNMP feature on the **Admin** tab.

The screenshot shows the 'Management Interface - Administration for Exinda SD-WAN' page. On the left, there is a network diagram with ports 1-6 and buttons for 'Reboot', 'Shutdown', and 'Reset'. Below that is a 'Backup/Restore Settings' section with 'Backup To File' (Download) and 'Restore From File' (Browse... Send) options. The main configuration table is as follows:

|                        |          |                |           |                    |           |
|------------------------|----------|----------------|-----------|--------------------|-----------|
| Web Interface Password | Disabled | Web Allow IPs  | 0.0.0.0/0 | LAN User Interface | HTTP only |
| WAN HTTP Port          | Disabled | WAN HTTPS Port | 8081      | VOM Config         | Disabled  |
| Maintenance Mode       | Enabled  | VRRP Config    | Disabled  | Force VLL Reroutes | Disabled  |
| Mail Notification      | Disabled | SNMP Config    | Disabled  | Ignore missing DNS | Disabled  |

The 'SNMP Config' pop-up window is open, showing the following settings:

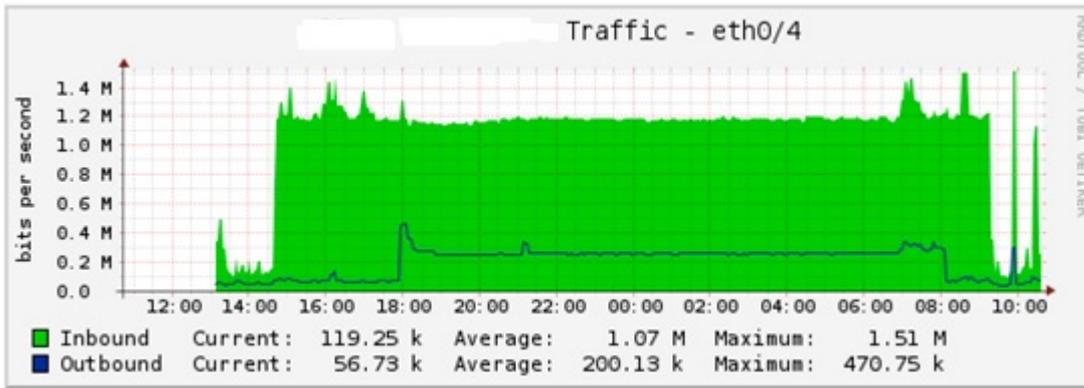
- Disabled
- Enabled
- Snmp read community: 3sLnzh1dtI9c \*
- Snmp write community: g380c4inImGe \*
- Trap server IP: 127.0.0.1 \*

Buttons for 'Cancel' and 'Apply' are at the bottom of the pop-up.

Screenshot 255: SNMP configuration

Asynchronous notification by using SNMP traps for WAN links getting **Connected** and **Disconnected** can also be pushed to a **Trap server IP**.

By using external SNMP management and graphing tools you can retrieve interface statistics from the GFI Exinda SD-WAN device. The following is an example graph for one of the WAN statistics using MRTG grapher:



Screenshot 256: Example Graph of WAN Statistic

Note that any host on the Exinda SD-WAN LAN can access the SNMP interface. For external network access to the SNMP interface the Exinda SD-WAN firewall needs to be configured to allow SNMP traffic from the outside that uses port 161. To do this, go to the **Firewall** tab in the Exinda SD-WAN Management Interface and add a rule that permits inbound traffic on port 161 for UDP protocol. In rare cases it may be necessary to also add a rule that permits inbound traffic on port 161 for TCP protocol.

### Ignore missing DNS

If this option is disabled, if your DHCP server fails to provide DNS servers the device attempts to re-DHCP onto the network. This handles cases where the DHCP server is in a bad state, but can prevent accessing an unconfigured server to configure it.



Screenshot 257: Ignore Missing DNS configuration

### Assigning a static IP address to a PC

It is likely that a search on the Internet easily yield instructions on how to configure the management computer (PC) with a static IP address. For convenience, instructions on how to do this for a Windows XP-based PC are given below.

To set a static IP address:

1. Open the Windows **Start** menu.
2. Open the **Control Panel**.
3. Classic view: Open **Network Connections**.  
Category view: Open **Network and Internet Connections**, and then **Network Connections**.
4. Double-click on the active **LAN or Internet connection**.
5. Click **Properties**. This opens the Local Area Connection Properties window.

6. In the General tab, highlight the **Internet Protocol (TCP/IP)**, and click **Properties**.



This opens the Internet Protocol (TCP/IP) Properties window.

7. In the General tab, click **Use the following IP address**, and enter:

- » IP address. The static IP address to be assigned to this computer (it is recommended that 192 . 168 . 254 . 10 be used unless changed the IP address of the Exinda SD-WAN has been reconfigured).
- » Subnet mask. Subnet mask used by the Exinda SD-WAN (usually this is 255 . 255 . 255 . 0).
- » Default gateway. IP address of the Exinda SD-WAN (usually this is 192 . 168 . 254 . 99).

8. In **Use the following DNS server addresses**, enter all the IP addresses for the DNS servers the Exinda SD-WAN uses (the address 192 . 168 . 254 . 99 can be used, or this can be left blank to configure the Exinda SD-WAN).

9. Click **OK**.

10. Click **OK** again to close each window.

11. Restart the computer.

12. Then, check the IP address again, to make sure that the changes were applied.

To check the IP address:

1. Open Windows **Start** menu.
2. Select **Run**. Type: command and click **OK**.
3. At the blinking cursor, type: **ipconfig /all** and press Enter.

4. Look for these entries near the end of the list:

» **DHCP Enabled.** “No” means the IP address is static. “Yes” means it is dynamic.

» **IP Address.** This is the current IP address.

5. To exit, at the blinking cursor, type: exit and press Enter.

```
C:\WINDOWS\System32\command.com
C:\DOCUME~1\DEFAULT>ipconfig /all
Windows IP Configuration

Host Name . . . . . : vaio
Primary Dns Suffix . . . . . :
Node Type . . . . . : Unknown
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No

Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . :
   Description . . . . . : Realtek RTL8139/810
Ethernet NIC
   Physical Address. . . . . : 08-00-46-1E-27-B7
   Dhcp Enabled. . . . . : No
   IP Address. . . . . : 192.168.1.50
   Subnet Mask . . . . . : 255.255.255.0
   Default Gateway . . . . . : 192.168.1.1
   DNS Servers . . . . . : 192.233.85.19
                           192.233.145.6
   NetBIOS over Tcpip. . . . . : Disabled

C:\DOCUME~1\DEFAULT>
```

### 0.17.4 Maintenance and usage

The Exinda SD-WAN is an appliance that is designed to run at all times. It does not need to be powered off or restarted for normal usage. For moving or storage, the Exinda SD-WAN may be powered off at any time, but it is recommended that the **SHUTDOWN** option be used on the Admin tab. Note that this puts the Exinda SD-WAN in a state whereby the power to the unit can be safely interrupted.

In addition to the **SHUTDOWN** option on the Admin tab, there are the **REBOOT** and **RESET** options. The **REBOOT** option causes the Exinda SD-WAN to reboot. The **RESET** option resets all configuration data in the Exinda SD-WAN Management Interface to their factory defaults.

The Exinda SD-WAN should be placed on a flat and dry surface in a well-ventilated area. Overheating shortens the life of the Exinda SD-WAN.

Opening the case of the Exinda SD-WAN, except under instruction of GFI, Inc., voids the warranty.

### Software upgrade

You can find the latest software version on the **Status** tab of the Exinda SD-WAN Management Interface and upgrade the software if required.

To determine the latest software version,

1. Go to the **Status** tab of the Exinda SD-WAN Management Interface.
2. Click **Firmware Upgrade** located next to the current **Software Revision**.

This causes the Exinda SD-WAN to access the internet and determine the latest software version available for the Exinda SD-WAN and display it at the top of the list.

3. Select the latest software version in the list and click **Update**.

A status bar is displayed indicating the progress of the software upgrade to the device. It may take several minutes for the upgrade to complete.

4. To apply the upgrade, click **Finalize** after the progress bar has completed.

The upgrade is effective upon the next device reboot.

After upgrading the system software, you can verify that the software has been upgraded by checking the version number in the **Status** tab. You should clear your browser cache to make sure that you are not viewing a cached page.

### Firmware reset with USB key

A USB key for resetting the software and settings is included with each unit. It is recommended that you keep the USB key in a safe and secure place, since it is needed to restore the unit to factory default condition, when necessary.

You should use the appropriate procedure as indicated to reset or reload the software into the Exinda SD-WAN unit.

When resetting the Exinda SD-WAN system software, it is possible to do only a software reset without resetting the factory defaults, if the USB key is removed at the appropriate time during the upgrade process. Alternatively, to avoid the USB key resetting the Exinda SD-WAN to factory settings and only reset the firmware, a file named `noreset` can be added to the USB key where the file can be of any type.

#### Firmware Reset with or without restore to factory default settings

1. Power off the device.
2. Unplug all USB attached devices from the Exinda SD-WAN, and insert the USB key that came with the unit into any USB slot on the front panel.
3. Power on the device.
4. The device loads the firmware from the USB key. When the firmware upload is complete, the Exinda SD-WAN beeps three times.

To avoid resetting to factory default settings, the Exinda SD-WAN may safely be powered off following the 3 beeps.

5. To reset to factory default settings, continue to wait for 1 minute after the beeps and the device powers itself off.
6. Remove the USB key and power the unit back on. It is now fully reset.

### 0.17.5 Abbreviations and acronyms

| Term | Description                        |
|------|------------------------------------|
| ADSL | Asymmetric Digital Subscriber Line |
| APN  | Access Point Network               |
| ARP  | Address Resolution Protocol        |

| Term     | Description                                                |
|----------|------------------------------------------------------------|
| CDMA     | Code Division Multiple Access                              |
| CIDR     | Classless Inter-Domain Routing                             |
| DHCP     | Dynamic Host Configuration Protocol                        |
| DMZ      | Demilitarized Zone                                         |
| DDNS     | Dynamic DNS                                                |
| DNS      | Domain Name Server                                         |
| DSCP     | Differentiated Services Code Point                         |
| DSL      | Digital Subscriber Line                                    |
| GSM      | Global System for Mobile Communications                    |
| HSPA     | High Speed Packet Access                                   |
| ICMP     | Internet Control Message Protocol                          |
| IP       | Internet protocol                                          |
| IPSec    | Internet Protocol Security                                 |
| LAN      | Local area network                                         |
| MAC      | Media Access Control                                       |
| MIB      | Management Information Base                                |
| MIBv2    | Management Information Base, version 2                     |
| MTU      | Maximum Transmission Unit                                  |
| NAT      | Network Address Translation                                |
| NTP      | Network Time Protocol                                      |
| PPPoE    | Point-to-Point Protocol over Ethernet                      |
| PPTP     | Point-to-Point Tunneling Protocol                          |
| QoS      | Quality of Service                                         |
| SD-WAN   | Software-Defined Wide Area Network                         |
| SMTP     | Simple Mail Transfer Protocol                              |
| SNMP     | Simple Network Management Protocol                         |
| TCP      | Transmission Control Protocol                              |
| UDP      | User Datagram Protocol                                     |
| UPnP     | Universal Plug and Play                                    |
| USB      | Universal Serial Bus                                       |
| VDSL     | Very High Bitrate Digital Subscriber Line                  |
| VOIP     | Voice over Internet Protocol                               |
| VOIP/SIP | Voice over Internet Protocol / Session Initiation Protocol |

| Term | Description                        |
|------|------------------------------------|
| VPN  | Virtual Private Network            |
| VRID | Virtual Router Identifier          |
| VRRP | Virtual Router Redundancy Protocol |
| WAN  | Wide Area Network                  |

### 0.17.6 Support information

If you purchased your Exinda SD-WAN directly from GFI, please visit the support portal at [Support](#).

You can access the knowledge base when you create your account.

# Settings

Learn how to configure Exinda Appliance to meet your requirements.

---

---

## 0.18 Network settings

Learn how to configure the network setting for your Exinda Appliance(s).

---

|                                                 |     |
|-------------------------------------------------|-----|
| 0.18.1 NIC configuration .....                  | 535 |
| 0.18.2 IP address configuration .....           | 537 |
| 0.18.3 Routes configuration .....               | 540 |
| 0.18.4 DNS and domain names configuration ..... | 541 |
| 0.18.5 HTTP proxy configuration .....           | 544 |
| 0.18.6 Email configuration .....                | 544 |
| 0.18.7 SNMP configuration .....                 | 547 |
| 0.18.8 Integrate with Active Directory .....    | 553 |
| 0.18.9 IPMI Configuration .....                 | 572 |
| 0.18.10 Overview of QoS by host .....           | 576 |

---

### 0.18.1 NIC configuration

The NIC settings page is used to set the speed, duplex, and MTU of the System NICs, to set the behavior of the bridges in the event of an appliance failure (for example, fail to wire), and to set the behavior of the second port in a bridge pair when the first port goes down (for example, link state mirroring).

#### Interface Settings

You need the Exinda appliance and devices that are connected to the appliance to have the same speed and duplex settings for their network interfaces. In most cases the default settings will work as the Exinda is setup to auto-negotiate. However, some equipment is not compatible with this.

If the appliance and connected devices are using different speeds and duplex settings, the devices may not communicate and traffic may be dropped. In this case, you may notice collisions, errors, packet loss, and network delays on the Exinda NICs, which will cause the System health status to show as "Warning" and the offending interface(s) will be highlighted.

To resolve this, check if the router or switch is hard-coded to a speed or duplex setting. If hard-coded, then either set all devices to auto-negotiate or set the Exinda device to the same speed and duplex mode.

#### NOTE

For further troubleshooting, click on the system warning or view the NIC Diagnostics by clicking on the View NIC Diagnostics link.

### View NIC Diagnostics...

| Interface | Media        | HW Address        | Speed | Duplex | MTU  | Link Status                                                    |
|-----------|--------------|-------------------|-------|--------|------|----------------------------------------------------------------|
| eth1      | Twisted Pair | 00:22:19:D4:8D:C4 | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 100Mb/s (auto), Duplex: Full (auto)  |
| eth2      | Twisted Pair | 00:22:19:D4:8D:C5 | Auto  | Auto   | 1500 | Admin UP, Link DOWN, Speed: UNKNOWN, Duplex: UNKNOWN           |
| eth10     | Twisted Pair | 00:E0:ED:13:73:C2 | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 100Mb/s (auto), Duplex: Full (auto)  |
| eth11     | Twisted Pair | 00:E0:ED:13:73:C3 | Auto  | Auto   | 1500 | Admin UP, Link UP, Speed: 1000Mb/s (auto), Duplex: Full (auto) |

Apply Changes

Screenshot 258: Set and view the speed, duplex, MTU of NIC interfaces

### Fail to Wire (bypass)

The Fail to Wire (bypass) settings control the behaviour of the Exinda appliance bridges in the event of failure, power outage or reboot. Depending on the hardware appliance and the type of interface cards installed, fail to wire or bypass settings may be configured globally or per bridge.

| Bridge | Status | Running Mode | Enable Failover                     | On Failover |
|--------|--------|--------------|-------------------------------------|-------------|
| br10   | Active | Active       | <input checked="" type="checkbox"/> | Bypass      |
| br20   | Active | Active       | <input checked="" type="checkbox"/> | Bypass      |
| br30   | Active | Active       | <input checked="" type="checkbox"/> | Bypass      |
| br40   | Active | Active       | <input checked="" type="checkbox"/> | Bypass      |

Apply Changes

Screenshot 259: Independently controllable bypass bridges

| Bridge  | Status | Running Mode | Enable Failover                     | On Failover |
|---------|--------|--------------|-------------------------------------|-------------|
| br0,br1 | Active | Active       | <input checked="" type="checkbox"/> | Bypass      |

Apply Changes

Screenshot 260: Globally controllable bypass bridges

### Link State Mirroring

With link state mirroring, the Exinda appliance will bring down the second port of a bridge if the first port goes down. This feature allows the Exinda appliance to sit between a WAN router and a switch without blocking detection of switch outages by the router. This is a global setting that is applied to all enabled bridges.

| Link State Mirroring            |
|---------------------------------|
| <input type="checkbox"/> Enable |

Apply Changes

Screenshot 261: Configure link state mirroring

Where do I find this configuration?

Go to **Configuration > System > Network > NICs**.

To configure the NIC interfaces

1. In the interface table, you will see:

- **Interface** - Each interface corresponds to a physical port. (e.g. eth1, eth2)
- **Media** - Indicates the interface media. The options are Twisted pair or Fibre.
- **HW Address** - Indicates the MAC address of the interface.

2. Specify the **Speed** and **Duplex** at which the Exinda will negotiate with neighboring equipment. Use **Auto** speed to allow the Exinda appliance to auto-negotiate the speed with neighboring equipment. Use **Auto** duplex to allow the Exinda appliance to auto-negotiate the duplex with neighboring equipment.

3. Specify the **MTU** (maximum transmission unit) size in bytes.

4. View the **Link Status**. The link status shows whether the interface is up/down, the link is up/down, as well as the speed/duplex that has been negotiated with the neighboring equipment.

5. Click **Apply Changes**.

To configure the failover mode of the bridges

1. For each bridge, check the **Enable Failover** checkbox and set the **On Failover** status. On failover, the bridge will be automatically moved to the specified mode.

- **Bypass** - The traffic passes through the Exinda appliance but is not intercepted or manipulated.
- **No-link** - The bridge interfaces are forced to link state down (as if the cables are not plugged into the interfaces).

2. Click **Apply Changes**.

To configure the current running mode of the bridges (e.g. active, bypass, no-link)

1. For each bridge, you will see:

- **Bridge** - The name of the bridge. Where available, bridges can be controlled independently; otherwise they will be controlled globally.
- **Status** - The current status of the bridge.

2. Specify the Running Mode of the bridge.

- **Active** - The bridge is active and traffic is being intercepted by the Exinda appliance.
- **Bypass** - The traffic passes through the Exinda appliance but is not intercepted or manipulated.
- **No-link** - The bridge interfaces are forced to link state down (as if the cables are not plugged into the interfaces).

2. Click **Apply Changes**.

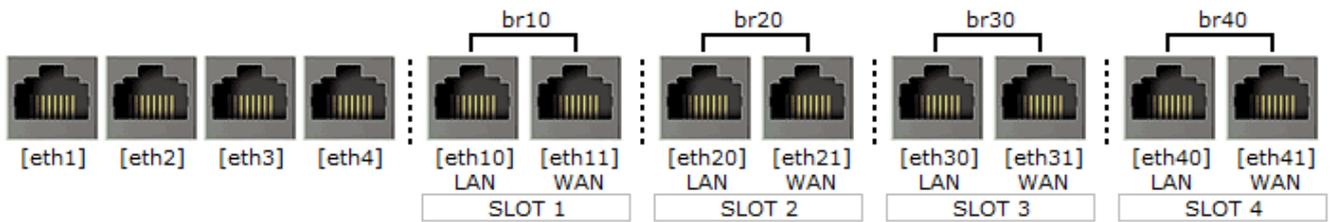
To configure link state mirroring

In the **Link State Mirroring** section, check the **Enable** checkbox.

## 0.18.2 IP address configuration

The Exinda appliance allows you to configure bridges and network interfaces as required. A bridge consists of a LAN and WAN interface. Bridges can be enabled, and roles can be assigned to an interface (Cluster, Mirror or WCCP) and IP settings applied.

The form displays an image showing the available physical interfaces, physical interface to I/O slot and physical interface to bridge assignments.



| Interface Settings                       |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| eth1                                     | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::a6ba:dbff:fe1e:9eec/64<br>Static Addresses: <input type="text" value="172.24.1.80"/> / <input type="text" value="24"/><br><input type="text"/><br>Comment: <input type="text"/> |
|                                          | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::a6ba:dbff:fe1e:9eee/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                        |
|                                          | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                           |
|                                          | Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP <input type="checkbox"/> PBR<br>Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                           |
| br10 <input checked="" type="checkbox"/> | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::2e0:edff:fe17:3870/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                                                                              |
|                                          | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::2e0:edff:fe16:c00e/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                                                                              |
| br30 <input checked="" type="checkbox"/> | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::2e0:edff:fe0e:9eea/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                                                                              |
|                                          | Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC<br>Dynamic Addresses: fe80::2e0:edff:fe1a:e8a4/64<br>Static Addresses: <input type="text"/> / <input type="text"/><br>Comment: <input type="text"/>                                                                                                                                                                                              |
| Gateway Settings                         |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| IPv4:                                    | <input type="text" value="172.24.1.1"/>                                                                                                                                                                                                                                                                                                                                                                                             |
| IPv6:                                    | <input type="text"/>                                                                                                                                                                                                                                                                                                                                                                                                                |

Interfaces that are not assigned to a bridge may have the following roles configured:

- » **Cluster**- One interface may be configured for Cluster internal use in High Availability (HA) environments. An interface and Cluster Master address should also be configured.
- » **Mirror**- One or more interfaces may be configured in Mirror mode. This mode of operation is used for out of path monitoring using a hub or switch mirror/SPAN port.
- » **WCCP** - One interface may be configured in WCCP mode. WCCP allows out of path Application Acceleration.
- » **PBR**- One interface may be configured in Policy-based Routing (PBR) mode. PBR allows for acceleration of only the types of traffic specified in the policy.

The DHCP option is enabled by default on the Exinda appliance. If a DHCP server is available, an IP address will be automatically assigned. If a DHCP address is not picked up, the Exinda defaults to the IP Address 172.14.1.57.

The VLAN configuration allows an An 802.1Q VLAN ID to be set on an interface. The VLAN ID can be between 1 and 4094.

| VLAN Settings |                                   |
|---------------|-----------------------------------|
| Interface:    | <input type="text" value="eth1"/> |
| ID:           | <input type="text"/>              |

The Cluster Master address is the external address used to access an appliance in HA environments.

| Cluster Master Settings |                                                                       |
|-------------------------|-----------------------------------------------------------------------|
| Interface:              | <input type="text" value="eth1"/>                                     |
| Master Address:         | <input type="text" value="0.0.0.0"/> / <input type="text" value="0"/> |

Further information on Clustering/HA, Mirroring and WCCP is available in the associated How To guides.

Where do I find this configuration?

Go to **Configuration > System > Network > IP Address**.

To enable a bridge

Select the brXX checkbox above the interface pair.

To configure an interface address and netmask automatically

1. For the given interface or bridge, select either the **DHCP** checkbox for IPv4 networks or **SLAAC** for IPv6 networks.
2. If **DHCP** is selected, an IP address will be automatically assigned.
3. If **SLAAC** is selected for IPv6 networks, the following additional options are shown:
  - **Privacy Address**- Enable SLAAC privacy extension. Selecting this option will periodically change the automatically assigned IPv6 address.
  - **Gateway**- Assign an IPv6 gateway dynamically.

Which interfaces are DHCP enabled for each model?

| Model                                                        | Factory default DHCP enabled interface |
|--------------------------------------------------------------|----------------------------------------|
| 2000, 2061, 4000                                             | br1                                    |
| 5000, 6000, 6010                                             | eth0                                   |
| 2060, 4010, 4060, 4061, 6060, 6062, 8060, 8062, 10060, 10062 | eth1                                   |

To configure a static address

1. Enter an IPv4 or IPv6 address and netmask.
2. You can optionally add a comment describing how the interface is to be used in the **Comment** field.

To configure the gateway settings

Enter the address of your network's default IPv4 and IPv6 gateways.

### 0.18.3 Routes configuration

Static routes may need to be defined when access to external networks cannot be reached via the default gateway. This may be necessary so the appliance can connect to services such as DNS or NTP.

Routing table entries are shown for IPv4 and IPv6 networks. The destination, gateway, interface, source and state is shown for each route. Routing table entries can have multiple sources:

|           |                                                      |
|-----------|------------------------------------------------------|
| static    | A manually configured route.                         |
| interface | Derived from the addresses assigned to an interface. |
| SLAAC     | Assigned from SLAAC autoconfiguration.               |
| DHCP      | Assigned from DHCP autoconfiguration.                |

| IPv4 routes              |               |              |           |           |                                     |
|--------------------------|---------------|--------------|-----------|-----------|-------------------------------------|
|                          | Destination   | Gateway      | Interface | Source    | Active                              |
| <input type="checkbox"/> | default       | 172.16.1.254 | eth1      | static    | <input checked="" type="checkbox"/> |
|                          | 172.16.0.0/23 | 0.0.0.0      | eth1      | interface | <input checked="" type="checkbox"/> |

Remove Selected

| IPv6 routes |                       |                          |           |                 |                                     |
|-------------|-----------------------|--------------------------|-----------|-----------------|-------------------------------------|
|             | Destination           | Gateway                  | Interface | Source          | Active                              |
|             | 2001:44b8:62:690::/64 | ::                       | eth1      | SLAAC interface | <input checked="" type="checkbox"/> |
|             | default               | fe80::210:f3ff:fe0e:f4d0 | eth1      | SLAAC           | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | br10      | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth2      | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth20     | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth21     | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | br12      | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | br20      | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | brvm2     | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth1      | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth10     | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth11     | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth12     | interface       | <input checked="" type="checkbox"/> |
|             | fe80::/64             | ::                       | eth13     | interface       | <input checked="" type="checkbox"/> |

Remove Selected

| Add New Static Route |                                             |
|----------------------|---------------------------------------------|
| Destination:         | <input type="text"/> / <input type="text"/> |
| Gateway (Next Hop):  | <input type="text"/>                        |

Add Route

Screenshot 262: Routes configuration

- Destination The IPv4 or IPv6 address and netmask of the destination
- Gateway (Next Hop) The IPv4 or IPv6 address of the gateway (next hop).

### 0.18.4 DNS and domain names configuration

The DNS page is used to set a host name for your Exinda appliance and to configure the location of your DNS server(s). You can also configure domain names that can be used to resolve hostnames in other configuration screens.

The Exinda appliance hostname should be unique on the network. The DNS server setting may be dynamic, configured by the DHCP server, or it could be configured by entering one or more IP addresses of your DNS server(s).

### Static and Dynamic Name Servers

| IP Address | Active                              | Source     |
|------------|-------------------------------------|------------|
| 10.1.0.2   | <input checked="" type="checkbox"/> | configured |

### System Host Name

|               |                                                   |
|---------------|---------------------------------------------------|
| Host Name     | <input type="text" value="weber-exinda-monitor"/> |
| Primary DNS   | <input type="text" value="10.1.0.2"/>             |
| Secondary DNS | <input type="text"/>                              |
| Tertiary DNS  | <input type="text"/>                              |

#### NOTE

A valid DNS server is required for Edge Cache, system alerts, scheduled reports, firmware updates, license updates, and Anonymous Proxy updates

| Static and Dynamic Domain Names      |                                     |            |
|--------------------------------------|-------------------------------------|------------|
| Domain                               | Active                              | Source     |
| <input type="radio"/> wat.exinda.com | <input checked="" type="checkbox"/> | configured |

Remove Selected

| Add New Domain Name |                      |
|---------------------|----------------------|
| Domain Name         | <input type="text"/> |

Add New Domain Name

Where do I find this configuration?

Go to **Configuration > System > Network > DNS**.

To configure the appliance's hostname

1. In the **System Host Name** section, in the **Host Name** field, type the name for this appliance.
2. Click **Apply Changes**.

How to know if the DNS was configured by the DHCP server?

In the Static and Dynamic Name Servers section, there will be an IP address where the source is indicated to be dynamic. Dynamic means it was configured by the DHCP server.

To configure the location of the DNS servers

1. In the **System Host Name** section, type the IP addresses of your DNS servers in one or more of the **Primary DNS** field, **Secondary DNS** field, and **Tertiary DNS** field.
2. Click **Apply Changes**. The IP addresses entered will appear in the Static and Dynamic Name Servers section as configured.

To add a domain name

1. In the **Add New Domain Name** area, type the new domain name.
2. Click **Add New Domain Name**. The domain name is added to the Static and Dynamic Domain Names list. All manually added domain names are static.

To remove a domain name

1. In the **Static and Dynamic Domain Names** list, select the domain to remove. Only manually added domain names can be removed.
2. Click **Remove Selected**.

## 0.18.5 HTTP proxy configuration

Specify a HTTP proxy if you would like the appliance to access Exinda's server via HTTP proxy. Access to Exinda's HTTP server is required for firmware updates, license updates, and Anonymous Proxy updates. If you have SDP enabled, please ensure your proxy supports HTTPS.

Where do I find this configuration?

Go to **Configuration > System > Network > HTTP Proxy**.

To configure access to Exinda's server via HTTP proxy

1. Specify the hostname or IP address and HTTP proxy port of the HTTP proxy. IPv4 or IPv6 addresses can be specified.
2. Select the type of authentication for the HTTP proxy.
3. Type the **Username** and **Password** for the HTTP proxy.
4. To verify SSL certificates, clear the **Do not verify SSL certificates** checkbox.
5. Click **Apply Changes**.

## 0.18.6 Email configuration

An SMTP server is required for sending email from the Exinda appliance. The appliance can email scheduled reports, system alerts, and auto-support notifications. Initially, you must configure the connection to the SMTP server, and then manage the users who receive the system notifications.

| SMTP Server         |                                                    |
|---------------------|----------------------------------------------------|
| SMTP Server Name    | <input type="text" value="smtp.wat.exinda.com"/>   |
| SMTP Server Port    | <input type="text" value="25"/>                    |
| "From" Address      | <input type="text" value="bob.loblaw@exinda.com"/> |
| SMTP Domain Name    | <input type="text" value="localdomain"/>           |
| SMTP Authentication | <input type="checkbox"/>                           |

| Notify Recipients                              |                                     |                                     |                                     |
|------------------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Email Address                                  | Verbose                             | Info Emails                         | Failure Emails                      |
| <input type="checkbox"/> antonio.cucci@abc.com | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| <input type="checkbox"/> joseph.king@abc.com   | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

| Add New Notify Recipients |                                     |
|---------------------------|-------------------------------------|
| Email Address             | <input type="text"/>                |
| Verbose Detail            | <input checked="" type="checkbox"/> |
| Info Emails               | <input checked="" type="checkbox"/> |
| Failure Emails            | <input checked="" type="checkbox"/> |

### Configuring SMTP server settings

Use the following instruction to configure the SMTP server settings.

1. Go to **Configuration > System > Network > Email > SMTP Server**.
2. In the **SMTP Server Name** field, type the name.

#### NOTE

You can use IPv4 or IPv6 addresses, or DNS names.

3. In the **SMTP Server Port** field, type the port number.

#### NOTE

The default port number is 25.

4. In the **"From" Address** field, type the email address from which the system alerts and report notifications should be sent.
5. If authentication is required, select the **SMTP Authentication** checkbox, and provide the **Username** and **Password**.
6. If necessary, select the **Use Secure Sockets Layer (SSL)** checkbox.
7. Click **Apply Changes**.

### Testing the SMTP configuration

Use the following instructions to test the SMTP configuration.

1. Go to **Configuration > System > Network > Email > Add New Notify Recipients**.
2. Add your own email address and click **Add New Recipient**. The list in the "Notify Recipients" section above updates.
3. In the **Notify Recipients** section, click **Send Test Email to All**.

### Related Topics

- » [Adding notification email recipients](#)
- » [Removing notification email recipients](#)

### Adding notification email recipients

Use the following instructions to add new notification email recipients.

1. Go to **Configuration > System > Network > Email > Add New Notify Recipients**.
2. In the **Email Address** field, type the email address.
3. Select the types of notifications the user should receive:
  - **Verbose Detail**—Send detailed event emails to the user.
  - **Info Emails**—Send informational emails to the user.
  - **Failure Emails**—Send failure emails to the recipient.
4. Click **Add New Recipient**. The new recipients are added to the Notify Recipients list above.

## NOTE

The types of emails being received by a user cannot be modified. To change which emails a user receives, you must first delete the user, and then add the email address again with the appropriate types of notifications selected.

### Related Topics

- » [Testing the SMTP configuration](#)
- » [Removing notification email recipients](#)

### Removing notification email recipients

Use the following instructions to remove users from the list of notification email recipients.

1. Go to **Configuration > System > Network > Email > Notify Recipients**.
2. In the list, select the user to be deleted.
3. Click **Remove Recipients**. The user is removed from the list, and will no longer receive email notifications.

### Related Topics

- » [Testing the SMTP configuration](#)
- » [Adding notification email recipients](#)

## 0.18.7 SNMP configuration

The Exinda appliance allows data export to SNMP systems. Configure the SNMP settings or download the Exinda SNMP MIB.

| SNMP Configuration        |                                                                                   |
|---------------------------|-----------------------------------------------------------------------------------|
| SNMP                      | <input checked="" type="checkbox"/> Enable                                        |
| SNMP Traps                | <input checked="" type="checkbox"/> Enable                                        |
| SNMP Multiple Communities | <input checked="" type="checkbox"/> Enable                                        |
| Sys Contact               | <input type="text"/>                                                              |
| Sys Location              | <input type="text"/>                                                              |
| Read-Only Community       | <input type="text" value="public"/>                                               |
| Default Trap Community    | <input type="text" value="public"/>                                               |
| Download SNMP MIB         |  |

Apply Changes

**NOTE**

To disable or enable SNMP traps for system alerts, see For more information, refer to [Alerts](#) (page 630)..

### Configuring SNMP

Use the following instructions to configure SNMP.

1. Go to **Configuration > System > Network > SNMP > SNMP Configuration**.

| SNMP Configuration        |                                                                                   |
|---------------------------|-----------------------------------------------------------------------------------|
| SNMP                      | <input checked="" type="checkbox"/> Enable                                        |
| SNMP Traps                | <input checked="" type="checkbox"/> Enable                                        |
| SNMP Multiple Communities | <input checked="" type="checkbox"/> Enable                                        |
| Sys Contact               | <input type="text"/>                                                              |
| Sys Location              | <input type="text"/>                                                              |
| Read-Only Community       | <input type="text" value="public"/>                                               |
| Default Trap Community    | <input type="text" value="public"/>                                               |
| Download SNMP MIB         |  |

Apply Changes

2. Enable the following, as needed:

- SNMP
- SNMP Traps
- SNMP Multiple Communities

#### NOTE

When the Multiple Communities option is disabled, the Community list area does not appear.

4. In the **Sys Contact** field, specify the syscontact variable in MIB-II.

5. In the **Sys Location** field, specify the syslocation variable in MIB-II.

6. Type the **Read-only** and **Default Trap** community string.

#### NOTE

When the Read-only community is changed to have a value that does not match an existing community, a new SNMP community is added to the list.

7. Click **Apply Changes**.

## Removing an unwanted SNMP Community

Use the following instructions to remove an unwanted SNMP community.

1. Go to **Configuration > System > Network > SNMP > List of configured SNMP Communities**.

|                          | Community | Access Type |
|--------------------------|-----------|-------------|
| <input type="checkbox"/> | public    | Read-only   |

Remove Selected

2. In the list of **SNMP Communities** area, select the checkbox next to community entry and click **Remove Selected**.

## Downloading the SNMP MIB file

Use the following instructions to download the SNMP MIB file. The file contains additional monitoring information.

1. Go to **Configuration > System > Network > SNMP**.
2. Under **SNMP Configuration**, click **Download SNMP MIB** . The EXINDA-MIB.txt file downloads to the location you specify.

## Changing SNMP authentication for Admin user

Use the following instructions to change the SNMP authentication for the Admin user.

1. Go to **Configuration > System > Network > SNMP > SNMP v3 Admin User**.

| SNMP v3 Admin User      |                                                  |
|-------------------------|--------------------------------------------------|
| Admin User              | <input type="checkbox"/> Enable                  |
| Authentication Type     | SHA1 ▾                                           |
| Privacy Type            | AES-128 ▾                                        |
| Authentication Password | <input type="text"/> (leave blank to not change) |
| Privacy Password        | <input type="text"/> (leave blank to not change) |

Apply Changes

2. If you need to enable **Admin User**, select the checkbox.
3. From the **Authentication Type** spin-box, select either SHA1 or MD5.
4. From the **Privacy Type** spin-box, select either AES-128 or DES.
5. If necessary, change the **Authentication Password** by typing the new password.
6. If necessary, change the **Privacy Password** by typing the new password.
7. Click **Apply Changes**.

### Temporarily stopping the sending of SNMP traps

Use the following instructions to disable the sending of SNMP traps to the sink server.

| Trap Sinks     |           |         |         |
|----------------|-----------|---------|---------|
| Host           | Community | Version | Enabled |
| No trap sinks. |           |         |         |

Remove Trap Sink    Enable Trap Sink    Disable Trap Sink

1. Go to **Configuration > System > Network > SNMP > Trap Sinks**.
2. In the list, select the checkbox for server and click **Disable Trap Sink**.
3. To re-enable the server, select the server from the list and click or **Enable Trap Sink**.

## Removing Trap Sink servers

Use the following instruction to remove a trap sink server.

1. Go to **Configuration > System > Network > SNMP**.

| Trap Sinks     |           |         |         |
|----------------|-----------|---------|---------|
| Host           | Community | Version | Enabled |
| No trap sinks. |           |         |         |



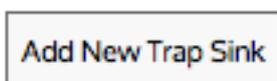
2. In the **Trap Sinks** area, select the server from the list and click **Remove Server**.

## Defining SNMP trap destinations

Use the following instructions to define where SNMP traps are sent.

1. Go to **Configuration > System > Network > SNMP**.

| Add New Trap Sink |                                      |
|-------------------|--------------------------------------|
| Server Address    | <input type="text"/>                 |
| Community         | <input type="text"/>                 |
| Trap Type         | v2c <input type="button" value="v"/> |



2. In the **Add New Trap Sink** area, specify the hostname or IP address of the SNMP trap sink server.

### TIP

You can specify IPv4 or IPv6 addresses, or a hostname.

3. Type the **Community** string for the SNMP trap sink server.
4. Select the appropriate SNMP trap type to send to the sink server.
5. Click **Add New Trap Sink**.

## 0.18.8 Integrate with Active Directory

### NOTE

You can configure the options in the Active Directory tab only after the Exinda AD Connector is installed and configured on a designated network server that has access to the Active Directory Server. You will see the Active Directory Server details on this tab only when the configuration is completed successfully.

Configuring Active Directory allows the Exinda Appliance to accept network users and groups from Active Directory (e.g logins, IP address, group membership) resulting in the ability to:

- » Expose Active Directory usernames in monitoring and reporting, no longer having to view users as IP addresses.
- » Use Active Directory groups and usernames in optimization policies, thereby implementing QoS and Optimization Polices based on individual users or entire groups.

To configure Active Directory, you need to install the Exinda AD Connector on a designated network server, configure various settings, and then proceed to configure the port and password settings on the Active Directory tab on each Exinda Appliance.

### How Active Directory Integration Works

Active Directory integration allows you to expose AD usernames within monitoring and reporting on the Exinda Appliance, rather than viewing the default IP Addresses. You can also use AD groups and usernames within optimization policies, allowing you to implement QoS and Optimization Policies based on individual users or entire groups.

Integration requires a proprietary Exinda AD Connector service installed onto a server in the network that has access to the Active Directory server. After configuration, the Connector functions as a gateway between the Active Directory Server and the Exinda Appliances to supply user and group information. As each user logs in using their Active Directory credentials, the information is gathered by the Connector and passed to the Exinda Appliances. Within the Monitor reports, IP Addresses are replaced by the user and group names where obtained from Active Directory.

### Integration Process

Complete the following tasks to connect the Exinda AD Connector to the Active Directory server, and to select the individual Exinda Appliances that will receive the AD information

### NOTE

Each installation of the Active Directory Connector can have a maximum of 20 Exinda Appliances connected to it. If there are more than 20 Exinda Appliances, you will need to install the connector on multiple Windows servers and divide the appliances across multiple instances of the Active Directory Connector. The instructions below step you through configuring a single Connector. Repeat these instructions if you are installing more than one instance of the Active Directory Connector.

1. Install the Exinda AD Connector. For more information, refer to [Install the Exinda AD Connector](#) (page 554).
  - a. Add the Exinda Appliances to the Exinda AD Connector. For more information, refer to [Add the Exinda Appliances to the Exinda AD Connector](#) (page 557).
  - b. Identify the Active Directory Server. For more information, refer to [Identify the Active Directory Server](#) (page 559).

c. Select the information sent between the Exinda appliance and the Active Directory server. For more information, refer to [Select the information sent between the Exinda appliance and the Active Directory server](#) (page 558).

d. The Exinda AD Connector port number. For more information, refer to [The Exinda AD Connector port number](#) (page 558).

2. Identifying users. For more information, refer to [Identifying users](#) (page 562).

3. Verify communication between the Active Directory server and the Exinda Appliance. For more information, refer to [Verify communication between the Active Directory server and the Exinda Appliance](#) (page 559).

#### NOTE

If you encounter any issues, see [Troubleshoot issues with Active Directory configuration](#).

## Configuration Options

After the integration process is successful, you can complete the following tasks to expose user names in monitoring reports, and to implement optimization policies based on user groups.

- » [View top Internal and External Users on the Network](#)
- » [Controlling Traffic Based on Users](#)

## Install the Exinda AD Connector

To integrate Active Directory with the Exinda Appliance, you need to install the Exinda AD Connector service on a Windows server that can then connect to the Active Directory server. Each Exinda AD Connector can talk to up to 20 Exinda appliances.

You can download the Active Directory Connector from the **Configuration > System > Network > Active Directory** tab on the Exinda Appliance. Click on the Microsoft Installer Executable link and save the installer to a location that can be accessed by all Windows servers in the network.

## Installation Requirements

- » The Exinda AD Connector is supported on the following platforms:
  - Windows Server 2003 SP2 - When the Active Directory server is running Windows Server 2003 R2, the Exinda AD Connector must be installed on the Active Directory server and cannot be installed on a remote server.
  - Windows Server 2008 SP2
  - Windows Server 2008 R2
  - Windows Server 2012
  - Windows Server 2012 R2
  - Windows Server 2016
- » The Exinda AD Connector requires .NET Framework 4.0.
- » Logon Auditing must be enabled on the Active Directory server to install the Exinda AD Connector.
- » The WMI service must be started on the Active Directory server and on the server where the Exinda AD Connector is installed.

» The Active Directory server and the server where the Exinda AD Connector is installed require the RPC Endpoint Mapper and LDAP ports open in your firewall. These ports are open by default. To verify your settings, see <http://support.microsoft.com/kb/179442>.

Complete the following tasks to continue...

Installing the Active Directory Connector

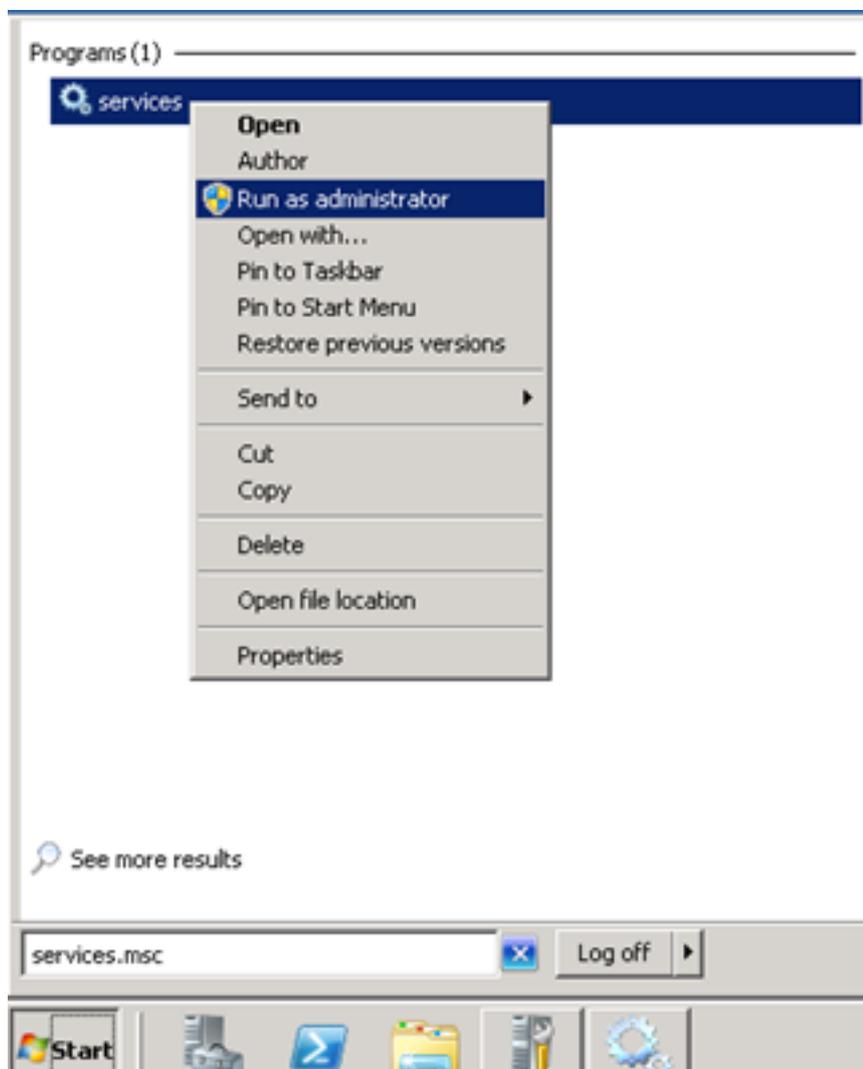
Providing the required permissions to the Exinda AD service

Providing the required permissions to the Exinda AD service

When installing the Exinda AD Connector on a server that is not a domain controller, ensure that the account in charge of running the service is an Active Directory domain admin account.

To provide the required permissions

1. Run **Services.msc** as an Administrator.



2. Find the entry for the **Exinda AD** service.

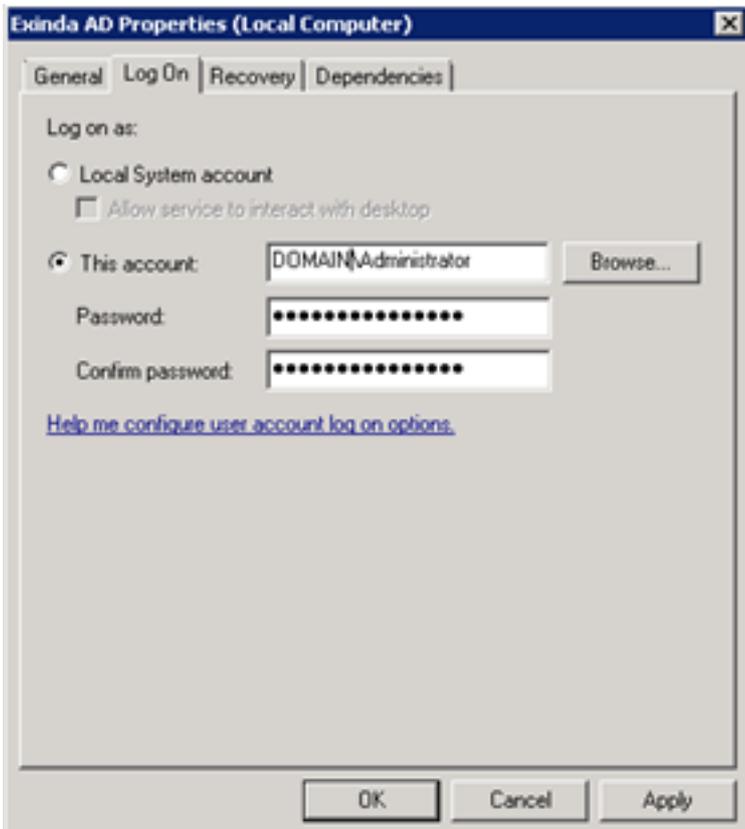
3. Right click on it and select **Properties**.

4. On the **Log On** tab, click **Browse** and select the domain and administrator account.

## NOTE

The domain and slash (\) are required.

5. Type the **Password** and confirm.



6. Click **OK** or **Apply** to save the changes.

7. Restart the service.

Complete the following tasks to continue...

Installing the Active Directory Connector

Installing the Exinda AD Connector

Use the following instructions to install the Exinda AD Connector.

Before you begin

Ensure that you have followed the Installation Requirements. For more information, refer to [Install the Exinda AD Connector](#) (page 554).

To install the Connector Service

1. On the server where the Exinda Active Directory Connector should be installed, run the installation file.
2. Read and accept the end-user license agreement, and then proceed through the screens, making the selections indicated below, and clicking Next where needed:

- Specify the directory where the Exinda Active Directory Connector should be installed.
- Select whether the Active Directory server is on **this server** or **another server**. If the connector is not installed on the server with Active Directory, type the IP address or hostname of the Active Directory server, and type the username and password of the Administrator account on the Active Directory server.

#### IMPORTANT

When the Active Directory server is running Windows Server 2003 R2, the Exinda Active Directory Connector must be installed on the Active Directory server and cannot be installed on a remote server.

#### IMPORTANT

When installing the Exinda AD Connector on a server that is not a domain controller, ensure that the account in charge of running the service is an Active Directory domain admin account. See *To ensure the Exinda AD service has the appropriate permissions* below.

- Optionally, type the Exinda appliance IP address or hostname, port number, and administrator password. This step is optional because you can add an Exinda Appliance after the Exinda Active Directory Connector is installed.
- In the **Include log entries newer than the specified age** field, specify the maximum age of log entries (in seconds) to be analyzed and sent to the Exinda Appliance when the Exinda Active Directory Connector service starts.

3. If any warnings are displayed, resolve the issues as specified in the dialog.

4. Click **Install**. Ensure **Launch Exinda Active Directory Connector** is selected, and click **Finish**.

After the installation is finished, the Exinda Active Directory Connector starts automatically and attempts to communicate with the configured Exinda appliance. When you first install the Exinda Active Directory Connector, it may take 24 hours or longer to obtain all user to IP address mappings as users progressively login.

Complete the following tasks to continue...

Providing the required permissions to the Exinda AD service

Add the Exinda Appliances to the Exinda AD Connector

Identify the Exinda Appliance using this Exinda AD Connector to retrieve user and group information.

#### NOTE

Each installation of the Active Directory Connector can have a maximum of 20 Exinda Appliances connected to it. If there are more than 20 Exinda Appliances, install the connector on multiple Windows servers and divide the appliances across multiple instances of the Active Directory Connector.

In the **Exinda AD Configuration Utility**, on the **Exinda Appliances** tab, type the IP Address and hostname of each appliance into an empty row. You will also need to type the Admin **password** for each appliance. The port number refers to the port that the Exinda AD Connector is using to communicate with the clients, and the Exinda Appliance must all use the same port number. The default port number of the Active Directory Client is 8015. For more information, refer to [The Exinda AD Connector port number](#) (page 558).

In the **Sync interval** field, identify how frequently the Exinda AD Connector contacts the Exinda Appliances to synchronize Active Directory user and group information. The default is 5 minutes.

## Related Topics

- » [The Active Directory Connector port number](#)

## The Exinda AD Connector port number

By default, port 8015 is used to communicate Active Directory information between the Exinda AD Connector and the connected Exinda Appliances. You should change the port number only if a conflict necessitates the change. If you change the port on the Connector, you must also change the port on each the Exinda Appliances.

### NOTE

Ensure that the firewall on the server running the Exinda AD Connector is configured to allow inbound and outbound traffic on the configured port.

## Changing the Exinda AD Connector port number

If necessary, you can change the AD Connector port number using the following instructions. Changing the port number requires that you do this on both the server hosting the AD Connector, and each of the Exinda appliances.

### To change the port number on the Exinda AD Connector

1. Launch the **Exinda AD Configuration Utility**.
2. Select the **Exinda Appliances** tab.
3. Type a new port number in the field. The default port number is 8015.

### To change the port number on each Exinda Appliance

1. Log into the Exinda Web UI.
2. Click **Configuration**, and from the System group, select **Network > Active Directory**.
3. Type the same port number you set above in the Exinda AD Configuration Utility.
4. Apply the changes.
5. Repeat these steps for each Exinda Appliance that communicates with this instance of the Exinda AD Connector.

### To determine if the port change was successful on the Exinda Appliance

Wait a few moments to ensure the information on the Active Directory tab updates with new information:

- » **IP Address**– The IP address of the server running the Exinda AD Connector.
- » **Windows Version**– The version of Windows on the Active Directory server.
- » **Version**– The Exinda AD Connector version.
- » **Agent Name**– The Exinda AD Connector name.
- » **Last Contact**– The last time the Active Directory server was contacted.

## Select the information sent between the Exinda appliance and the Active Directory server

Specify what information is sent between the Active Directory server and the Exinda appliance. When you first install the Exinda AD Connector, it may take a while to complete all user to IP address mappings as each user needs to logon.

## NOTE

User accounts that have been disabled on the Active Directory server are not included in the data sent to the Exinda Appliances

1. In the Exinda AD Connector, switch to the **AD Server** tab.
2. To send a list of users and groups to Exinda appliances when the service starts, select **Send Active Directory user and group information to Exinda appliances**. The list of users and groups that is sent to the appliance can be used to create user or group-based policy. If this is not selected, only logged on users will be available to your Exinda appliances. Information about groups will not be available. This information is obtained through an LDAP query against the Active Directory server.

## CAUTION

If there are multiple domain controllers, Sending the users/groups to the Exinda appliances on startup should only be selected on one of the domain controllers. For more information, refer to [Exinda Appliance Reboots Every Night](#) (page 733)..

3. To include user names in monitoring reports, allow the login history to be analyzed.
  - a. To enable this option, select **Analyze login history and send to Exinda appliance**. This information is obtained through a Windows Event Log query against the Active Directory server.
  - b. In the **Include log entries newer than the specified age** field, specify the maximum age of log entries (in seconds) to be analyzed and sent to the Exinda Appliance when the Exinda AD Connector service starts.
4. Click **OK**.

## Identify the Active Directory Server

The Exinda AD Connector can be installed on any server in the network that has access to the Active Directory server. If the Connector is installed somewhere other than on the Active Directory server, you must specify the location and authentication credentials of the Active Directory server.

## NOTE

You need to complete these instructions only if the Exinda AD Connector is NOT installed on the Active Directory server.

1. Launch the Exinda AD configuration utility, and switch to the AD Server tab.
2. Select **another server**, and then type the **IP Address** or the **hostname** of the Active Directory server.
3. To authenticate against the server, type the **username** and **password** of the Administrator account on the Active Directory server.

## Verify communication between the Active Directory server and the Exinda Appliance

To ensure the communication between the Active Directory server and the Exinda Appliance is successful, you can quickly check the Active Directory tab on the Exinda Appliance. Login to the Exinda Web UI. Click **Configuration**, and from the System group, select **Network > Active Directory**.

## NOTE

User accounts that have been disabled on the Active Directory server are not included in the data sent to the Exinda Appliances

Verify the Active Directory server is listed, and that the service is **Running**.

When the Exinda Appliance successfully communicates with the Active Directory Client, the following information is displayed in the table:

- » **Agent Name**– The Active Directory server name.
- » **IP Address**– The IP address of the Active Directory server.
- » **Version**– The Exinda Active Directory Windows client version.
- » **Windows Version**– The Active Directory server Windows version.
- » **Last Contact**– The last time the Active Directory server was contacted.

If the service is not visible on the list, run the Event Viewer program on your Active Directory server, and examine Windows logs:

1. From the **Start** menu select **Control Panel > Administrative Tools**.
2. Double-click **Services**, and verify the status of the **Exinda AD** service. If the service is stopped, restart the service.
3. In the **Windows Logs > Application** area, a "Service started successfully" message should be displayed from Exinda Networks Active Directory Connector.

If communications between the Active Director and the Exinda Appliance are failing, an error message from the Exinda Networks Active Directory Connector appears in these logs.

## Request updated user and group information from the Active Directory server

If the list of users and groups using the Active Directory client appears to be out of date, erase all username to IP address mappings and refresh the list sent from the Active Directory server.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Network**, and switch to the **Active Directory** tab.
6. To clear user, group, and login data from the appliance and requests an update from the Active Directory clients click **Renumerate**.

## Change the state of the Exinda AD Connector

Temporarily stop or disable the Active Directory integration to help with troubleshooting and to avoid errors when modifying the Exinda AD Connector settings.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Network**, and switch to the **Active Directory** tab.
6. Modify the state of the Active Directory service.

- To temporarily stop the Exinda AD Connector, click **Stop**.
- If you are experiencing issues with the Exinda AD Connector, **Restart** the service.
- If you no longer need the Exinda AD Connector running, click **Disable**.
- If the service has been disabled, to start it again click **Enable**.

### Exclude specific usernames from reports

You may have user accounts that should not be linked to IP addresses when reporting on the Exinda appliance, such as the account used for signing SMB traffic. Configure the Exinda AD Connector to prevent the IP address to username mapping being sent to the Exinda Appliance.

#### Before you begin...

You need to understand the process of:

- » Requesting updated user and group from Active Directory. For more information, refer to [Request updated user and group information from the Active Directory server](#) (page 560).
- » Restarting the Active Directory service. For more information, refer to [Change the state of the Exinda AD Connector](#) (page 560).

#### To exclude usernames

1. From the **Start** menu, click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
2. Select the **Excluded Users** tab.
3. Click in the **Ignored User**s area and type the full username of each user to ignore. Usernames are case sensitive. If the Active Directory has the user Domain/Test.User, and the excluded list has the user as Domain/test.user, the traffic is not excluded.

#### NOTE

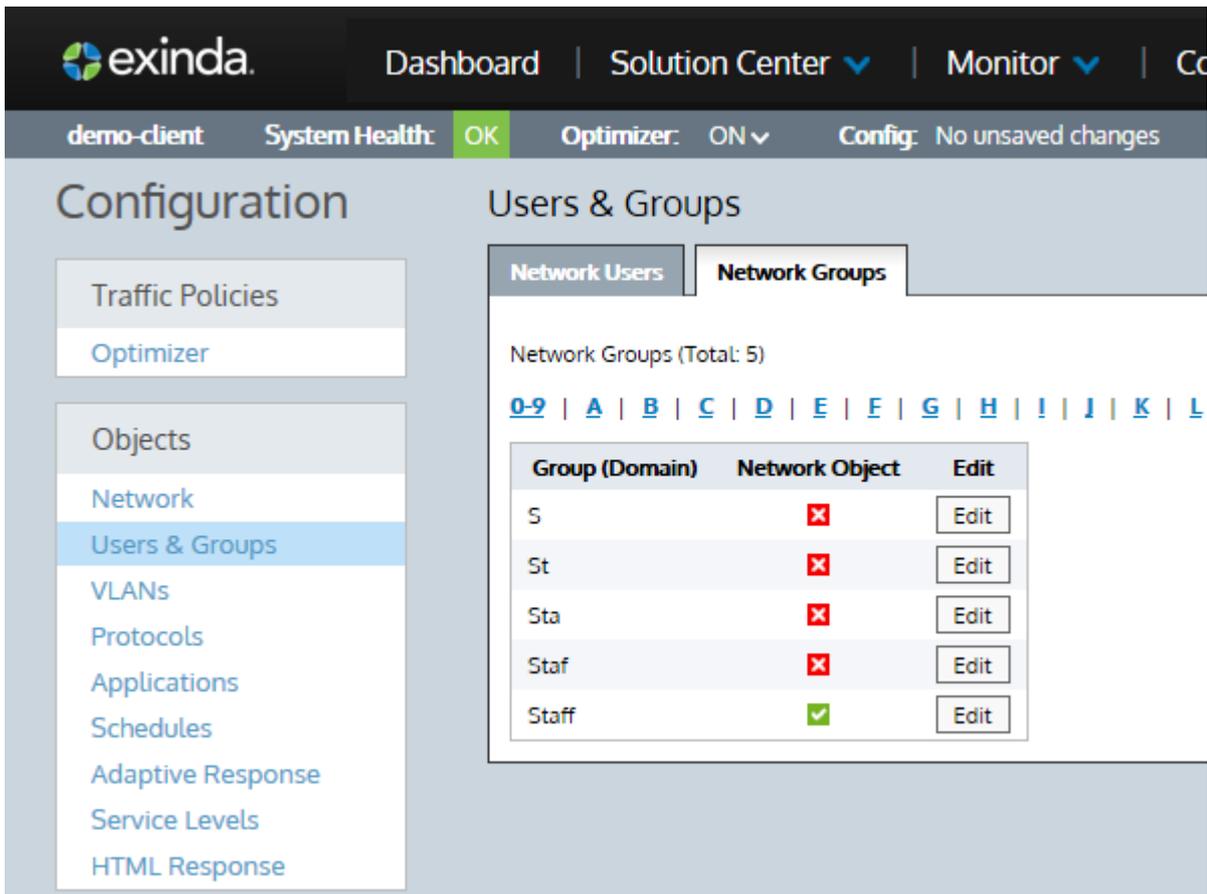
Regardless of the case of usernames in Active Directory, the Exinda Appliance displays the usernames with the first name capitalized and the surname in lower case; for example Domain/Test.user. Do not use the value in the Exinda Appliance when adding a username to the Excluded list.

4. Click **Apply**.
5. Request updated user and group information from the Active Directory server.
6. Restart the Active Directory service.

### Use Adaptive Response with Active Directory

In the last example, a static Network Object was used as the source of IPs. It is also possible to use a Dynamic Network Object mapped from an Active Directory group as a source.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > Objects > Users & Groups**.
6. Click **Edit** to edit the required user group.



7. Check **Map to Network Object** and **Ignore Domain** options.

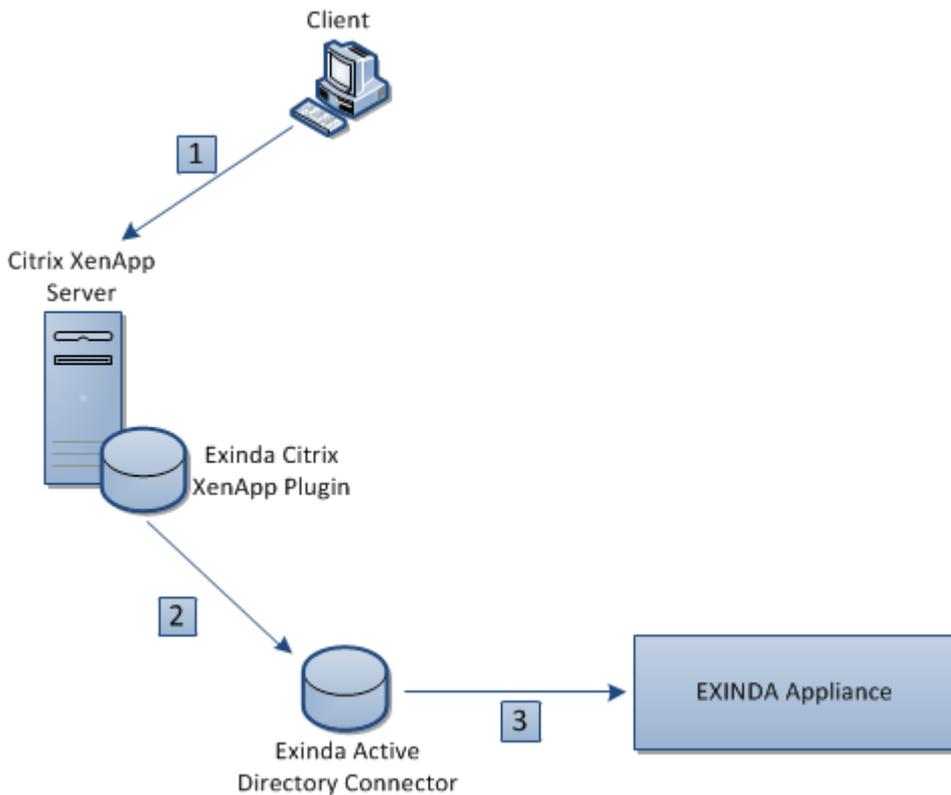
8. Click **Apply**.

A Network Object named similar to the user group name' is created that contains all IPs in the Active Directory 'Student' group. This Network Object can be used when creating an Adaptive Response rule exactly as for the previous example.

### Identifying users

A Citrix XenApp server hosts a virtual desktop with pre-installed software that users with the correct credentials can access as needed. This allows the company to provide access to commonly used software without having to maintain and upgrade installations on each client computer in the network.

Because the Citrix XenApp server is treated as a single IP address by the Exinda appliance, and the IP address of the clients connecting to the server are ignored, the Exinda Appliance cannot include the names of users who are accessing the applications on the XenApp server.



When a user on a client computer logs into a Citrix XenApp server (1), their IP address and user name are captured by the Exinda Citrix XenApp Plugin and sent on to the Exinda AD Connector (2). The connector then sends the user name and IP address of the XenApp user to the Exinda Appliance to include in reports (3).

Install and configure the Exinda Citrix XenApp Plugin to identify activity by specific users on the XenApp server.

### Install the Exinda Citrix XenApp Plugin

The Exinda Citrix XenApp Plugin sends the IP address and username of the user using the application on the XenApp server to the Exinda AD Connector so the user names can be displayed in reports on the Exinda Appliances. The Exinda Citrix XenApp Plugin must be installed on each Citrix XenApp server in the network.

#### NOTE

The Exinda Citrix XenApp Plugin is supported on Citrix XenApp Servers version 6.0.

1. Download the installer the Exinda Appliance.
  - a. Click **Configuration > System > Network**, and switch to the **Active Directory** tab.
  - b. Download the **Microsoft Installer Executable**.
2. Save the Exinda Citrix XenApp Plugin install to a location that can be accessed by the Citrix XenApp server.
3. On the server where the Exinda Citrix XenApp Plugin should be installed, locate and double-click installation file.
4. At the Welcome dialog, click **Next**.
5. Specify the directory where the Exinda Citrix XenApp Plugin should be installed and click **Next**.
6. Read the End-User License Agreement. Select **I Agree** and click **Next**.

7. To confirm the installation, click **Next**. The Exinda Citrix XenApp Plugin is installed.
8. When the installation is completed, click **Close**.

### Add the Exinda AD Connector to the Exinda Citrix XenApp Plugin

To ensure that user activity on the Citrix XenApp server is reported on the Exinda appliance, add the connection details for the Exinda AD Connector to the Exinda Citrix XenApp Plugin.

1. Open the Exinda Citrix XenApp Plugin.
2. Select the **Synchronization** tab and double-click in the **Location** area of the first blank line.
3. Type the IP address or hostname and port number of the computer where the Exinda AD Connector is installed.

#### NOTE

The port number used to communicate between the Exinda AD Connector and the Exinda Citrix XenApp Plugin cannot be the same as the port number used to communicate between the Exinda AD Connector and the Exinda Appliances.

4. In the **Sync Interval** field, define how frequently the Exinda AD Connector sends XenApp server user information to the Exinda AD Connector. The default is 1 minute.
5. Click **Apply**.

### Capture the Exinda Citrix XenApp Plugin Activity in a Log File

Depending on the logging level selected, the Exinda Citrix XenApp Plugin records various types of data in a log file. The available log levels include Error, Warning, Info, and Verbose. By default, the log sensitivity is Warning. The location of the log file and the level of detail recorded in the log file are configurable.

1. Open the Exinda Citrix XenApp Plugin.
2. On the **AD Server** tab, specify the location where log files should be stored.
3. Switch to the **Console** tab and select the level of messages that are recorded in the log file from the **Log Sensitivity** list.
4. Click **Apply**.
5. To view the contents of the log, on the **Console** tab click **Open Log**.

### Change the Exinda Citrix XenApp Plugin Port Number

Identify the port on which the Exinda AD Connector is communicating to the connected Exinda Citrix XenApp Plugins. The default port number is 8016.

Step 1: Change the port number on the Exinda Citrix XenApp Plugin.

1. From the **Start** menu click **All Programs > Exinda Networks > Exinda Citrix XenApp Plugin Configuration**.
2. Switch to the **Synchronization** tab.
3. Double-click the port number for the appropriate Exinda AD Connector and type the new port number in the field.
4. Click **OK**.

Step 2: Change the port number on the Exinda AD Connector.

1. From the **Start** menu click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
2. Switch to the **XenApp** tab.

3. Type the port number in the field.
4. Click **OK**.

### Request Updated User Information from the Exinda Citrix XenApp Plugin

If the synchronizations of the user data between the Exinda Citrix XenApp Plugin and the Exinda AD Connector are infrequent, trigger the Exinda Citrix XenApp Plugin to send the data to the Exinda AD Connector immediately.

1. From the **Start** menu, click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
2. Switch to the **XenApp** tab.
3. Click **Renumerate**.

The latest data is sent from the Exinda Citrix XenApp Plugin to the Exinda AD Connector.

### Adding a new application

Use the following instruction to add a new application.

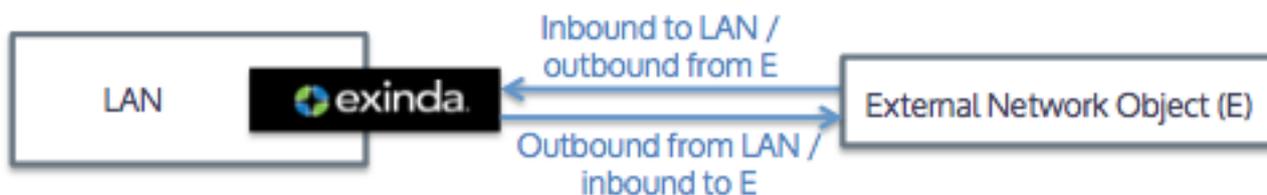
1. Click **Configuration > Objects > Applications > Applications**.
2. In the **Add New Application** area, type a name for the new application.
3. Define an application to be based on one of the following:

- L7 signature
- L7 signature + ports or protocols
- Network object + ports or protocols
- Network object
- Ports or protocols

#### NOTE

Network objects cannot be used in conjunction with a layer 7 signature.

3. Select the **Network Object** for the application. If the network object is internal, then traffic inbound to the LAN with the network object as a destination will be matched to this application, and traffic outbound from the LAN with the network object as the source will be matched to this application. If the network object is external, then traffic inbound to the LAN with the network object as a source will be matched to this application, and traffic outbound from the LAN with the network object as the destination will be matched to this application.



4. Select the **L7 Signature** for the application. Some layer 7 signatures have additional options that allow you to define application objects based on specific parts of that L7 signature. If a layer 7 signature is selected, specify the parameters for the signature.

### EXAMPLE

To create an application object that matches traffic to and from the Exinda.com website, in the **L7 Signature** field, select **http --->**, **host**, and type **exinda.com**.

5. In the **Ports/Protocols** controls, specify either TCP ports/port ranges, UDP ports/port ranges, or a layer 3 protocol. Multiple ports and port ranges can be specified at the same time by comma separating values.
6. Click the **Add New Application** button.

### Related topic

[What L7 signature options are there?](#)

[Example: How to create a custom application based on the HTTPS protocol](#)

### What L7 signature options are there?

Some Layer 7 signatures have additional options that allow you to define application objects based on specific parts of that L7 Signature. When configuring new application object, the L7 signatures followed by '--->' in the drop-down list have additional options. Most provide options that you simply select from. Some require a selection plus additional information. The following table explains the various options that require more than simply picking an option.

### NOTE

Citrix-based sub-types are no longer supported.

| Layer 7 Signature      | Sub-Type | Description                                                                                                                 |
|------------------------|----------|-----------------------------------------------------------------------------------------------------------------------------|
| (direct download link) | host     | Allows you to define an Application Object based on the 'host' field in the HTTP header.                                    |
| flash                  | host     | Allows you to define an Application Object based on the 'host' field in the HTTP header (where flash is running over http). |

| Layer 7 Signature                                                           | Sub-Type                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
|-----------------------------------------------------------------------------|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|--------|---------|--------|------------------------|-------------------------|----------------|----------------------------|--------------------------|--------------------|----------------------------|--------------------------|----------------------------|----------------------------|--------------------------|----------------------------------------------------------------------|----------------------------|------------------------|-----------------------------------------------------------------------------|----------------------------|
| http                                                                        | content_type               | Allows you to define an Application Object based on the 'content-type' field in the HTTP header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
|                                                                             | file                       | Allows you to define an Application Object based on the filename requested in the HTTP URL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
|                                                                             | host                       | Allows you to define an Application Object based on the 'host' field in the HTTP header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
|                                                                             | method                     | Allows you to define an Application Object based on the HTTP method (e.g. GET PUT HEAD DELETE).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
|                                                                             | user_agent                 | Allows you to define an Application Object based on the 'user-agent' field in the HTTP header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
|                                                                             | advanced                   | <p>Define custom criteria with the following syntax:</p> <ul style="list-style-type: none"> <li>» A string literal is enclosed in quotes (").</li> <li>» A backslash can be included in the string by escaping it with another backslash (\).</li> <li>» Keywords are bare (common_name) with no quotes.</li> <li>» Keywords are bare (host) with no quotes.</li> <li>» Grouping is supporting using parenthesis</li> <li>» Operators supported are or and and and has higher precedence than or</li> <li>» The comparison operators that are available are:</li> </ul> <table border="1"> <thead> <tr> <th>Description</th> <th>Syntax</th> <th>Example</th> </tr> </thead> <tbody> <tr> <td>equals</td> <td>&lt;keyword&gt;<br/>= &lt;value&gt;</td> <td>host =<br/>"example.com"</td> </tr> <tr> <td>does not equal</td> <td>&lt;keyword&gt;<br/>!=<br/>&lt;value&gt;</td> <td>host !=<br/>"example.com"</td> </tr> <tr> <td>contains substring</td> <td>&lt;keyword&gt;<br/>=%<br/>&lt;value&gt;</td> <td>host =%<br/>"example.com"</td> </tr> <tr> <td>does not contain substring</td> <td>&lt;keyword&gt;<br/>!%<br/>&lt;value&gt;</td> <td>host !%<br/>"example.com"</td> </tr> <tr> <td>Right side is a regular expression and it matches the full left side</td> <td>&lt;keyword&gt;<br/>=~<br/>&lt;value&gt;</td> <td>host =~<br/>"example.*"</td> </tr> <tr> <td>Right side is a regular expression and it does not match the full left side</td> <td>&lt;keyword&gt;<br/>!~<br/>&lt;value&gt;</td> <td>host !~<br/>"example.*"</td> </tr> </tbody> </table> <ul style="list-style-type: none"> <li>» Regular expressions use the perl syntax</li> <li>» The keywords for HTTP are: host, file, user_agent, content_type, method, content_len and encoding</li> </ul> <p>Examples:</p> <ul style="list-style-type: none"> <li>» (url =% "index" or file =% "login") and host =% "example.org" and content_type.case = "MyContentType"</li> <li>» (host =% "facebook.com" and file !=% "cgi-bin/abcd") or host =% "facebook2.com"</li> </ul> | Description | Syntax | Example | equals | <keyword><br>= <value> | host =<br>"example.com" | does not equal | <keyword><br>!=<br><value> | host !=<br>"example.com" | contains substring | <keyword><br>=%<br><value> | host =%<br>"example.com" | does not contain substring | <keyword><br>!%<br><value> | host !%<br>"example.com" | Right side is a regular expression and it matches the full left side | <keyword><br>=~<br><value> | host =~<br>"example.*" | Right side is a regular expression and it does not match the full left side | <keyword><br>!~<br><value> |
| Description                                                                 | Syntax                     | Example                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
| equals                                                                      | <keyword><br>= <value>     | host =<br>"example.com"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
| does not equal                                                              | <keyword><br>!=<br><value> | host !=<br>"example.com"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
| contains substring                                                          | <keyword><br>=%<br><value> | host =%<br>"example.com"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
| does not contain substring                                                  | <keyword><br>!%<br><value> | host !%<br>"example.com"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
| Right side is a regular expression and it matches the full left side        | <keyword><br>=~<br><value> | host =~<br>"example.*"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |
| Right side is a regular expression and it does not match the full left side | <keyword><br>!~<br><value> | host !~<br>"example.*"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                        |                                                                             |                            |

| Layer 7 Signature                                                           | Sub-Type                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
|-----------------------------------------------------------------------------|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|--------|---------|--------|------------------------|-------------------------|----------------|----------------------------|--------------------------|--------------------|----------------------------|--------------------------|----------------------------|----------------------------|--------------------------|----------------------------------------------------------------------|----------------------------|---------------------------|-----------------------------------------------------------------------------|----------------------------|
| mpeg                                                                        | host                       | Allows you to define an Application Object based on the 'host' field in the HTTP header (where mpeg is running over http).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| quicktime                                                                   | host                       | Allows you to define an Application Object based on the 'host' field in the HTTP header (where quicktime is running over http).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| silverlight                                                                 | host                       | Allows you to define an Application Object based on the 'host' field in the HTTP header (where silverlight is running over http).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| ssl                                                                         | common_name                | Allows you to define an Application Object based on the 'common name' field in the SSL certificate.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
|                                                                             | advanced                   | <p>Define custom criteria with the following syntax:</p> <ul style="list-style-type: none"> <li>» A string literal is enclosed in quotes (").</li> <li>» Internal quotes can be escaped with the backslash (\) character.</li> <li>» A backslash can be included in the string by escaping it with another backslash (\).</li> <li>» Keywords are bare (common_name) with no quotes.</li> <li>» Grouping is supporting using parenthesis</li> <li>» Operators supported are OR and AND. AND has higher precedence than OR.</li> <li>» The keywords for SSL are common_name (cn) and organization_name (o)</li> <li>» The comparison operators that are available are:</li> </ul> <table border="1"> <thead> <tr> <th>Description</th> <th>Syntax</th> <th>Example</th> </tr> </thead> <tbody> <tr> <td>equals</td> <td>&lt;keyword&gt;<br/>= &lt;value&gt;</td> <td>common_name =<br/>"John"</td> </tr> <tr> <td>does not equal</td> <td>&lt;keyword&gt;<br/>!=<br/>&lt;value&gt;</td> <td>common_name !=<br/>"John"</td> </tr> <tr> <td>contains substring</td> <td>&lt;keyword&gt;<br/>=%<br/>&lt;value&gt;</td> <td>common_name =%<br/>"John"</td> </tr> <tr> <td>does not contain substring</td> <td>&lt;keyword&gt;<br/>!%<br/>&lt;value&gt;</td> <td>common_name !%<br/>"John"</td> </tr> <tr> <td>Right side is a regular expression and it matches the full left side</td> <td>&lt;keyword&gt;<br/>=~<br/>&lt;value&gt;</td> <td>common_name =~<br/>"John*"</td> </tr> <tr> <td>Right side is a regular expression and it does not match the full left side</td> <td>&lt;keyword&gt;<br/>!~<br/>&lt;value&gt;</td> <td>common_name !~<br/>"John*"</td> </tr> </tbody> </table> <ul style="list-style-type: none"> <li>» Regular expressions use the perl syntax</li> </ul> | Description | Syntax | Example | equals | <keyword><br>= <value> | common_name =<br>"John" | does not equal | <keyword><br>!=<br><value> | common_name !=<br>"John" | contains substring | <keyword><br>=%<br><value> | common_name =%<br>"John" | does not contain substring | <keyword><br>!%<br><value> | common_name !%<br>"John" | Right side is a regular expression and it matches the full left side | <keyword><br>=~<br><value> | common_name =~<br>"John*" | Right side is a regular expression and it does not match the full left side | <keyword><br>!~<br><value> |
| Description                                                                 | Syntax                     | Example                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| equals                                                                      | <keyword><br>= <value>     | common_name =<br>"John"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| does not equal                                                              | <keyword><br>!=<br><value> | common_name !=<br>"John"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| contains substring                                                          | <keyword><br>=%<br><value> | common_name =%<br>"John"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| does not contain substring                                                  | <keyword><br>!%<br><value> | common_name !%<br>"John"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| Right side is a regular expression and it matches the full left side        | <keyword><br>=~<br><value> | common_name =~<br>"John*"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| Right side is a regular expression and it does not match the full left side | <keyword><br>!~<br><value> | common_name !~<br>"John*"                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
|                                                                             | organization_name          | Allows you to define an Application Object based on the 'organization' name field in the SSL certificate.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
|                                                                             | spdy                       | This field should remain empty as any values typed here are ignored.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| rtp                                                                         | codec                      | Allows you to define an Application Object based on the 'codec' used in a RTP stream.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |
| windowsmedia                                                                | host                       | Allows you to define an Application Object based on the 'host' field in the HTTP header (where windowsmedia is running over http).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |             |        |         |        |                        |                         |                |                            |                          |                    |                            |                          |                            |                            |                          |                                                                      |                            |                           |                                                                             |                            |

## Example: How to create a custom application based on the HTTPS protocol

Get the common name of the (https) SaaS site and create an application using the ssl L7 signature with the common name.

1. Go to the site that you are interested in.
2. In the address bar of most browsers, click on https or the lock symbol.
3. Show the certificate details.
4. Copy the common name shown in the certificate details.
5. Go to **Configuration > Objects > Applications**.
6. In the **L7 Signature** field, select 'ssl --->'
7. In the field beside the L7 Signature, select **common name**.
8. Enter the common name of the site that you got from the certificate in the browser.

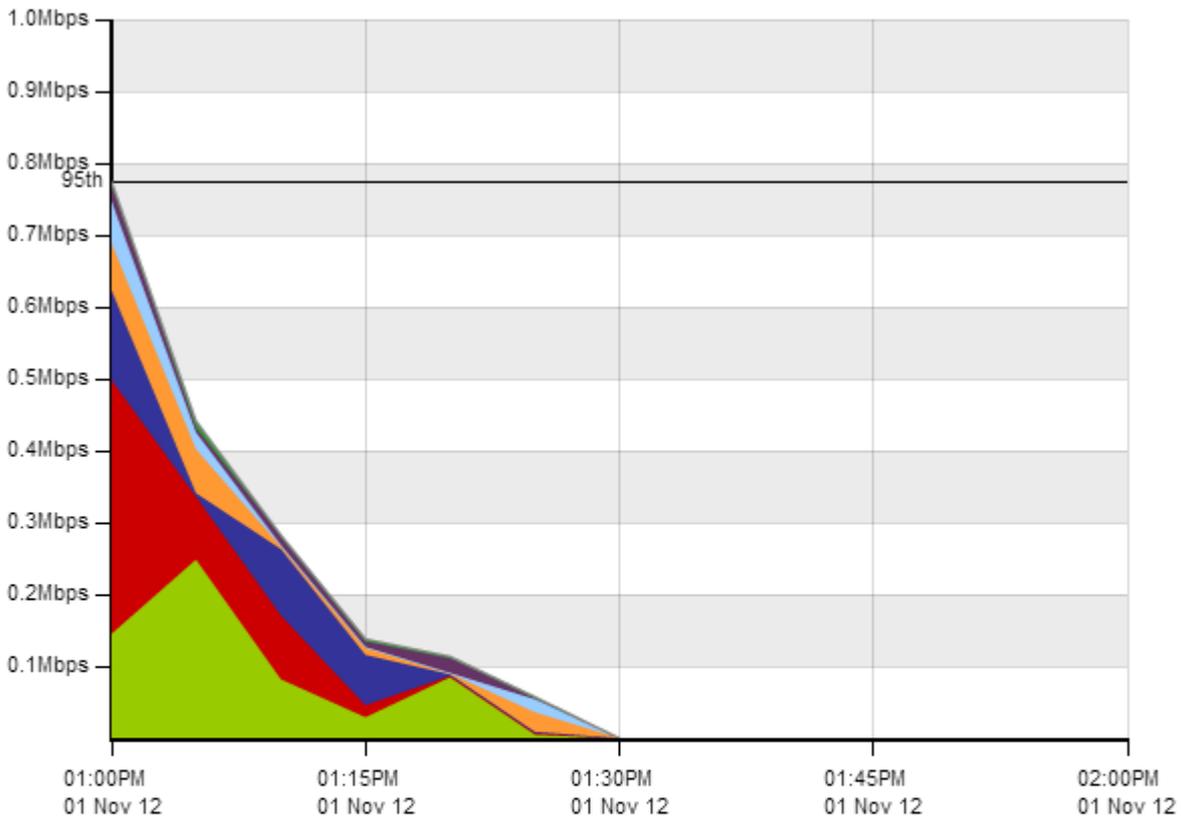
## Top Internal and External Users on the Network

The Network - Users (Internal) and Users (External) reports displays the top users sending traffic through the network.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Monitor > Network**.
6. In the Select Graph to Display list, select **Users - Internal** or **Users - External**.
7. Set the Time Period Reflected in the Report. For more information, refer to [Setting the time period for a report](#) (page 284). After the date range is select, the graphs and charts are immediately updated.
8. Remove specific types of traffic from the graph by deselecting their checkbox in the legend below the graph.
9. To determine what the size of your WAN link should be configured to, from the **Select Percentile Marker to Display** select **95th**.

Use the 95th percentile mark for throughput speed to configure your WAN link.

Throughput for Top 10 Inbound Users - Internal LAN



|                                     | Name           | Total Data (MB) | Throughput Max (Mbps) | Throughput Avg (Mbps) |
|-------------------------------------|----------------|-----------------|-----------------------|-----------------------|
| <input checked="" type="checkbox"/> | EXANET\Brad    | 8.866           | 0.249                 | 0.020                 |
| <input checked="" type="checkbox"/> | EXANET\Dale    | 14.765          | 0.354                 | 0.033                 |
| <input checked="" type="checkbox"/> | EXANET\Jan     | 9.689           | 0.125                 | 0.022                 |
| <input checked="" type="checkbox"/> | EXANET\Micheal | 4.834           | 0.065                 | 0.011                 |
| <input checked="" type="checkbox"/> | EXANET\Ian     | 0.228           | 0.008                 | 0.001                 |
| <input checked="" type="checkbox"/> | EXANET\Vince   | 0.152           | 0.002                 | 0.000                 |

### Dual Bridge Bypass

When mapping your circuit to multiple bridges there are two scenarios to consider: backhauling traffic through your Exinda appliance and link aggregation through your Exinda appliance. The dual bridge bypass feature ensures that traffic is processed properly in these scenarios. Note that "dual" really means "multi" as traffic can pass through two or more bridges. There are separate dual bridge bypass settings for acceleration and for monitoring. Note that by default both acceleration and monitor for dual bridge bypass are enabled.

Backhauled traffic flows between distributed sites and the Internet via a centralized backbone, which is typically at headquarters. This means that traffic may go through your Exinda appliance at headquarters twice: from the source through the Exinda appliance, turning around at a router, back through the Exinda appliance, and on to the destination. This is problematic for accelerated traffic because you do not want to re-accelerate the traffic. The dual bridge bypass feature allows each bridge to treat traffic differently, so that the traffic is accelerated on one bridge on the way in and bypasses the acceleration handling from the second bridge on the way out.

Link aggregation maps multiple bridges to act as a single link. You may have a single router that is passing data through multiple bridges, where some packets may go over one bridge and some over the other. Or you may have multiple

upstream service providers with asymmetric routes, where traffic may go out one bridge and return via another bridge. In this case, dual bridge bypass needs to be disabled so that traffic is handled the same regardless of which bridge it arrives on.

For acceleration, dual bridge bypass controls where packets are intercepted and passed to acceleration and edge cache.

» When enabled, all acceleration processing is performed on one bridge only. This is desirable for backhaul settings but will cause traffic problems if there is asymmetric routing or link aggregation in use. Consider traffic passing through the appliance from the WAN to the Internet as: WAN <-> br10 <-> router <-> br20 <-> internet

- For an incoming accelerated connection, the acceleration processing happens on the first WAN interface to see the SYN. In this case, where the connection is going to the internet, the accelerated connection would be processed on br10. That same traffic when seen on br20 would then be passed though untouched (the traffic 'bypasses' br20).
- For an outgoing accelerated connection, the acceleration processing happens on the first bridge with a matching acceleration policy. In this case, where the connection is from the Internet, if there is an acceleration policy on br20, the acceleration policy happens on br20 and the traffic would be untouched on br10. Or if the acceleration policy was on br10, then the traffic would be passed through br20 untouched and accelerated on br10.

» When disabled, acceleration processing is performed on whatever bridge it arrives on.

This is desirable for asymmetric routing or link aggregation, where you want traffic to switch bridges. However, in an accelerated backhauled scenario, the accelerated traffic cannot switch bridges and still function properly. The SYN from the client must be processed on the same bridge as the SYN/ACK from the server.

For monitoring, dual bridge bypass controls how flows are tracked and how flows appear in the real-time monitor.

» When enabled, the same flow is tracked separately on each bridge. This allows you to see the different policies and transfer rates that are being applied on each bridge in the real-time monitor.

This may be desirable when you have dual bridge bypass for acceleration enabled so that you can see all the policies applied to the flows. If traffic is both accelerated and edge cached you see the two separate flows in real-time, but they will not be coloured yellow and blue to indicate acceleration and edge cache. They will both be coloured the one of the two colours.

If you see the (A) symbol in the real-time monitor and you get asymmetric route warnings, either the appliance is not seeing half of the flow or the other half of the flow is on another bridge, where the flow is bypassing acceleration processing on the second bridge and therefore you need to disable dual bridge bypass for acceleration for it to work properly.

If you have two circuits configured on the Exinda such that traffic between the headquarters appliance and the branch appliance goes through one circuit and the traffic between the headquarters appliance and the internet goes through a second circuit, then the traffic for a single flow will be counted in the appropriate virtual circuit for each circuit.

» When disabled, the same flow through two different bridges will be shown as one merged flow. The real-time monitor will only show the policy of the last packet processed.

When doing load balancing, failover, link bonding, or seeing asymmetric routes (either locally or in a HA cluster), you will want to turn off the monitor dual bridge bypass function to get back to a merged view of all flows.

## Modifying the Dual Bridge Bypass settings

To modify acceleration settings:

1. Go to **Configuration > System > Optimization > TCP**.
2. Modify the **Dual-bridge bypass** setting. If checked, it is enabled.

To modify monitoring settings:

1. Go to **Configuration > System > Setup > Monitoring**.
2. In the **Monitoring Options** section, modify the **Dual-bridge bypass** setting. If checked, it is enabled.

## Controlling Traffic based on Users and Groups

After reviewing the traffic patterns of the users, it may be necessary to implement optimization policies to ensure a positive user experience for key applications or traffic types. By limiting the traffic usage for a specific group of users, network availability can be increased for other user groups.

### NOTE

Active Directory must be configured before optimization policies can target specific users and groups. For more information, refer to [Integrate with Active Directory](#) (page 553).

## Optimize Traffic Based on Users and Groups

Create policies that affect the traffic based on the source or destination host.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > Optimizer > Policies**.
6. Type a name for the policy.
7. Set the required bandwidth and acceleration parameters.
8. In the Filter Rules area, select the network user or network group object in the Host source and destination fields, and specify the ToS/DSCP or Application traffic to be affected.
9. Click **Create New Policy**.

### NOTE

Once the desired policies are in place on all Exinda appliances, restart the Optimizer. On the appliance status bar, click **Restart**.

**Optimizer Status : On (Restart / Stop)**

## 0.18.9 IPMI Configuration

The Intelligent Platform Management Interface (IPMI) is a specification for remote server management. It is a way to manage a computer by using a network connection to the hardware rather than to an operating system or login shell. This allows you to manage a computer that may be powered off or is otherwise unresponsive. An IPMI-enabled server contains a separate, dedicated micro-controller, called Baseboard Management Controller (BMC), to perform operations such as remote power off power on, or to access the console (serial-over-LAN). The data is made available over the LAN using either a shared or dedicated NIC. The IPMI interface must be assigned a dedicated IPv4 address.

To configure the appliance to be managed via IPMI, IPMI must be enabled and the IPv4 address and gateway must be specified, as well as the administrator credentials. You can manually set these values or you may use DHCP to configure the IPv4 address. Note that although the IPMI LAN interface may be the same as the appliance's eth1 hardware interface, a separate IPv4 address is required.

The appliance can be managed on the specified IPv4 address either from another Exinda appliance or by a connection to the specified IPv4 address (e.g. <http://10.0.100.238>, assuming a web interface is supported). When a web interface is not supported, a command line tool such as `ipmitool` may be used.

If your management network is on a separate VLAN and you want the IPMI interface to also be on that VLAN, you can enable VLAN support and specify an 802.1Q VLAN ID.

| IPMI Network Settings |                                                                             |
|-----------------------|-----------------------------------------------------------------------------|
| Enable                | <input checked="" type="checkbox"/>                                         |
| DHCPv4                | <input type="checkbox"/>                                                    |
| IPv4 Address          | <input type="text" value="10.0.100.238"/> / <input type="text" value="16"/> |
| IPv4 Gateway          | <input type="text" value="10.0.0.1"/>                                       |
| VLAN Enable           | <input type="checkbox"/>                                                    |
| VLAN ID               | <input type="text" value="1"/>                                              |
| Admin User            | admin                                                                       |

Screenshot 263: Configure the appliance to be managed via IPMI

Where do I find this configuration?

Go to **Configuration > Objects > Adaptive Response > Network Objects**.

To configure IPMI settings

Using a static IP address and gateway:

1. To enable IPMI management, select the **Enable** checkbox.
2. Enter the **IPv4 Address** and netmask for the IPMI interface.
3. Enter the **IPv4 Gateway**.
4. Enable and type the **VLAN ID**, if desired. The VLAN ID should be an 802.1Q VLAN ID.
5. Click **Apply Changes**.

Using DHCP:

| IPMI Network Settings |                                     |
|-----------------------|-------------------------------------|
| Enable                | <input checked="" type="checkbox"/> |
| DHCPv4                | <input checked="" type="checkbox"/> |
| IPv4 Address          | 0.0.0.0/0                           |
| IPv4 Gateway          | 0.0.0.0                             |
| VLAN Enable           | <input type="checkbox"/>            |
| VLAN ID               | <input type="text" value="1"/>      |
| Admin User            | admin                               |

Apply Changes

1. To enable IPMI management, select the **Enable** checkbox.
2. Select the **DHCPv4** checkbox.
3. Enable and type the **VLAN ID**, if desired. The VLAN ID should be an 802.1Q VLAN ID.
4. Click **Apply Changes**.

To change the IPMI username and password

| Change IPMI Administrator Details |                                    |
|-----------------------------------|------------------------------------|
| Administrator User Name           | <input type="text" value="admin"/> |
| New Password                      | <input type="text"/>               |
| Confirm Password                  | <input type="text"/>               |

Change Details

1. Enter the **Administrator User Name**.
2. Enter the **New Password** and then type it again in the **Confirm Password** field.

### 3. Click **Change Details**.

Which Exinda models support IPMI?

The table below shows the Exinda models that support IPMI and their capabilities.

| Model | NIC           | User Interface |
|-------|---------------|----------------|
| 306X  | Shared (eth1) | CLI            |
| 406X  | Shared (eth1) | CLI            |
| 606X  | Shared (eth1) | CLI, Web       |
| 806X  | Shared (eth1) | CLI, Web       |
| 1006X | Dedicated     | CLI, Web       |

To manage the power settings on another IPMI-enabled Exinda Appliance

Go to [Configuration > System > Tools > IPMI](#).

#### Related topics

[Manage Power Settings on an IPMI Enabled Appliance](#)

[Manage Power Settings on an IPMI Enabled Appliance](#)

To use an Exinda appliance to manage the power settings of another appliance that has IPMI enabled, navigate to **Configuration > System > Tools > IPMI**.

The screenshot shows a web interface for IPMI configuration. At the top is a section titled "Power Control Options" with a dropdown menu labeled "Command" set to "Get Status". Below this is a section titled "Remote IPMI Login Details" containing three input fields: "IPv4 Address", "Username", and "Password". At the bottom of the form is a button labeled "Do Power Action".

Enter the IPv4 address of the remote appliance and authentication details. Select an operation from the dropdown list and click on the 'Do Power Action' button.

#### EXAMPLE

Power cycle the Exinda appliance with IPMI address 192.168.110.61:

```
(config) # ipmi power address 192.168.110.61 username admin password exinda control cycle
```

```
Chassis Power Control: Cycle
```

## EXAMPLE

Show the current power state of the Exinda appliance with IPMI address 192.168.110.61:

```
(config) # show ipmi power address 192.168.110.61 username admin password exinda
```

## 0.18.10 Overview of QoS by host

Per Host Quality of Service (QoS) allows you to manage traffic congestion by policing the bandwidth available to each host in your network. You can allocate a minimum amount of bandwidth for critical applications, such as VoIP and Citrix, for every host in your network. You can also restrict the bandwidth that each host can utilize for recreational purposes. All out-of-path interfaces are included in the QoS calculations.

The Exinda Appliance enables greater system throughput, up to 10GB, by using multiple queues to handle the traffic. The multiple queues are based on the licensed bandwidth, but the multiple queues are used when the licensed bandwidth exceeds 1.8GB per second.

## NOTE

Per Host QoS can be integrated with Active Directory so bandwidth management can be tailored to users or groups. For more information about integrating the Exinda Appliance with Active Directory, see [Integrate with Active Directory](#).

### Set a per-host limit on bandwidth usage

Per Host QoS is applied at the Virtual Circuit level. It is disabled by default. A Virtual Circuit with Per Host QoS enabled is called a Dynamic Virtual Circuit (DVC).

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > Optimizer**.
6. Click **Create New Virtual Circuit**.
7. In the **Virtual Circuit Name** field, type a name for the virtual circuit.
8. In the **Virtual Circuit Bandwidth** field, type the amount of bandwidth to be used by the virtual circuit.
9. To enable Per Host QoS, select the **Dynamic Virtual Circuit** checkbox. The dialog box expands to show the Dynamic Options.
10. In the **Per Host Bandwidth** field, set the amount of bandwidth (in kB/s or percentage of the virtual circuit bandwidth) that each host will receive. This bandwidth is guaranteed, so it will be available to each host, if required.

## NOTE

If you select **Automatically Share**, the amount of bandwidth each host receives calculated by dividing the Virtual Circuit guaranteed bandwidth by the number of active hosts.

11. In the **Per Host Max Bandwidth** field, set the maximum amount of bandwidth (in kB/s or percentage of the virtual circuit bandwidth) that each host can burst to. If **No Bursting Allowed** is selected, each host gets only the bandwidth that they have been guaranteed.
12. Set the location of the hosts to which to allocate bandwidth.

## NOTE

Internal Hosts are those that are on the LAN side of the appliance. External Hosts are those that are on the WAN side of the appliance.

13. Set the maximum number of hosts that can use the Dynamic Virtual Circuit. Any host that becomes active after the maximum number of hosts is exceeded do not fall into this Virtual Circuit.

## NOTE

If you select **Auto**, the maximum number of hosts is calculated by assuming each host gets its guaranteed bandwidth.

## NOTE

» There is a system limit of 325,00 hosts that can fall into each Dynamic Virtual Circuit. This may occur if the Virtual Circuit has more than 300 Mbps of bandwidth. When this limit is exceeded, hosts fall into the next applicable Virtual Circuit.

» When Per Host QoS is enabled, a further level of traffic shaping is introduced. Traffic is first shaped at the Host level, then at the Policy level. The bandwidth allocated will be the minimum of the two levels.

## Examples

The following examples describe various Dynamic Virtual Circuit configurations.

**Name:** Example 1

**Bandwidth:** 1024 kb/s

**Direction:** Both

**Network Object:** Internal Users

**Dynamic Virtual Circuits Enabled:**

Yes

**Per Host**

**Bandwidth:** Auto

**Per User Max**

**Bandwidth:** 100%

**Host Location:**

Internal

**Max Hosts:** Auto

Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance:

» If there is 1 user, the user gets the full 1024 kb/s.

» If there are 2 users, they each get 512 kb/s and can burst up to the full 1024 kb/s (if the other user is not using their guaranteed 512 kb/s).

» If there are 10 users, they each get 102 kb/s and can burst up to the full 1024 kb/s (if the other users are not using their guaranteed 102 kb/s).

|                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Name:</b> Example 2</p> <p><b>Bandwidth:</b> 1024kbps</p> <p><b>Direction:</b> Both</p> <p><b>Network Object:</b> Internal Users</p> <p><b>Dynamic Virtual Circuits Enabled:</b> Yes</p> <p><b>Per Host Bandwidth:</b> 10%</p> <p><b>Per User Max Bandwidth:</b> No</p> <p><b>Host Location:</b> Internal</p> <p><b>Max Hosts:</b> Auto</p>                                    | <p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <ul style="list-style-type: none"> <li>» If there is 1 user, the user gets 102kbps and cannot burst.</li> <li>» If there are 10 users, they each get 102kbps and cannot burst.</li> <li>» If there are 100 users, the first 10 users each get 102kbps and cannot burst. The remaining 90 users will not match this VC.</li> </ul>                                                                                                                                                                                       |
| <p><b>Name:</b> Example 3</p> <p><b>Bandwidth:</b> 1024kbps</p> <p><b>Direction:</b> Both</p> <p><b>Network Object:</b> Internal Users</p> <p><b>Dynamic Virtual Circuits Enabled:</b> Yes</p> <p><b>Per Host Bandwidth:</b> 64kbps</p> <p><b>Per User Max Bandwidth:</b> 50%</p> <p><b>Host Location:</b> Internal</p> <p><b>Max Hosts:</b> 16</p>                                  | <p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance:</p> <ul style="list-style-type: none"> <li>» If there is 1 user, the user gets 64kbps and can burst up to 512kbps.</li> <li>» If there are 16 users, they each get 64kbps and can burst up to 512kbps (if the other users are not using their guaranteed 64kbps).</li> <li>» If there are 30 users, the first 16 users each get 64kbps and can burst up to 512kbps (if the other users are not using their guaranteed 64kbps). The remaining 14 users will not match this VC.</li> </ul>                                    |
| <p><b>Name:</b> Example 4</p> <p><b>Bandwidth:</b> 1024kbps</p> <p><b>Direction:</b> Both</p> <p><b>Network Object:</b> Internal Users</p> <p><b>Application:</b> Citrix</p> <p><b>Dynamic Virtual Circuits Enabled:</b> Yes</p> <p><b>Per Host Bandwidth:</b> 64kbps</p> <p><b>Per User Max Bandwidth:</b> No</p> <p><b>Host Location:</b> Internal</p> <p><b>Max Hosts:</b> 16</p> | <p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance. "Citrix" is an Application that defines Citrix traffic. This VC will match all Internal User's Citrix traffic:</p> <ul style="list-style-type: none"> <li>» If there is 1 user, the user gets 64kbps for their Citrix traffic and cannot burst.</li> <li>» If there are 16 users, they each get 64kbps for their Citrix traffic and cannot burst.</li> <li>» If there are 30 users, the first 16 users each get 64kbps for their Citrix traffic and cannot burst. The remaining 14 users will not match this VC.</li> </ul> |

## Specify when multi-queue is activated

The Exinda Appliance enables greater system throughput, up to 10GB, by using multiple queues to handle the traffic. Configure the appliance to switch from using a single-queue to using multiple queues when the specified bandwidth is reached.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.

3. Click **Login**.

4. Click **Configuration > System > Tools > Console**.

5. Type the appliance username and password at the prompts. Do one of the following:

- To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`

The `hostname #` prompt appears.

- To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`

The `hostname (config) #` prompt appears.

7. Ensure the bridge and policy-based routing interfaces use auto-license mode. **Auto-license** is the default mode.

```
(config)# bridge <bridge-name> mq mode auto-license
```

```
(config)# pbr interface <interface-name> mq mode auto-license
```

8. Specify at what bandwidth usage the auto-license switches from single-queue to multi-queue.

```
(config)# bridge <bridge-name> mq switch-bandwidth <bandwidth>
```

```
(config)# pbr interface <interface-name> mq switch-bandwidth <bandwidth>
```

### View the number of hosts on a Dynamic Virtual Circuit

The Optimizer Shaping Report shows how each Circuit, Virtual Circuit, Dynamic Virtual Circuit and Policy performs over time. You can see how well your Policies are performing and exactly how much bandwidth each Policy is served. The "Average Rate" is the average policy throughput for the time specified in the time range. The "Current Rate" is the policy throughput averaged over the last 20 seconds.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).

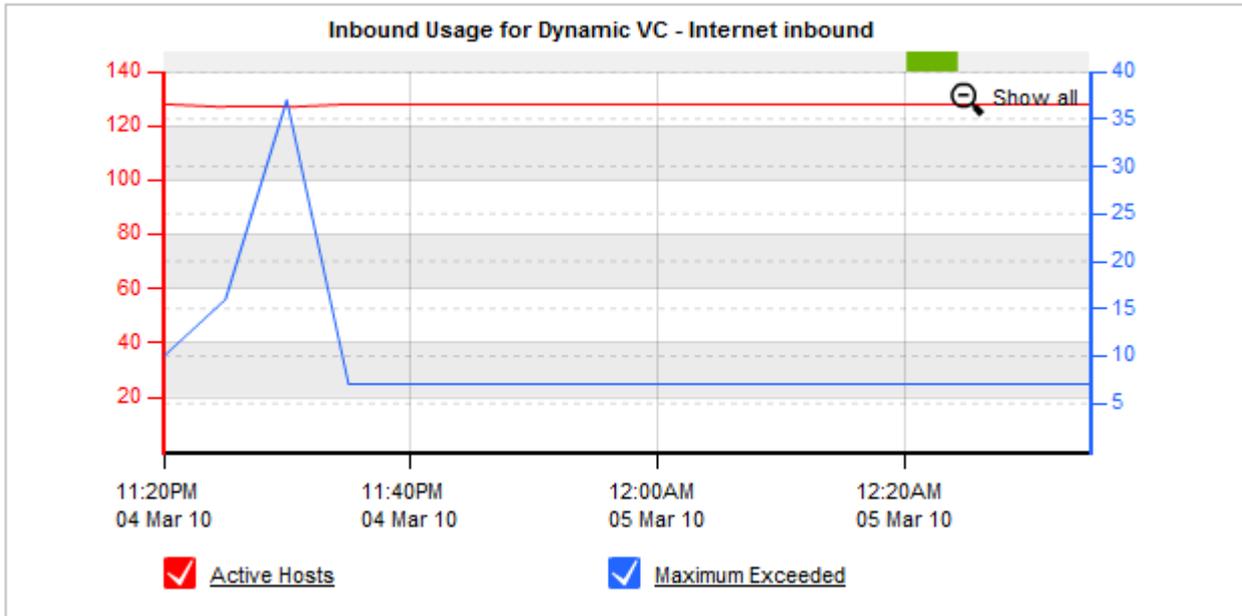
2. Key-in the **User** and **Password**.

3. Click **Login**.

5. Click **Monitor > Control > Policies**.

6. Filter the charts by selecting the relevant **Circuit**, **Virtual Circuit**, and **Policy**. The charts are updated immediately to reflect these choices.

When the Virtual Circuit selected is a Dynamic Virtual Circuit, the following graph is displayed above the throughput graph.



The number of **Active Hosts** for the selected Dynamic Virtual Circuit is represented by the red line.

The number of hosts that have exceeded maximum allowed hosts for this Dynamic Virtual Circuit is represented by the blue **Maximum Exceeded** line.

### Limit Bandwidth Per Host

#### EXAMPLE

Limit bandwidth to 100 kbps for each internal host.

| Edit Virtual Circuit      |                                                                                                                             |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . <input type="text" value="10"/>                                                                                        |
| Virtual Circuit Name      | <input type="text" value="WAN"/>                                                                                            |
| Schedule                  | <input type="text" value="ALWAYS"/>                                                                                         |
| Bandwidth Options         |                                                                                                                             |
| Virtual Circuit Bandwidth | <input type="text" value="50000"/> <input type="text" value="kbps"/>                                                        |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual                                                  |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                                                                         |
| Dynamic Options           |                                                                                                                             |
| Per Host Bandwidth        | <input checked="" type="checkbox"/> Automatically Share<br><input type="text" value="0"/> <input type="text" value="kbps"/> |
| Per Host Max Bandwidth    | <input type="checkbox"/> No Bursting Allowed<br><input type="text" value="100"/> <input type="text" value="kbps"/>          |
| Host Location             | <input type="text" value="Internal"/>                                                                                       |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br><input type="text" value="0"/>                                                  |
| Connection Options        |                                                                                                                             |
| Connection Limit          | <input type="text"/>                                                                                                        |
| Filter Options            |                                                                                                                             |
| VLAN Object               | <input type="text" value="ALL"/>                                                                                            |
| Network Object            | <input type="text" value="ALL"/>                                                                                            |
| Application               | <input type="text" value="ALL"/>                                                                                            |
| Direction                 | <input type="text" value="Both"/>                                                                                           |

In this Dynamic Virtual Circuit, each host is limited to a maximum bandwidth of 100 kbps.

With Max Hosts set to "Auto", a maximum of 5000 hosts can fall into this Dynamic Virtual Circuit. This is calculated by assuming each host is entitled to a minimum bandwidth of 10 kbps as "Automatically Share" is selected.

Optimizer

Optimizer Policies Wizard

|                                                                                                                |                                                                 | Operations                                                  |
|----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|-------------------------------------------------------------|
| Circuit 10 - Default (50000 kbps)                                                                              |                                                                 | <a href="#">Create New Virtual Circuit...</a> --Actions-- ▼ |
| Dynamic Virtual Circuit 10 - WAN (50000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'ALL') |                                                                 | <a href="#">Create New Policy...</a> --Actions-- ▼          |
| <input checked="" type="checkbox"/>                                                                            | 20 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5) | --Actions-- ▼                                               |
| Order:                                                                                                         | Policy: ALL - Accelerate                                        | <input type="button" value="Add To 'WAN'"/>                 |
| <a href="#">Create New Policy...</a>                                                                           |                                                                 |                                                             |
| <a href="#">Create New Virtual Circuit...</a>                                                                  |                                                                 |                                                             |
| <a href="#">Create New Circuit...</a>                                                                          |                                                                 |                                                             |

### EXAMPLE

Limit bandwidth to 100 kbps for each internal host. Further limit P2P traffic to a maximum of 32 kbps across ALL hosts.

In this example, each host will receive between 10 and 100 kbps. In addition, P2P traffic summed across all hosts is capped at 32 kbps, with a guaranteed rate of 16 kbps. To further illustrate this example, suppose there are 100 active users, all using P2P applications on the WAN. The per host bandwidth is 100 kbps, but the P2P policy caps bandwidth at 32 kbps which will be fairly shared between each user. So we would expect to see P2P traffic per user at approx 320 bps.

Optimizer

Optimizer Policies Wizard

|                                                                                                                |                                                                 | Operations                                                  |
|----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|-------------------------------------------------------------|
| Circuit 10 - Default (50000 kbps)                                                                              |                                                                 | <a href="#">Create New Virtual Circuit...</a> --Actions-- ▼ |
| Dynamic Virtual Circuit 10 - WAN (50000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'ALL') |                                                                 | <a href="#">Create New Policy...</a> --Actions-- ▼          |
| <input checked="" type="checkbox"/>                                                                            | 10 P2P (Optimize 16 kbps - 32 kbps, Priority 10)                | --Actions-- ▼                                               |
| <input checked="" type="checkbox"/>                                                                            | 20 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5) | --Actions-- ▼                                               |
| Order:                                                                                                         | Policy: ALL - Accelerate                                        | <input type="button" value="Add To 'WAN'"/>                 |
| <a href="#">Create New Policy...</a>                                                                           |                                                                 |                                                             |
| <a href="#">Create New Virtual Circuit...</a>                                                                  |                                                                 |                                                             |
| <a href="#">Create New Circuit...</a>                                                                          |                                                                 |                                                             |

## Limit Application Bandwidth

### EXAMPLE

Limit P2P to 20 kbps.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 5                                                                     |
| Virtual Circuit Name      | P2P                                                                        |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 5000 kbps ▼                                                                |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input checked="" type="checkbox"/> Automatically Share<br>0 kbps ▼        |
| Per Host Max Bandwidth    | <input type="checkbox"/> No Bursting Allowed<br>20 kbps ▼                  |
| Host Location             | Internal ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | ALL ▼                                                                      |
| Application               | P2P ▼                                                                      |
| Direction                 | Both ▼                                                                     |

In the P2P Dynamic Virtual Circuit, each host is limited to 20 kbps of P2P traffic. With **Max Hosts** set to **Auto**, a maximum of 500 hosts can fall into this Dynamic Virtual Circuit. Additional hosts will share bandwidth allocated in the P2P Overflow Virtual Circuit.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . <input type="text" value="10"/>                                       |
| Virtual Circuit Name      | <input type="text" value="P2P Overflow"/>                                  |
| Schedule                  | <input type="text" value="ALWAYS"/> ▼                                      |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | <input type="text" value="100"/> <input type="text" value="kbps"/> ▼       |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input type="checkbox"/>                                                   |
| Connection Options        |                                                                            |
| Connection Limit          | <input type="text"/>                                                       |
| Filter Options            |                                                                            |
| VLAN Object               | <input type="text" value="ALL"/> ▼                                         |
| Network Object            | <input type="text" value="ALL"/> ▼                                         |
| Application               | <input type="text" value="P2P"/> ▼                                         |
| Direction                 | <input type="text" value="Both"/> ▼                                        |

Screenshot 264: P2P Overflow Virtual Circuit

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 25                                                                    |
| Virtual Circuit Name      | WAN                                                                        |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 45000 kbps ▼                                                               |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input checked="" type="checkbox"/> Automatically Share<br>0 kbps ▼        |
| Per Host Max Bandwidth    | <input type="checkbox"/> No Bursting Allowed<br>100 kbps ▼                 |
| Host Location             | Internal ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | ALL ▼                                                                      |
| Application               | ALL ▼                                                                      |
| Direction                 | Both ▼                                                                     |

Screenshot 265: Dynamic Virtual Circuit To Share Remaining Bandwidth

Create a Dynamic Virtual Circuit using the remaining bandwidth. Each user is limited to a maximum bandwidth of 100 kbps for all other applications.

Optimizer ?

Optimizer Policies Wizard

|                                                                                                                           |                                                                               | Operations                                                  |
|---------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|-------------------------------------------------------------|
| Circuit 10 - Default (50000 kbps)                                                                                         |                                                                               | <a href="#">Create New Virtual Circuit...</a> --Actions-- ▼ |
| Dynamic Virtual Circuit 5 - P2P (5000 kbps [auto kbps - 20 kbps per user / auto users max] 'P2P' traffic to / from 'ALL') |                                                                               | <a href="#">Create New Policy...</a> --Actions-- ▼          |
| <input checked="" type="checkbox"/>                                                                                       | 10 ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)               | --Actions-- ▼                                               |
| Order:                                                                                                                    | Policy: ALL - Accelerate <input type="button" value="Add To 'P2P'"/>          |                                                             |
| <a href="#">Create New Policy...</a>                                                                                      |                                                                               |                                                             |
| Virtual Circuit 10 - P2P Overflow (100 kbps 'P2P' traffic to / from 'ALL')                                                |                                                                               | <a href="#">Create New Policy...</a> --Actions-- ▼          |
| <input checked="" type="checkbox"/>                                                                                       | 10 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)               | --Actions-- ▼                                               |
| Order:                                                                                                                    | Policy: ALL - Accelerate <input type="button" value="Add To 'P2P Overflow'"/> |                                                             |
| <a href="#">Create New Policy...</a>                                                                                      |                                                                               |                                                             |
| Dynamic Virtual Circuit 25 - WAN (45000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'ALL')            |                                                                               | <a href="#">Create New Policy...</a> --Actions-- ▼          |
| <input checked="" type="checkbox"/>                                                                                       | 20 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)               | --Actions-- ▼                                               |
| Order:                                                                                                                    | Policy: ALL - Accelerate <input type="button" value="Add To 'WAN'"/>          |                                                             |
| <a href="#">Create New Policy...</a>                                                                                      |                                                                               |                                                             |
| <a href="#">Create New Virtual Circuit...</a>                                                                             |                                                                               |                                                             |
| <a href="#">Create New Circuit...</a>                                                                                     |                                                                               |                                                             |

## Guarantee Application Bandwidth

### EXAMPLE

Guarantee 30 kbps per host, for the Citrix application.

Citrix typically requires 20 to 30 kbps of bandwidth to work effectively.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 5                                                                     |
| Virtual Circuit Name      | Citrix                                                                     |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 10000 kbps ▼                                                               |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input type="checkbox"/> Automatically Share<br>30 kbps ▼                  |
| Per Host Max Bandwidth    | <input type="checkbox"/> No Bursting Allowed<br>100 kbps ▼                 |
| Host Location             | Internal ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | ALL ▼                                                                      |
| Application               | Citrix ▼                                                                   |
| Direction                 | Both ▼                                                                     |

Screenshot 266: Citrix Dynamic Virtual Circuit

In this example, each user is guaranteed 30 kbps for Citrix. Furthermore, each user can burst up to 100% of the Dynamic Virtual Circuit bandwidth.

With **Max Hosts** set to **Auto**, a maximum of 333 hosts can fall into this Dynamic Virtual Circuit. Additional hosts will share bandwidth allocated in the second Dynamic Virtual Circuit.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 25                                                                    |
| Virtual Circuit Name      | WAN                                                                        |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 45000 kbps ▼                                                               |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input checked="" type="checkbox"/> Automatically Share<br>0 kbps ▼        |
| Per Host Max Bandwidth    | <input checked="" type="checkbox"/> No Bursting Allowed<br>0 kbps ▼        |
| Host Location             | External ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | ALL ▼                                                                      |
| Application               | ALL ▼                                                                      |
| Direction                 | Both ▼                                                                     |

Screenshot 267: Dynamic Virtual Circuit For Remaining Bandwidth

The **WAN** Dynamic Virtual Circuit has **Per Host Bandwidth** set to **Automatically Share**. Each user will be allocated a percentage of the Dynamic Virtual Circuit bandwidth. This is calculated by dividing the Dynamic Virtual Circuit bandwidth by the number of active hosts.

Optimizer

Optimizer Policies Wizard

|                                                                                                                                 |    |                                                              | Operations                                                  |
|---------------------------------------------------------------------------------------------------------------------------------|----|--------------------------------------------------------------|-------------------------------------------------------------|
| Circuit 10 - Default (50000 kbps)                                                                                               |    |                                                              | <a href="#">Create New Virtual Circuit...</a> --Actions-- ▾ |
| Dynamic Virtual Circuit 5 - Citrix (10000 kbps [30 kbps - 100 kbps per user / auto users max] 'Citrix' traffic to / from 'ALL') |    |                                                              | <a href="#">Create New Policy...</a> --Actions-- ▾          |
| <input checked="" type="checkbox"/>                                                                                             | 10 | ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5) | --Actions-- ▾                                               |
| Order:                                                                                                                          |    | Policy: ALL - Accelerate ▾ Add To 'Citrix'                   |                                                             |
| <a href="#">Create New Policy...</a>                                                                                            |    |                                                              |                                                             |
| Dynamic Virtual Circuit 25 - WAN (45000 kbps [auto kbps per user / auto users max] to / from 'ALL')                             |    |                                                              | <a href="#">Create New Policy...</a> --Actions-- ▾          |
| <input checked="" type="checkbox"/>                                                                                             | 20 | ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5) | --Actions-- ▾                                               |
| Order:                                                                                                                          |    | Policy: ALL - Accelerate ▾ Add To 'WAN'                      |                                                             |
| <a href="#">Create New Policy...</a>                                                                                            |    |                                                              |                                                             |
| <a href="#">Create New Virtual Circuit...</a>                                                                                   |    |                                                              |                                                             |
| <a href="#">Create New Circuit...</a>                                                                                           |    |                                                              |                                                             |

## Per Host QoS with Active Directory

### EXAMPLE

Restrict users in the Active Directory 'Students' group to 100 kbps.

1. Using the Web UI - Advanced Mode, go to **Configuration > Objects > Users & Groups**. Edit the "Students (DEV)" group.

exinda. Dashboard | Solution Center | Monitor | Configuration

exinda-0c47ca System Health: **OK** Optimizer: OFF Config: No unsaved changes v7.4.1 (4048)

## Configuration

- Traffic Policies
  - Optimizer
- Objects
  - Network
  - Users & Groups**
  - VLANs
  - Protocols
  - Applications
  - Schedules
  - Adaptive Response
  - Service Levels
  - HTML Response
- System
  - Basic Install Wizard
  - Network
  - Setup
  - Optimization
  - Certificates
  - Authentication
  - Logging
  - Diagnostics
  - Maintenance
  - Tools

## Users & Groups

Network Users | **Network Groups**

Network Groups (Total: 9)

[0-9](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#)

| Group (Domain)                               | Network Object                      | Edit                                |
|----------------------------------------------|-------------------------------------|-------------------------------------|
| Administrators (DEV)                         | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Denied rodc password replication group (DEV) | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Domain admins (DEV)                          | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Domain users (DEV)                           | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Enterprise admins (DEV)                      | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Group policy creator owners (DEV)            | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Schema admins (DEV)                          | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Students (DEV)                               | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |
| Users (DEV)                                  | <input checked="" type="checkbox"/> | <input type="button" value="Edit"/> |

Screenshot 268: Active Directory Groups

- Each host is limited to a maximum bandwidth of 100 kbps. With **Per Host Bandwidth** set to **Automatically Share**, a maximum of 400 hosts can fall into this Dynamic Virtual Circuit.
- Create a Network Object from the Active Directory group.

exinda. Dashboard | Solution Center | Monitor | Configuration

exlnda-0c47ca System Health: OK Optimizer: OFF Config: No unsaved changes v7.4.1 (4048)

### Configuration

- Traffic Policies
- Optimizer
- Objects
  - Network
  - Users & Groups
  - VLANs
  - Protocols
  - Applications
  - Schedules
  - Adaptive Response
  - Service Levels
  - HTML Response

### Edit Group

Network Users | Network Groups

Logins from users in the network group(s) DEV\Students will be mapped to the Students network object

Name: DEV\Students

Map to Network Object:

Ignore Domain:

Apply Cancel

Screenshot 269: Map AD Group 'Students' To Network Object 'Students'

4. The Network Object "Students" can now be used in a Dynamic Virtual Circuit.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 5                                                                     |
| Virtual Circuit Name      | Students                                                                   |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 24000 kbps ▼                                                               |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input checked="" type="checkbox"/> Automatically Share<br>0 kbps ▼        |
| Per Host Max Bandwidth    | <input type="checkbox"/> No Bursting Allowed<br>100 kbps ▼                 |
| Host Location             | Internal ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | Students ▼                                                                 |
| Application               | ALL ▼                                                                      |
| Direction                 | Both ▼                                                                     |

Screenshot 270: Students Dynamic Virtual Circuit

5. Each host in the "Students" Network Object is limited to 100 kbps. With Max Hosts set to "Auto", a maximum of 2400 hosts can fall into this Dynamic Virtual Circuit. Additional hosts will share bandwidth allocated in the "Students Overflow" Virtual Circuit.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . <input type="text" value="10"/>                                       |
| Virtual Circuit Name      | <input type="text" value="Students Overflow"/>                             |
| Schedule                  | <input type="text" value="ALWAYS"/> ▼                                      |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | <input type="text" value="1000"/> <input type="text" value="kbps"/> ▼      |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input type="checkbox"/>                                                   |
| Connection Options        |                                                                            |
| Connection Limit          | <input type="text"/>                                                       |
| Filter Options            |                                                                            |
| VLAN Object               | <input type="text" value="ALL"/> ▼                                         |
| Network Object            | <input type="text" value="ALL"/> ▼                                         |
| Application               | <input type="text" value="ALL"/> ▼                                         |
| Direction                 | <input type="text" value="Both"/> ▼                                        |

Screenshot 271: Students Overflow Virtual Circuit

| Add New Virtual Circuit   |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 15                                                                    |
| Virtual Circuit Name      | WAN                                                                        |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 25000 kbps ▼                                                               |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input type="checkbox"/> Automatically Share<br>100 kbps ▼                 |
| Per Host Max Bandwidth    | <input checked="" type="checkbox"/> No Bursting Allowed<br>0 kbps ▼        |
| Host Location             | Internal ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | ALL ▼                                                                      |
| Application               | ALL ▼                                                                      |
| Direction                 | Both ▼                                                                     |

Screenshot 272: Dynamic Virtual Circuit For Remaining Bandwidth

6. Another Dynamic Virtual Circuit can be created to share the remaining bandwidth for other hosts. In this example, each host is guaranteed 100 kbps with **No Bursting Allowed**.

The screenshot shows the 'Optimizer' interface with three tabs: 'Optimizer', 'Policies', and 'Wizard'. The 'Policies' tab is active, displaying a list of network policies. Each policy row includes a checkbox, a priority value (10), a policy name, a description, an 'Order' field, a 'Policy' dropdown menu (set to 'ALL - Accelerate'), an 'Add To' button, and an 'Operations' dropdown menu (set to '--Actions--').

| Policy Name                                                                                                             | Description                                                     | Order | Policy           | Add To                     | Operations  |
|-------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|-------|------------------|----------------------------|-------------|
| Circuit 10 - Default (50000 kbps)                                                                                       |                                                                 |       |                  |                            | --Actions-- |
| Dynamic Virtual Circuit 5 - Students (24000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'Students') |                                                                 |       |                  |                            | --Actions-- |
| <input checked="" type="checkbox"/>                                                                                     | 10 ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7) |       | ALL - Accelerate | Add To 'Students'          | --Actions-- |
| Virtual Circuit 10 - Students Overflow (1000 kbps to / from 'ALL')                                                      |                                                                 |       |                  |                            | --Actions-- |
| <input checked="" type="checkbox"/>                                                                                     | 10 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5) |       | ALL - Accelerate | Add To 'Students Overflow' | --Actions-- |
| Dynamic Virtual Circuit 15 - WAN (25000 kbps [100 kbps per user / auto users max] to / from 'ALL')                      |                                                                 |       |                  |                            | --Actions-- |
| <input checked="" type="checkbox"/>                                                                                     | 10 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5) |       | ALL - Accelerate | Add To 'WAN'               | --Actions-- |

## Per Host QoS for Adaptive Response

### EXAMPLE

Restrict users in the Active Directory 'Students' group to 100 kbps, once the user has downloaded 100 MB per day.

1. Create the Network Object **Students** based on the Active Directory **Students** group as shown in the previous topic.
2. Using the Web UI - Advanced Mode, go to **Configuration > Objects > Adaptive Response**.
3. Create a new Adaptive Response rule based on the **Students** Network Object. Each host is allowed to download 100 MB per day before being placed into the **Students\_Shaped** Network Object.

exinda. Dashboard | Solution Center | Monitor | Configuration

exinda-0c47ca System Health: OK Optimizer: Restart Required Config: No unsaved changes v7.4.1 (4048)

## Configuration

- Traffic Policies
  - Optimizer
- Objects
  - Network
  - Users & Groups
  - VLANs
  - Protocols
  - Applications
  - Schedules
  - Adaptive Response**
  - Service Levels
  - HTML Response
- System
  - Basic Install Wizard
  - Network
  - Setup
  - Optimization
  - Certificates
  - Authentication
  - Logging
  - Diagnostics
  - Maintenance
  - Tools

## Adaptive Response

Adaptive Response Limits are rules which are used to create and populate network objects based on amount of data transferred. These [dynamic network](#)

### Add New AR Limit

Name:

Source Network Object:

Destination Network Object:

Duration:

Direction:

Limit Type:

Amount (MB):

Time (Min):

Enable:

| Name          | Source Network | Destination Network | Duration | Direction | Amount (MB) | Time (Min) | Enabled | Edit | Delete |
|---------------|----------------|---------------------|----------|-----------|-------------|------------|---------|------|--------|
| No AR Limits. |                |                     |          |           |             |            |         |      |        |

Screenshot 273: Create Adaptive Response Object

4. Create a Dynamic Virtual Circuit, with **Network Object** set to "**Students\_Shaped**". Hosts matching this Network Object will fall into this Dynamic Virtual Circuit.
5. Each host is limited to a maximum bandwidth of 100 kbps. With **Per Host Bandwidth** set to **Automatically Share**, a maximum of 400 hosts can fall into this Dynamic Virtual Circuit.

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 5                                                                     |
| Virtual Circuit Name      | Students                                                                   |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 4000 kbps ▼                                                                |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input checked="" type="checkbox"/>                                        |
| Dynamic Options           |                                                                            |
| Per Host Bandwidth        | <input checked="" type="checkbox"/> Automatically Share<br>0 kbps ▼        |
| Per Host Max Bandwidth    | <input type="checkbox"/> No Bursting Allowed<br>100 kbps ▼                 |
| Host Location             | Internal ▼                                                                 |
| Max Hosts                 | <input checked="" type="checkbox"/> Auto<br>0                              |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | Students_Shaped ▼                                                          |
| Application               | ALL ▼                                                                      |
| Direction                 | Both ▼                                                                     |

Screenshot 274: Students Dynamic Virtual Circuit

6. Additional hosts will share bandwidth allocated in the Students Overflow Virtual Circuit.

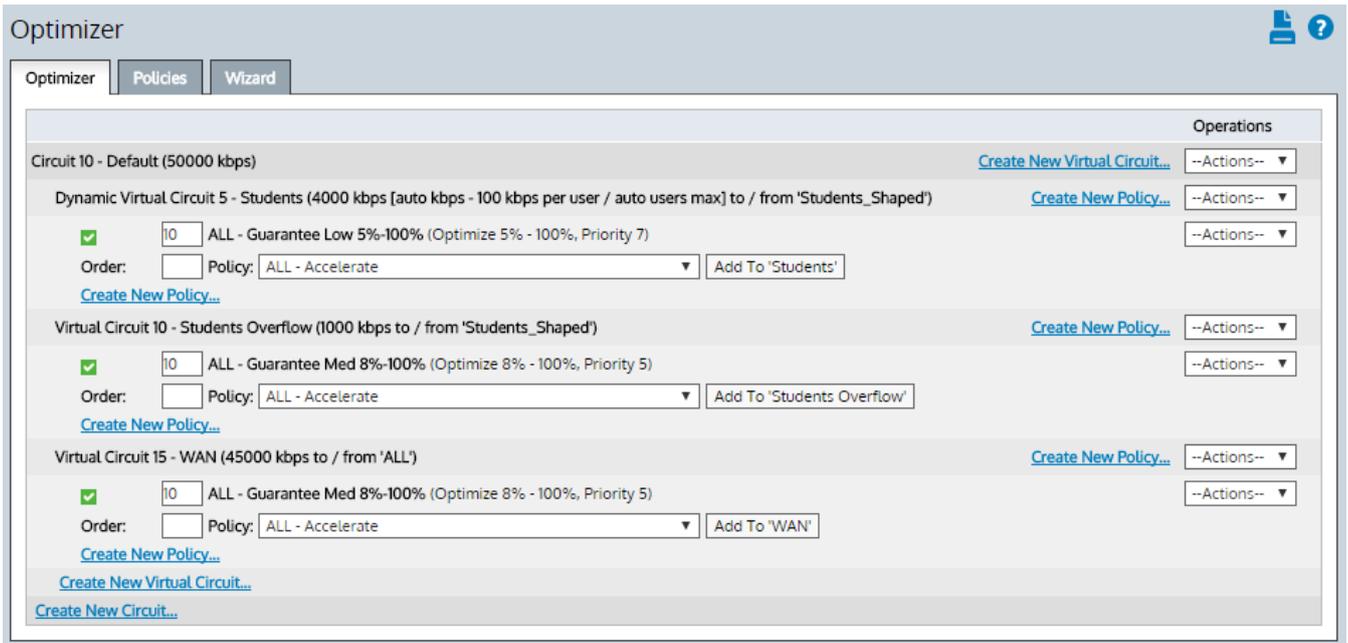
| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . <input type="text" value="10"/>                                       |
| Virtual Circuit Name      | <input type="text" value="Students Overflow"/>                             |
| Schedule                  | <input type="text" value="ALWAYS"/>                                        |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | <input type="text" value="1000"/> <input type="text" value="kbps"/>        |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input type="checkbox"/>                                                   |
| Connection Options        |                                                                            |
| Connection Limit          | <input type="text"/>                                                       |
| Filter Options            |                                                                            |
| VLAN Object               | <input type="text" value="ALL"/>                                           |
| Network Object            | <input type="text" value="Students_Shaped"/>                               |
| Application               | <input type="text" value="ALL"/>                                           |
| Direction                 | <input type="text" value="Both"/>                                          |

Screenshot 275: Students Overflow Virtual Circuit

| Edit Virtual Circuit      |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| Virtual Circuit Number    | 10 . 15                                                                    |
| Virtual Circuit Name      | WAN                                                                        |
| Schedule                  | ALWAYS ▼                                                                   |
| Bandwidth Options         |                                                                            |
| Virtual Circuit Bandwidth | 45000 kbps ▼                                                               |
| Oversubscription          | <input checked="" type="radio"/> Automatic<br><input type="radio"/> Manual |
| Dynamic Virtual Circuit   | <input type="checkbox"/>                                                   |
| Connection Options        |                                                                            |
| Connection Limit          |                                                                            |
| Filter Options            |                                                                            |
| VLAN Object               | ALL ▼                                                                      |
| Network Object            | ALL ▼                                                                      |
| Application               | ALL ▼                                                                      |
| Direction                 | Both ▼                                                                     |

Screenshot 276: Virtual Circuit To Share Remaining Bandwidth

7. Other users and students who have not used their 100MB daily quota will share 45 Mbps of bandwidth in the WAN Virtual Circuit.



## 0.19 System Setup

Learn how to set up your Exinda Appliance(s). The configuration information provided focuses on the appliance and not upon the Exinda firmware.

|                                          |     |
|------------------------------------------|-----|
| 0.19.1 Date and Time Configuration ..... | 600 |
| 0.19.2 UI Access Configuration .....     | 603 |
| 0.19.3 SDP Configuration .....           | 605 |
| 0.19.4 Configure SQL Access .....        | 605 |
| 0.19.5 Monitoring Configuration .....    | 621 |
| 0.19.6 Netflow Configuration .....       | 625 |
| 0.19.7 Create a Scheduled Job .....      | 628 |
| 0.19.8 Alerts .....                      | 630 |
| 0.19.9 Control Configuration .....       | 633 |
| 0.19.10 Disk Storage Explained .....     | 633 |
| 0.19.11 Virtualization .....             | 640 |

### 0.19.1 Date and Time Configuration

It is important to accurately set the date and time of your Exinda appliance so that all time-based functions use the appropriate time. It is highly recommended to set the date and time using an NTP server. This is particularly important if you have multiple Exinda appliances where you need to correlate or aggregate the monitoring data or if you need the exported NetFlow records to be synchronized with NetFlow records from other network appliances.

It is recommended to set the date and time using one or more NTP servers. The more NTP servers that are configured, the more accurate the time will be. It is generally accepted that four NTP servers is the optimal number of servers for

extremely accurate time. To explain why four NTP servers is considered an optimal number of servers, see <http://www.ntp.org/ntpfaq/NTP-s-algo-real.htm>

A great source for NTP servers is the NTP Pool Project at <http://www.pool.ntp.org/en/use.html>

The date and time setting has implications for the following functions:

- » Monitoring data have time stamps and the monitor charts will be reported relative to these time stamps
- » Exported NetFlow records have time stamps
- » Schedule-based policies need to take effect at the appropriate times
- » Scheduled events, such as scheduled reports or scheduled jobs, need to occur at appropriate times.

Note that when the current time on the appliance is out of sync with the date-time provided by the NTP servers, the NTP servers will slowly adjust the clock time. If the appliance's time is significantly out of sync with the NTP servers (say, 1000 seconds or approximately 15 minutes), then it is recommended that you force the appliance to jump to the correct time using the `ntpdate` command from the command line.

Where do I find this configuration?

Go to **Configuration > System > Setup > Date and Time**.

To set the date & time using an NTP server

| NTP Servers              |                |         |                                     |
|--------------------------|----------------|---------|-------------------------------------|
| Server                   | Version        | Enabled |                                     |
| <input type="checkbox"/> | 0.pool.ntp.org | 4       | <input checked="" type="checkbox"/> |
| <input type="checkbox"/> | 1.pool.ntp.org | 4       | <input checked="" type="checkbox"/> |
| <input type="checkbox"/> | 2.pool.ntp.org | 4       | <input checked="" type="checkbox"/> |
| <input type="checkbox"/> | 3.pool.ntp.org | 4       | <input checked="" type="checkbox"/> |

### Add New NTP Server

Server address

Version

Enabled

1. Add one or more NTP servers by entering the IP address or hostname of the NTP server, the version of NTP supported by the server, and enabling it by checking the enable checkbox in the Add New NTP Server area, Only hostnames and IPv4 addresses are supported.

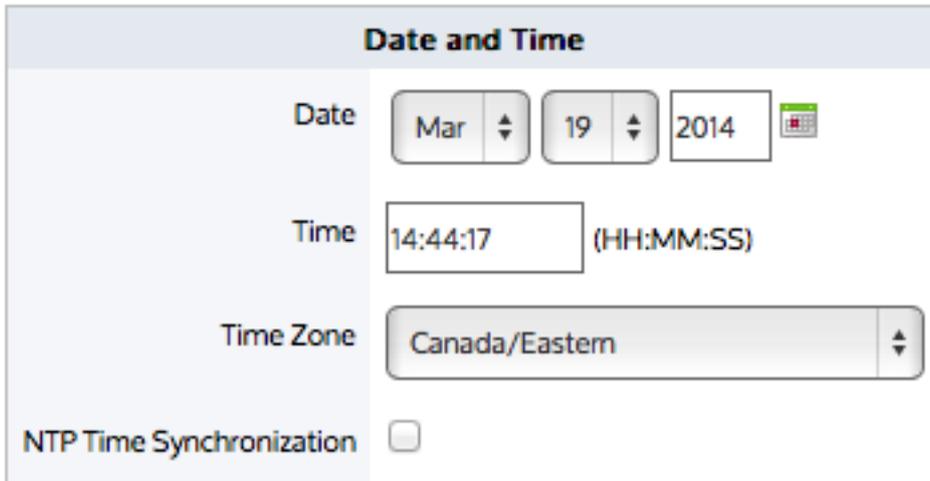
2. In the Date and Time area, select the **NTP Time Synchronization** checkbox and **Apply Changes**.

**NTP Time Synchronization**

The change is applied only if you accept the Restart Message to restart the UI.

Any of the NTP servers can be disabled, re-enabled, or removed by clicking the appropriate button - **Disable Server**, **Enable Server**, **Remove Server**.

To set the date & time manually.



1. In the Date and Time area, ensure that the **NTP Time Synchronization** checkbox is not checked.
2. Set the desired date, time, and timezone and click **Apply Changes**.

#### WARNING

If you change the time manually, you will be prompted to restart the UI. If you do not accept the Restart, the configuration change is not applied.

If the NTP Time Synchronization checkbox is checked, then the manual date-time setting will not be applied.

To force a time reset when the time is significantly out of sync

On the command line, type: `ntpdate <ntp-server-address>`

- `<ntp-server-address>` - The location of an NTP server specified as hostname or IPv4/IPv6 address.

This command is similar to the deprecated `ntpdate` command.

## 0.19.2 UI Access Configuration

The Access page allows you to specify how long the appliance web user interface can be idle before the user is automatically logged out. Similarly you can specify how long the CLI can be idle before being logged out. You can specify whether to enable http or https access and which port numbers to use. If desired, you can also choose to disable the web UI altogether. You can specify whether CLI access is through telnet or SSH.

#### NOTE

Once you disable the Web UI, you can only re-enable it via the CLI.

| Web UI Options      |                                            |
|---------------------|--------------------------------------------|
| Web UI              | <input checked="" type="checkbox"/> Enable |
| Auto Logout Timeout | <input type="text" value="0"/> minutes     |
| HTTP Access         | <input type="checkbox"/> Enable            |
| HTTP Port           | <input type="text" value="80"/>            |
| HTTPS Access        | <input checked="" type="checkbox"/> Enable |
| HTTPS Port          | <input type="text" value="443"/>           |
| Web Session Renewal | <input type="text" value="60"/> minutes    |
| Web Session Timeout | <input type="text" value="1440"/> minutes  |

[Apply Changes](#)

Screenshot 277: Web UI options for setting HTTP or HTTPS access, auto-logout time period, and disabling the Web UI

| CLI Options         |                                            |
|---------------------|--------------------------------------------|
| Auto Logout Timeout | <input type="text" value="900"/> seconds   |
| Telnet Access       | <input type="checkbox"/> Enable            |
| SSH Access          | <input checked="" type="checkbox"/> Enable |
| SSH Version         | <input type="text" value="SSH v2 or v1"/>  |

[Apply Changes](#)

Screenshot 278: CLI options for setting Telnet or SSH access, and the auto-logout timeout period

To configure the Web UI to auto-logout after a specified idle period

1. Ensure the **Web UI Enable** checkbox is checked.
2. Set the **Auto Logout Timeout** period to the specified number of minutes that the user can be idle before the user gets automatically logged out. To configure the system to never automatically log out, set the field to **0** minutes. It is not recommended to change the values in the **Web Session Renewal** field or the **Web Session Timeout** field.
3. Click **Apply Changes**.

To enable HTTP or HTTPS web access

1. Ensure the **Web UI Enable** checkbox is checked.
2. To enable HTTP access, check the **HTTP Access** checkbox and specify the **HTTP Port** number to use. The default port number is 80.
3. To enable HTTPS access, check the **HTTPS Access** checkbox and specify the **HTTPS Port** number to use. The default port number is 443.
4. Click **Apply Changes**.

To disable the Web UI

1. Uncheck the **Web UI Enable** checkbox.
2. Click **Apply Changes**.

To re-enable the Web UI

From the CLI type: `web enable`

To configure CLI to be accessed via Telnet or SSH

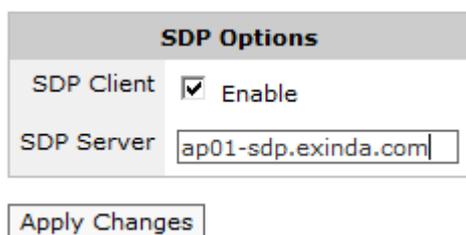
1. To enable Telnet access, check the **Telnet Access** checkbox.
2. To enable SSH access, check the **SSH Access** checkbox and select the **SSH Version** to use.
3. Click **Apply Changes**.

To configure CLI to auto-logout after a specified idle period

1. Set the **Auto Logout Timeout** period to the specified number of seconds that the user can be idle before the user gets automatically logged out. To configure the system to never automatically log out, set the field to **0** minutes.
2. Click **Apply Changes**.

### 0.19.3 SDP Configuration

The SDP options form is used to enable the SDP service on the Exinda appliance and change the SDP server address if required. This will enable communication between the Exinda appliance and the specified SDP server. A SDP subscription is required in order to use the Exinda SDP feature.



| SDP Options |                                                  |
|-------------|--------------------------------------------------|
| SDP Client  | <input checked="" type="checkbox"/> Enable       |
| SDP Server  | <input type="text" value="ap01-sdp.exinda.com"/> |

For further information, consult the SDP User Manual.

### 0.19.4 Configure SQL Access

The SQL Access feature on an Exinda appliance provides access to the traffic monitoring database from any ODBC compliant application.

In order to use this feature, SQL access needs to be configured on the Exinda appliance, and an ODBC driver needs to be installed and configured on a client. ODBC aware applications running on the client will then be able to query the Exinda appliance's internal monitoring database.

This How to Guide explains how to configure the Exinda appliance to accept remote SQL connections, as well as setting up the ODBC driver on Windows 8 and Windows 10 clients.

#### Download the ODBC Driver

Download the ODBC driver version that corresponds to your client operating system. Follow the instructions on this site for installing the ODBC driver on your client operating system.

The ODBC driver can be downloaded from:

<http://dev.mysql.com/downloads/connector/odbc/>

#### Set Remote SQL Options

To allow the Exinda appliance to accept remote SQL connections from an external ODBC connector, you must configure the settings in **Configuration > System > Setup > SQL Access**.

- » **Remote SQL:** Select this option to allow the Exinda appliance to accept remote SQL connections from external ODBC connectors.
- » **Allow access from (Hostname or IP):** Use this option to restrict the hosts that can connect to the SQL database. Specify '%' to allow any hosts to connect or type an IP address or Hostname of a specify host to restrict access.
- » **Username:** Specify a username to use for authentication (E.g. 'database').
- » **Password:** Specify a password to use for authentication.
- » **Confirm Password:** Retype the password specified above.

Apply the changes. The SQL access will be made available immediately. A successfully configured appliance would look something like:

| Remote SQL Options                 |                                                                  |
|------------------------------------|------------------------------------------------------------------|
| Remote SQL                         | <input checked="" type="checkbox"/> Enable                       |
| Allow access from (Hostname or IP) | <input style="width: 150px;" type="text" value="%"/> (% = 'any') |
| Username                           | <input style="width: 100px;" type="text" value="database"/>      |
| Password                           | <input style="width: 100px;" type="password"/>                   |
| Confirm Password                   | <input style="width: 100px;" type="password"/>                   |

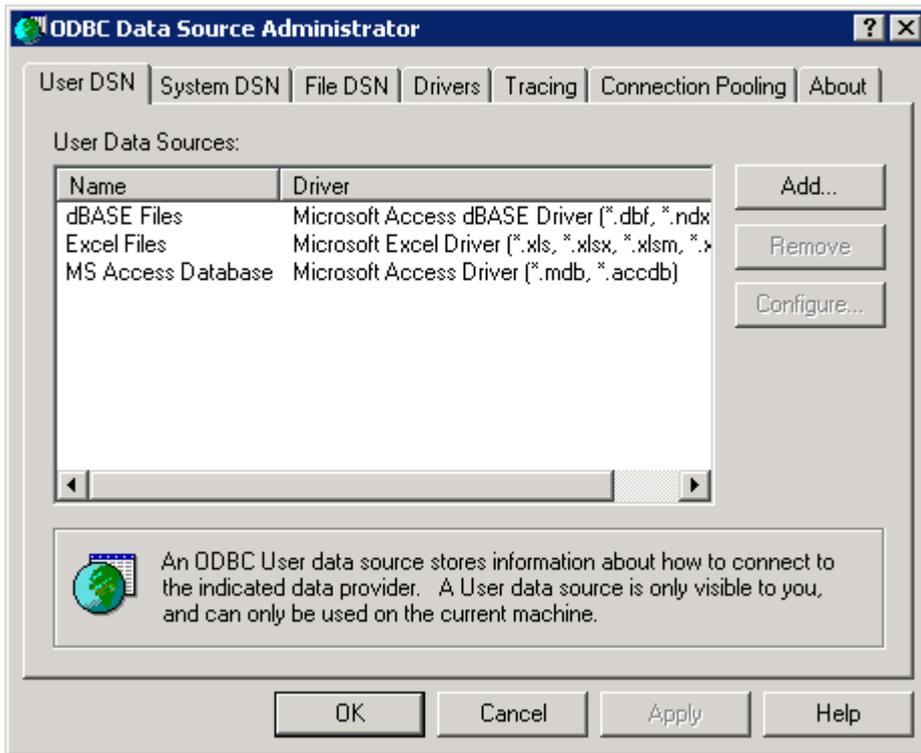
Once remote SQL access has been configured on the Exinda appliance, the next step is to create an ODBC data source on the client.

See the following for more information on creating ODBC sources:

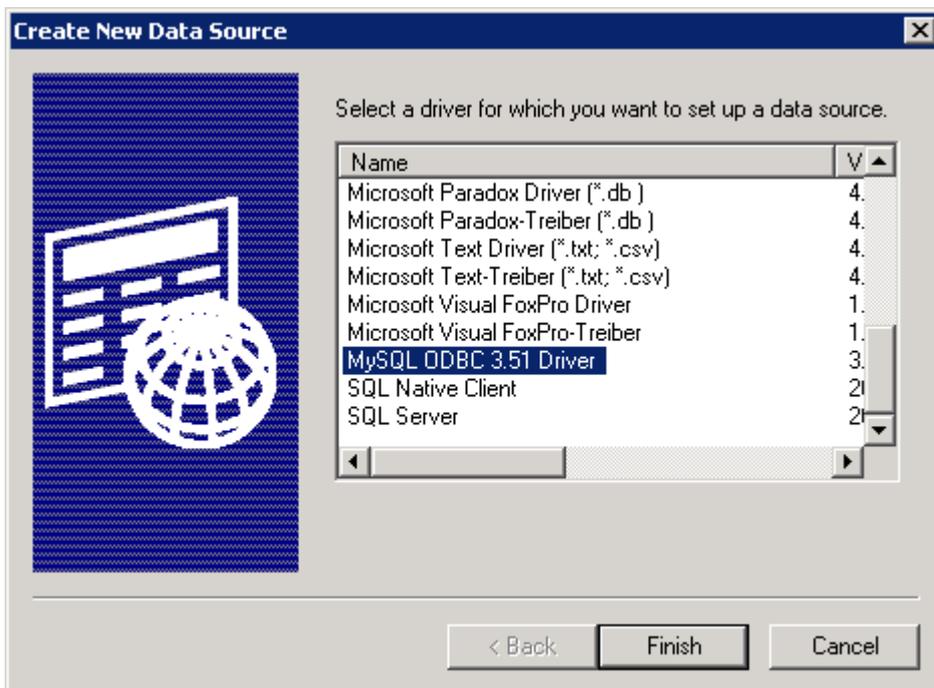
- » [Create ODBC Data Source on Windows XP](#)
- » [Create ODBC Data Source on Windows 7](#)

#### Create ODBC Data Source on Windows XP

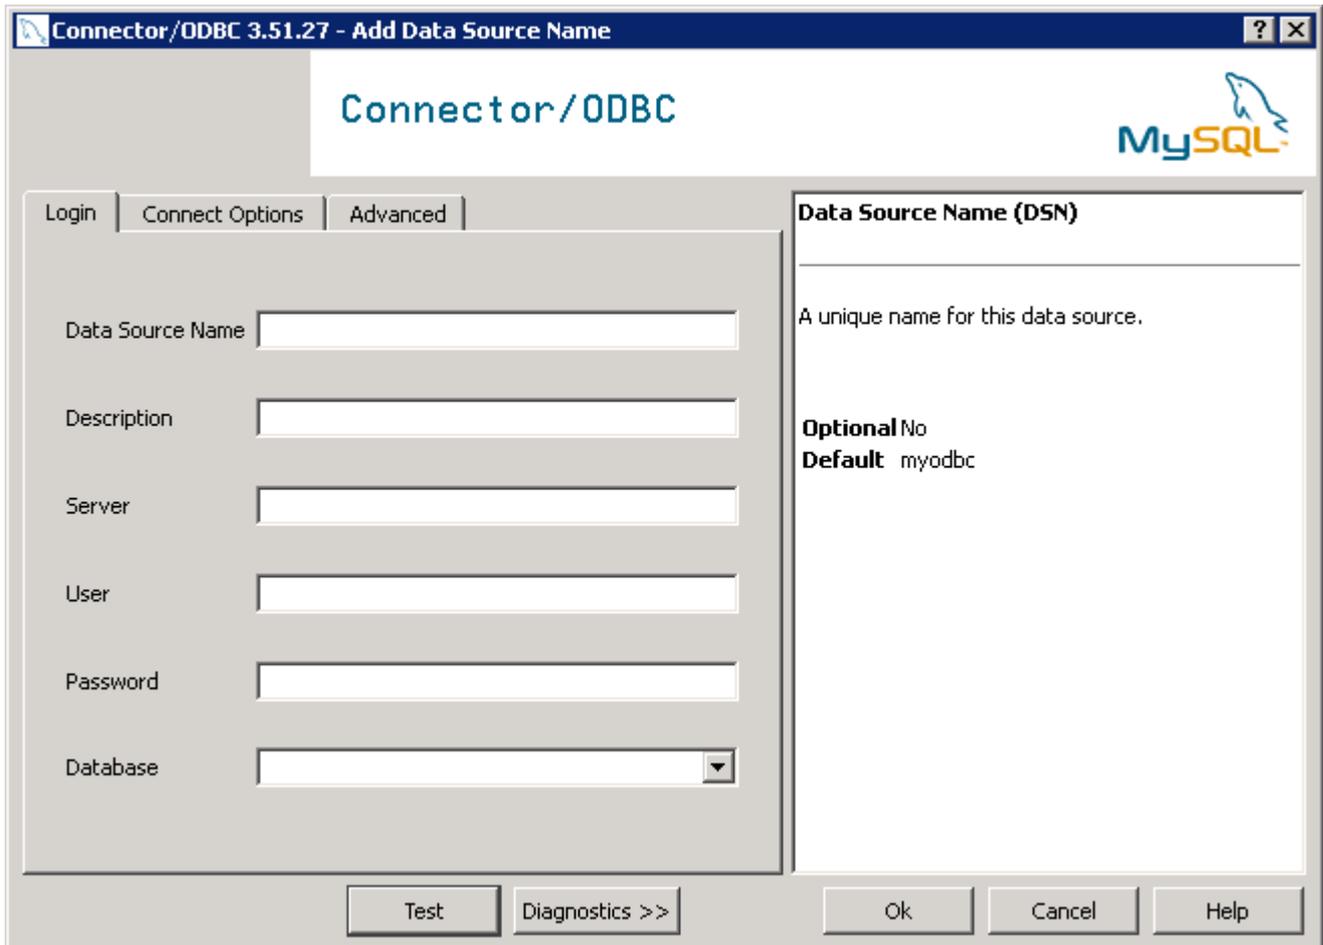
1. Open **Administrative Tools** and select **Data Sources (ODBC)**. You should be presented with the following dialog.



2. Select the **User DSN** tab or the **System DSN** tab depending on whether you wish the SQL data to be made available to only the current user (User DSN) or all users (System DSN). Then click **Add...** This will start a wizard that allows you to create a new data source.

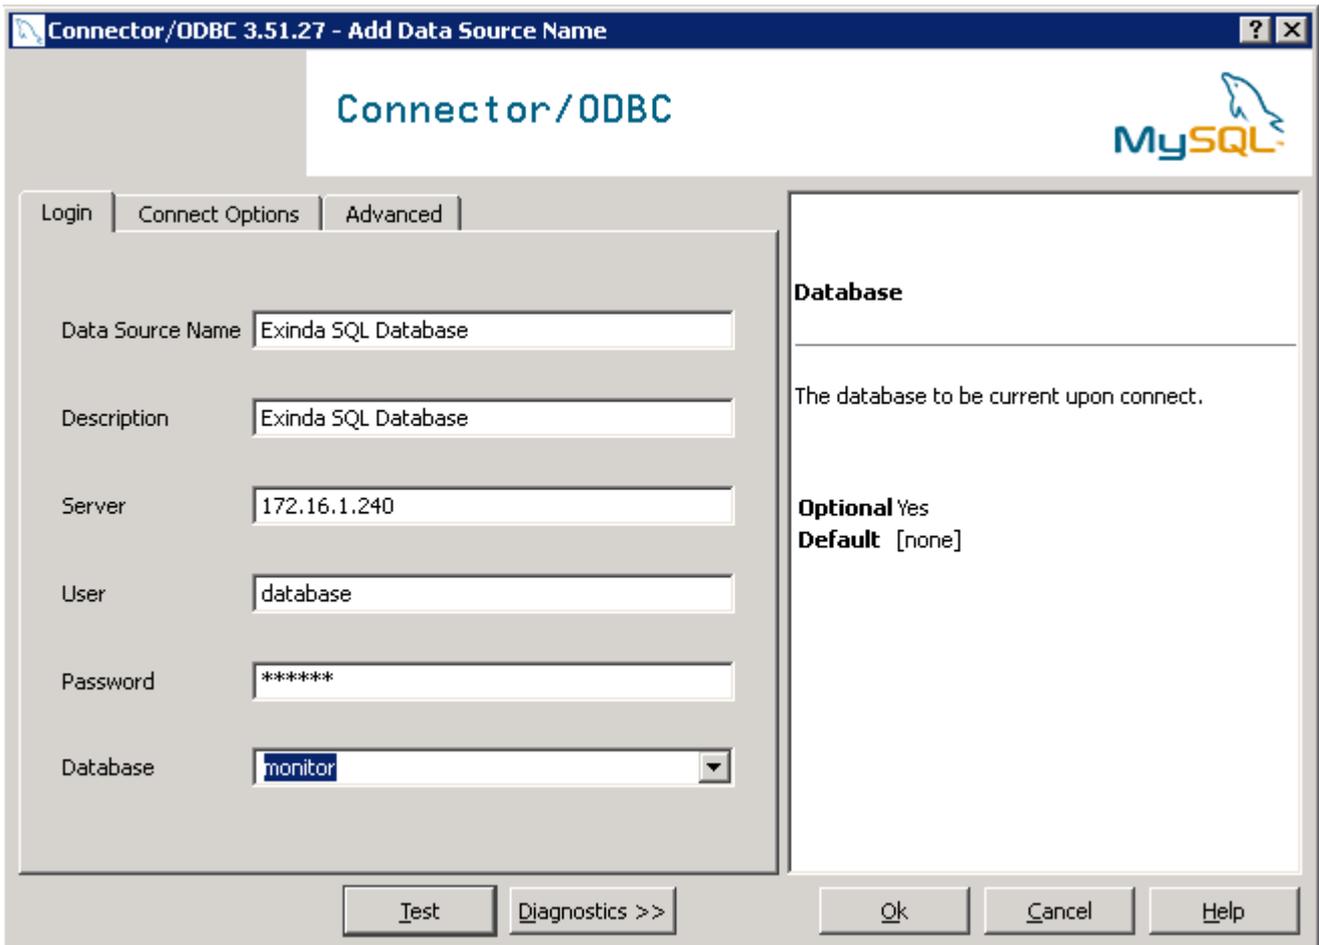


3. Select **MySQL ODBC Driver** and click **Finish**. You will be prompted to enter details about the SQL access using the form below:

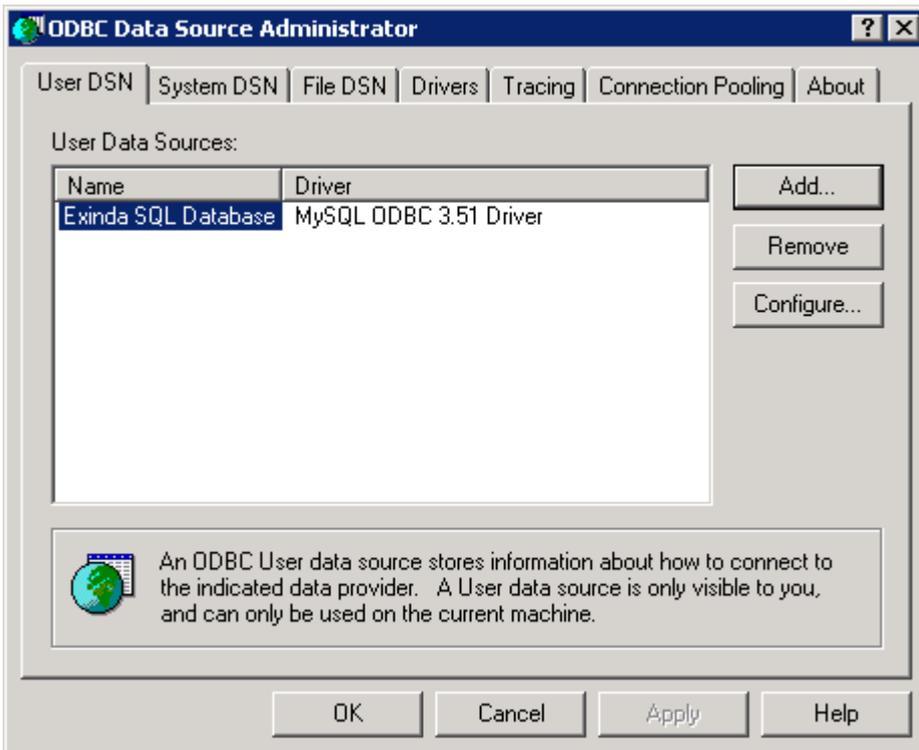


|                                |                                                                                                                                                                                                          |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Source Name / Description | Enter a descriptive name for the DSN. E.g. 'Exinda SQL Database'.                                                                                                                                        |
| Server                         | Enter the IP address of the Exinda appliance.                                                                                                                                                            |
| User                           | Enter the username you specified when enabling SQL access on the Exinda appliance.                                                                                                                       |
| Password                       | Enter the password you specified when enabling SQL access on the Exinda appliance.                                                                                                                       |
| Database                       | Once the above fields are configured, press the 'Test' button. If the connection attempt is successful, the 'Database' drop down will be populated with a list of available databases. Select 'monitor'. |

Here is what a successful configuration looks like:

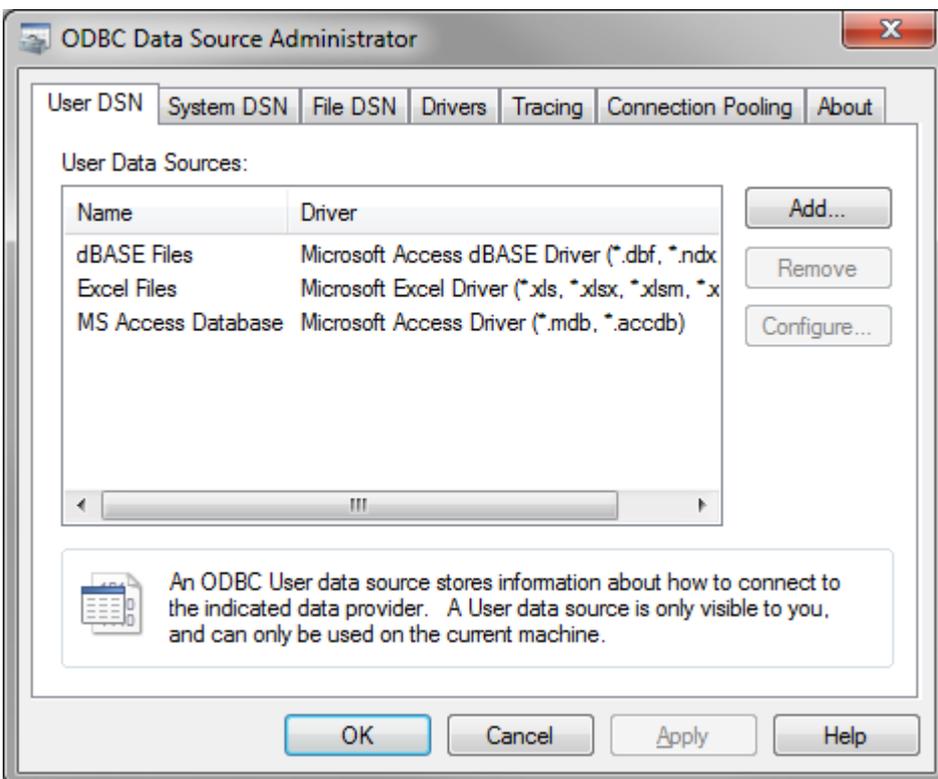


Click **OK**. This will add the 'Exinda SQL Database' to the list of available data sources that can be used by 3rd party applications on this client.

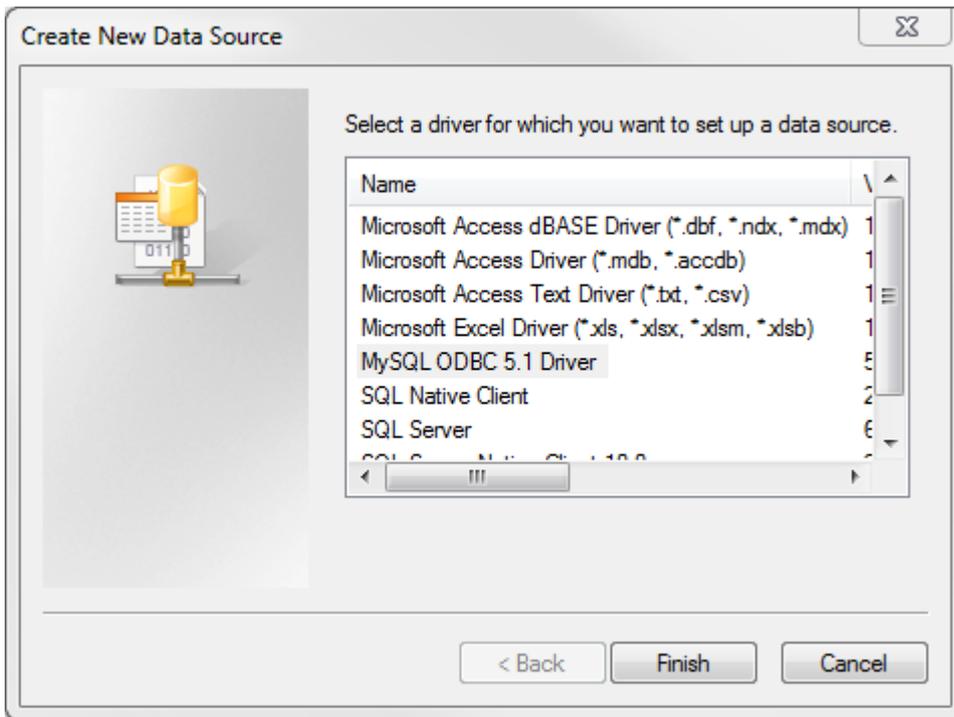


### Create ODBC Data Source on Windows 7

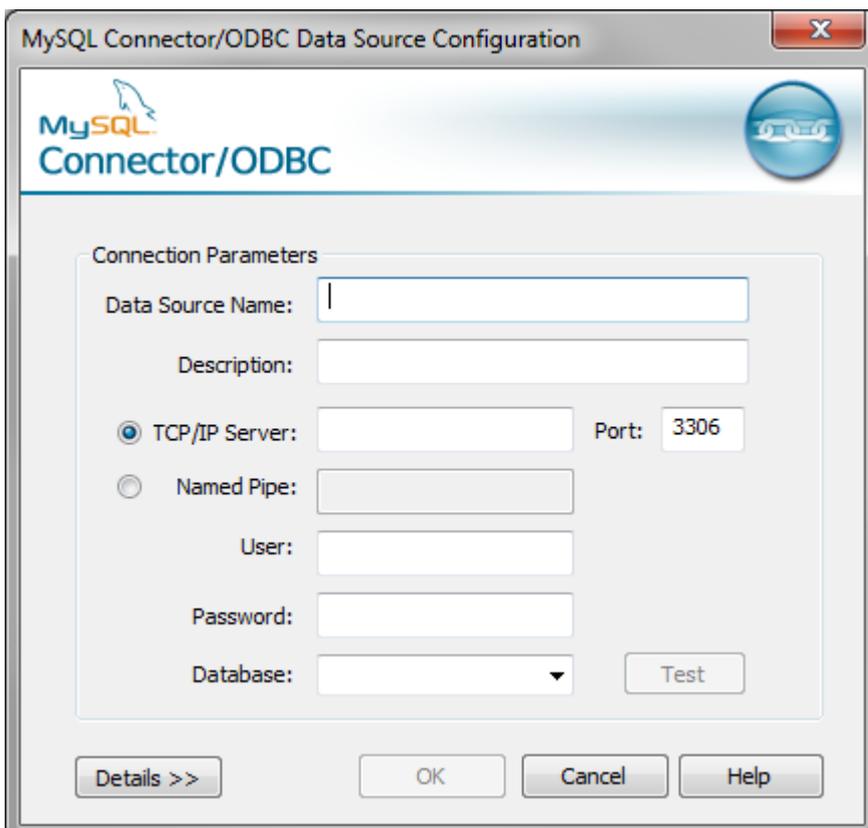
Open **Administrative Tools** and select **Data Sources (ODBC)**. You should be presented with the following dialog.



Select the **User DSN** tab or the **System DSN** tab depending on whether you wish the SQL data to be made available to only the current user (User DSN) or all users (System DSN). Then click **Add...** This will start a wizard that allows you to create a new data source.



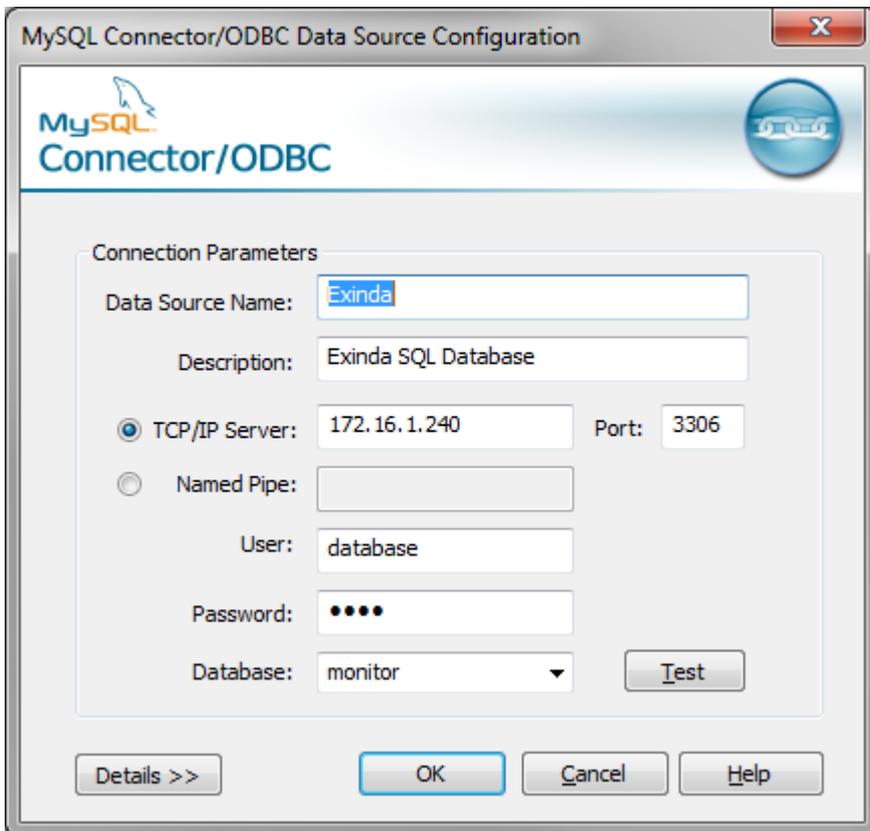
Select **MySQL ODBC Driver** and click **Finish**. You will be prompted to enter details about the SQL access using the form below:



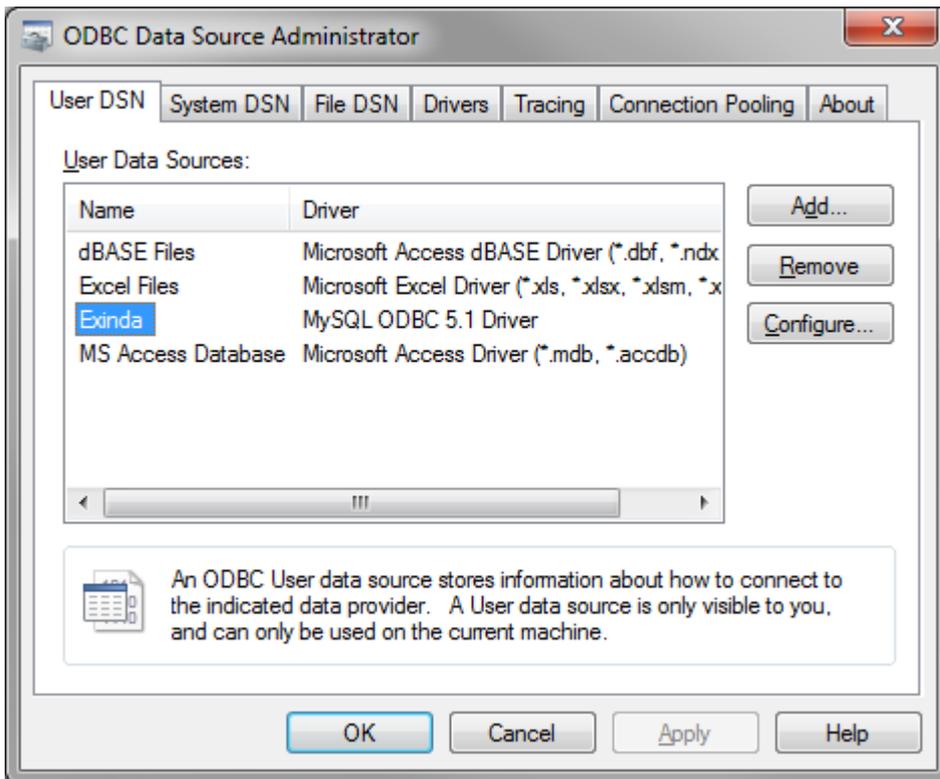
Data Source Name / Description      Enter a descriptive name for the DSN. E.g. 'Exinda SQL Database'.

|          |                                                                                                                                                                                                          |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Server   | Enter the IP address of the Exinda appliance.                                                                                                                                                            |
| User     | Enter the username you specified when enabling SQL access on the Exinda appliance.                                                                                                                       |
| Password | Enter the password you specified when enabling SQL access on the Exinda appliance.                                                                                                                       |
| Database | Once the above fields are configured, press the 'Test' button. If the connection attempt is successful, the 'Database' drop down will be populated with a list of available databases. Select 'monitor'. |

Here is what a successful configuration looks like:



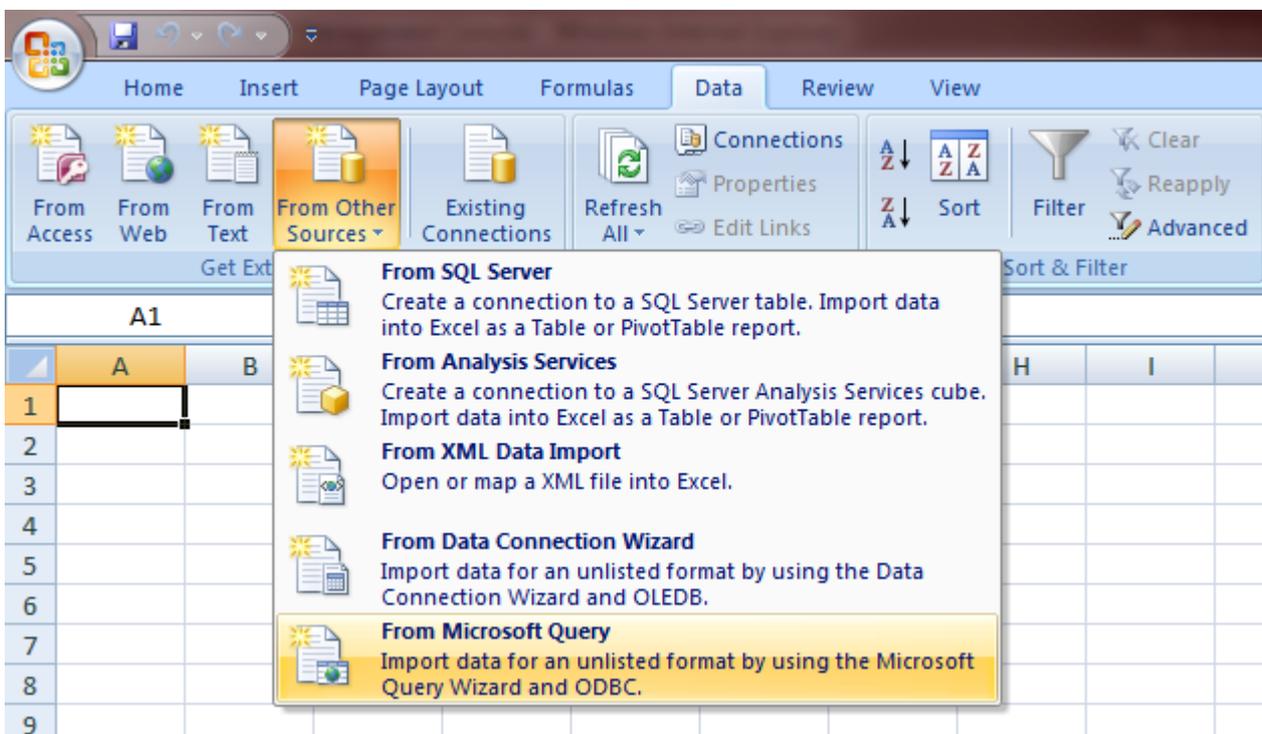
Click **OK**. This will add the 'Exinda SQL Database' to the list of available data sources that can be used by 3rd party applications on this client.



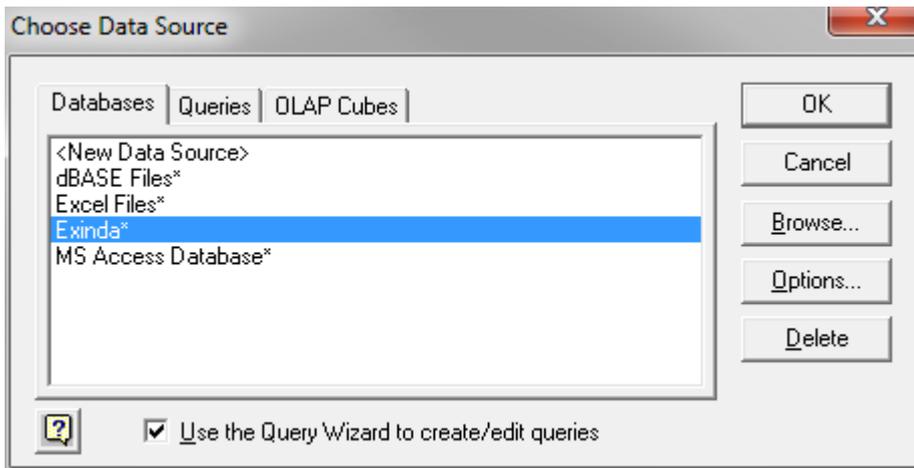
### View SQL Access data in Microsoft Excel

You will need a 3rd party application that is capable of accessing data from ODBC data sources. For the purposes of this How to Guide, we will use Microsoft Excel as an example.

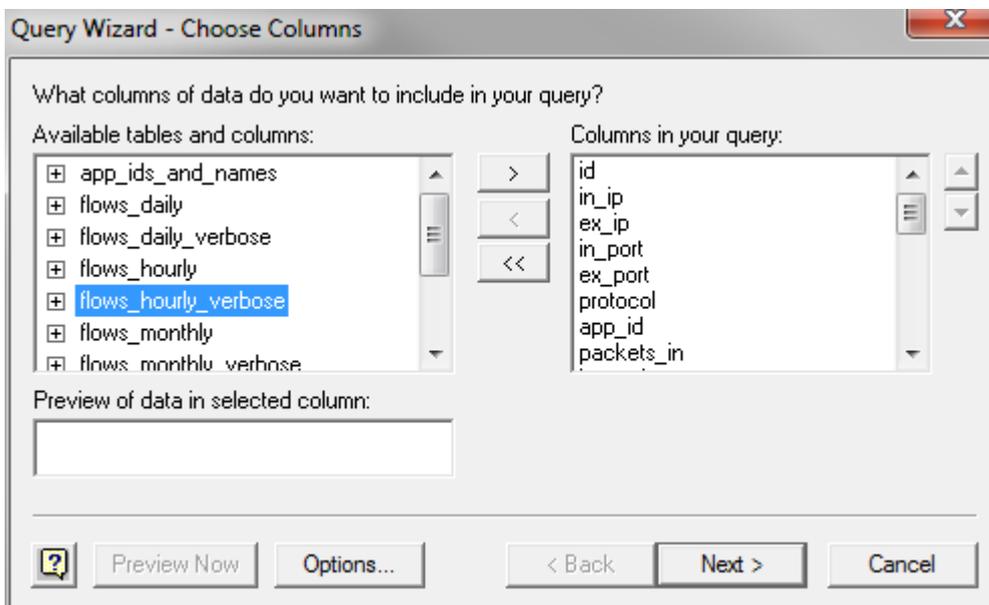
From the **Data** tab in Excel, select **From Other Sources > From Microsoft Query**.



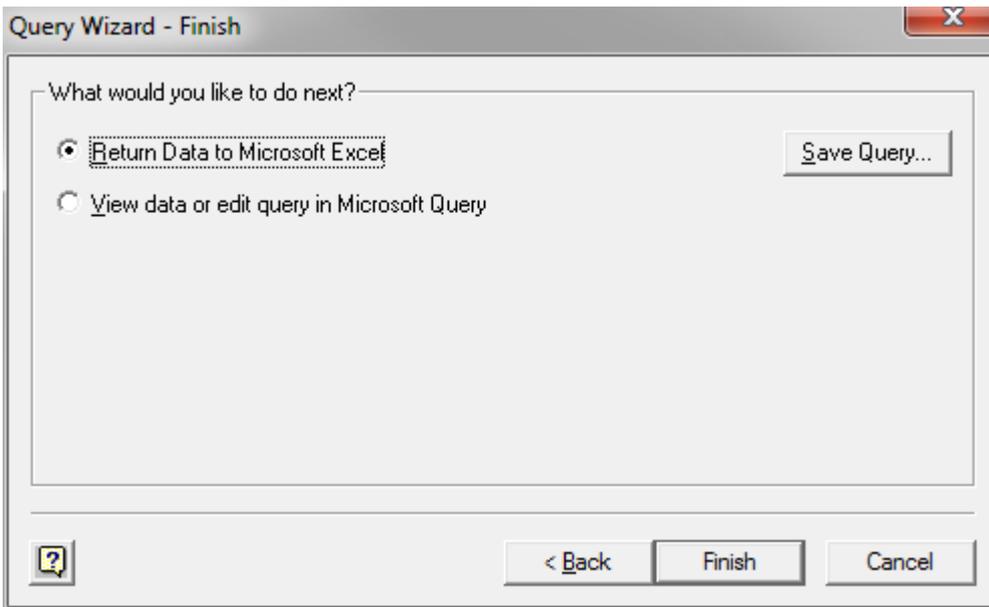
You will be presented with a dialog box that allows you to select the DSN you created in the previous chapter.



Select the **Exinda SQL Database** DSN. This will allow you to choose from the available tables and select the columns to query. Select a table and click the > button to move that table's fields into the list of columns to query.



Click through the wizard, optionally specifying columns to filter or sort by. Then click Finish to return the data to Excel.



The Exinda appliance will now be queried and the data will be returned to the Excel spreadsheet.

|    | A       | B          | C          | D       | E       | F        | G      | H          | I        | J           | K         | L           |
|----|---------|------------|------------|---------|---------|----------|--------|------------|----------|-------------|-----------|-------------|
| 1  | id      | in_ip      | ex_ip      | in_port | ex_port | protocol | app_id | packets_in | bytes_in | packets_out | bytes_out | max_tput_in |
| 2  | 2714022 | 2886729828 | 3197021980 | 0       | 0       | 17       | 222    | 0          | 0        | 6           | 1581      | 0           |
| 3  | 2714021 | 2886729850 | 2523226833 | 0       | 0       | 6        | 201    | 6          | 1104     | 6           | 1621      | 883         |
| 4  | 2714020 | 2886729972 | 3339138632 | 0       | 0       | 6        | 201    | 12         | 3324     | 12          | 1666      | 1329        |
| 5  | 2714019 | 2886729939 | 3494527776 | 0       | 0       | 1        | 201    | 22         | 1760     | 0           | 0         | 448         |
| 6  | 2714018 | 2886729972 | 1249745235 | 0       | 0       | 6        | 207    | 16         | 3184     | 19          | 3825      | 1165        |
| 7  | 2714017 | 2886729877 | 1494265866 | 0       | 0       | 6        | 201    | 7          | 1942     | 13          | 1539      | 1553        |
| 8  | 2714016 | 2886729939 | 3339139912 | 0       | 0       | 6        | 201    | 6          | 2129     | 6           | 877       | 1703        |
| 9  | 2714015 | 2886729939 | 1113983841 | 0       | 0       | 6        | 207    | 7          | 2162     | 8           | 1909      | 1729        |
| 10 | 2714014 | 2886729972 | 1249733985 | 0       | 0       | 6        | 201    | 6          | 1104     | 8           | 2283      | 883         |
| 11 | 2714013 | 2886729882 | 3413282335 | 0       | 0       | 6        | 222    | 119        | 12450    | 114         | 11368     | 919         |
| 12 | 2714012 | 2886729888 | 3510548001 | 0       | 0       | 6        | 201    | 4          | 2359     | 5           | 1317      | 1887        |
| 13 | 2714011 | 2886729828 | 3416333846 | 0       | 0       | 6        | 222    | 211        | 18338    | 241         | 21137     | 896         |
| 14 | 2714010 | 2886730069 | 2149463094 | 0       | 0       | 6        | 201    | 36         | 5620     | 44          | 3580      | 593         |
| 15 | 2714009 | 2886729850 | 2523226710 | 0       | 0       | 6        | 201    | 89         | 70476    | 85          | 14272     | 11439       |
| 16 | 2714008 | 2886729882 | 3406878235 | 0       | 0       | 6        | 201    | 24         | 2839     | 18          | 1330      | 2271        |
| 17 | 2714007 | 2886729855 | 1114779712 | 0       | 0       | 6        | 201    | 6          | 3055     | 7           | 775       | 2444        |
| 18 | 2714006 | 2886729855 | 3452668776 | 0       | 0       | 6        | 201    | 90         | 47511    | 90          | 10534     | 2546        |
| 19 | 2714005 | 2886729888 | 3452668776 | 0       | 0       | 6        | 201    | 6          | 3183     | 7           | 743       | 2546        |
| 20 | 2714004 | 2886729939 | 3494527776 | 0       | 0       | 6        | 201    | 19         | 2552     | 37          | 3483      | 530         |
| 21 | 2714003 | 2886729974 | 2827985172 | 0       | 0       | 6        | 207    | 37         | 7416     | 36          | 4420      | 1507        |
| 22 | 2714002 | 2886729888 | 3539452941 | 0       | 0       | 6        | 201    | 6          | 1131     | 8           | 3813      | 904         |

## SQL Schema

There are a total of 10 tables available for access via SQL.

| Name         | Description                                                                                                   |
|--------------|---------------------------------------------------------------------------------------------------------------|
| flows_hourly | Flow records at an hourly resolution, that is, information for each flow is stored hourly, on the hour.       |
| flows_daily  | Flow records at daily resolution, that is, information for each flow is stored daily, on the day at midnight. |

| Name                 | Description                                                                                                                                                                       |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flows_monthly        | Flow records at monthly resolution, that is, information for each flow is stored monthly, on the 1st day of the month at midnight.                                                |
| urls_hourly          | URL records for each flow record that contain 1 or more urls at hourly resolution, that is, information for each url is stored hourly, on the hour.                               |
| urls_daily           | URL records for each flow record that contain 1 or more urls at daily resolution, that is, information for each url is stored daily, on the day at midnight.                      |
| urls_monthly         | URL records for each flow record that contain 1 or more urls at monthly resolution, that is, information for each url is stored monthly, on the 1st day of the month at midnight. |
| app_ids_and_names    | Application records. The record contains a name, id and a flag to indicate if the application has been deleted. Deleted applications are used when labeling historical data.      |
| summary_applications | Flow records summarized by application. Each record contains information gathered over a 5 minute period.                                                                         |
| summary_hosts_ex     | Flow records summarized by external host. Each record contains information gathered over a 5 minute period.                                                                       |
| summary_hosts_in     | Flow records summarized by internal host. Each record contains information gathered over a 5 minute period.                                                                       |

## flows Table

The following table describes the schema of the flows\_\* SQL tables.

| Field      | Type                    | Description                                                                                                                                                                                   |
|------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id         | unsigned 32-bit integer | A unique id that defines this record. This is the primary key.                                                                                                                                |
| in_ip      | binary (128 bit)        | A 16 byte (128 bit) representation of the internal IPv6 address (the IP address on the LAN side of the Exinda appliance) of the flow. IPv4 addresses are represented as IPv4 mapped format.   |
| ex_ip      | binary (128 bit)        | A 16 byte (128 bit) representation of the external IPv6 address (the IP address on the WAN side of the Exinda appliance) of the flow. IPv4 addresses are represented as IPv4 mapped format.   |
| in_port    | unsigned 24-bit integer | The TCP or UDP port number on the internal side (the LAN side of the Exinda appliance) of the flow. <sup>1</sup>                                                                              |
| ex_port    | unsigned 24-bit integer | The TCP or UDP port number on the external side (the WAN side of the Exinda appliance) of the flow. <sup>1</sup>                                                                              |
| protocol   | unsigned 24-bit integer | The IANA assigned IP protocol number of the flow. See <a href="http://www.iana.org/assignments/protocol-numbers/">http://www.iana.org/assignments/protocol-numbers/</a> for more information. |
| app_id     | unsigned 24-bit integer | The internal Exinda Application ID assigned to this flow. This represents Exinda's classification of the flow - 0 means unclassified.                                                         |
| packets_in | unsigned 64-bit integer | The number of inbound (WAN -> LAN) packets recorded for this flow over the sample period.                                                                                                     |

| Field          | Type                          | Description                                                                                                                                |
|----------------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| bytes_in       | unsigned<br>64-bit<br>integer | The number of inbound (WAN -> LAN) bytes recorded for this flow over the sample period.                                                    |
| packets_out    | unsigned<br>64-bit<br>integer | The number of outbound (LAN -> WAN) packets recorded for this flow over the sample period.                                                 |
| bytes_out      | unsigned<br>64-bit<br>integer | The number of outbound (LAN -> WAN) bytes recorded for this flow over the sample period.                                                   |
| max_tput_in    | unsigned<br>64-bit<br>integer | The maximum inbound (WAN -> LAN) throughput observed for this flow during the sample period.                                               |
| max_tput_out   | unsigned<br>64-bit<br>integer | The maximum outbound (LAN -> WAN) throughput observed for this flow during the sample period.                                              |
| intervals_in   | unsigned<br>24-bit<br>integer | The number of 10 second intervals there was inbound (WAN -> LAN) traffic observed for this flow during the sample period (bps).            |
| intervals_out  | unsigned<br>24-bit<br>integer | The number of 10 second intervals there was outbound (LAN -> WAN) traffic observed for this flow during the sample period (bps).           |
| timestamp      | unsigned<br>32-bit<br>integer | A UNIX timestamp (number of seconds since epoch - 1st Jan 1970) that represents the start of the sample period.                            |
| in_username    | string                        | A string representation of the username that was assigned to the internal IP of this flow when it was created (if available).              |
| ex_username    | string                        | A string representation of the username that was assigned to the external IP of this flow when it was created (if available). <sup>1</sup> |
| rtt            | unsigned<br>32-bit<br>integer | Round Trip Time in milliseconds. A measure of the time a packet takes to leave a device, cross a network and return. <sup>2</sup>          |
| network_delay  | unsigned<br>32-bit<br>integer | A normalized measure of the time taken for transaction data to traverse the network. <sup>2</sup>                                          |
| network_jitter | unsigned<br>32-bit<br>integer | A normalized measure of the network_delay variability. <sup>2</sup>                                                                        |
| server_delay   | unsigned<br>32-bit<br>integer | A normalized measure of the time taken for a server to respond to a transaction request. <sup>2</sup>                                      |
| bytes_lost_in  | unsigned<br>64-bit<br>integer | The number of bytes lost due to retransmissions (WAN -> LAN). <sup>2</sup>                                                                 |
| bytes_lost_out | unsigned<br>64-bit<br>integer | The number of bytes lost due to retransmissions (LAN -> WAN). <sup>2</sup>                                                                 |

| Field | Type                    | Description                                                                              |
|-------|-------------------------|------------------------------------------------------------------------------------------|
| aps   | unsigned 64-bit integer | Application Performance Score. A measure of an applications performance on the network.2 |

» `in_port` and `ex_port` are only defined when the IP protocol is TCP (6) or UDP (17) and the Exinda was unable to classify the flow (so the `app_id` is 0).

» For more information, refer to [Using Application Performance reports](#) (page 332).

The `flows_*` tables are available as views that represent the binary IPv6 addresses in string format. The views tables are `flows*_verbose` (e.g. `flows_hourly_verbose`). The fields are identical to the above except for the following:

| Field              | Type   | Description                                                                                                                                                                           |
|--------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>in_ip</code> | string | A string representation of the internal address (the IP address on the LAN side of the Exinda appliance) of the flow. IPv4 mapped IPv6 addresses are represented as IPv4 dotted quad. |
| <code>ex_ip</code> | string | A string representation of the external address (the IP address on the WAN side of the Exinda appliance) of the flow. IPv4 mapped IPv6 addresses are represented as IPv4 dotted quad. |

### app\_ids\_and\_names Table

The following table describes the schema of the `app_ids_and_names` SQL table.

| Field                     | Type                    | Description                                                                                |
|---------------------------|-------------------------|--------------------------------------------------------------------------------------------|
| <code>app_id</code>       | unsigned 24-bit integer | A unique id that defines the Application. This is the primary key.                         |
| <code>app_name</code>     | string                  | The Application name (e.g HTTP, Hotmail)                                                   |
| <code>deleted_flag</code> | unsigned 8-bit integer  | A flag indicating if the Application has been deleted from the appliance (0 = no, 1 = yes) |

### urls Table

The following table describes the schema of the `urls_*` SQL tables.

| Field                    | Type                    | Description                                                                                                                                                                   |
|--------------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>id</code>          | unsigned 32-bit integer | This id references an id in the corresponding parent <code>flows_*</code> table. There can be multiple url records referencing the same flow id, so this field is not unique. |
| <code>url</code>         | string                  | The URL (host) extracted from the HTTP header of the parent flow.                                                                                                             |
| <code>packets_in</code>  | unsigned 64-bit integer | The number of inbound (WAN -> LAN) packets recorded for this URL over the sample period.                                                                                      |
| <code>bytes_in</code>    | unsigned 64-bit integer | The number of inbound (WAN -> LAN) bytes recorded for this URL over the sample period.                                                                                        |
| <code>packets_out</code> | unsigned 64-bit integer | The number of outbound (LAN -> WAN) packets recorded for this URL over the sample period.                                                                                     |
| <code>bytes_out</code>   | unsigned 64-bit integer | The number of outbound (LAN -> WAN) bytes recorded for this URL over the sample period.                                                                                       |
| <code>max_tput_in</code> | unsigned 64-bit integer | The maximum inbound (WAN -> LAN) throughput observed for this URL during the sample period.                                                                                   |

| Field         | Type                    | Description                                                                                                               |
|---------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------|
| max_tput_out  | unsigned 64-bit integer | The maximum outbound (LAN -> WAN) throughput observed for this URL during the sample period.                              |
| intervals_in  | unsigned 16-bit integer | The number of 10 second intervals there was inbound (WAN -> LAN) traffic observed for this URL during the sample period.  |
| intervals_out | unsigned 16-bit integer | The number of 10 second intervals there was outbound (LAN -> WAN) traffic observed for this URL during the sample period. |

#### NOTE

IDs are only consistent across the same sample periods. For example, IDs in the `urls_hourly` table only reference IDs in the `flows_hourly` table.

### summary\_applications Table

The `summary_application` table summarizes the aggregated data from the Exinda. The following table describes the schema of the `summary_applications` SQL table.

| Field         | Type                    | Description                                                                                                                                                                                   |
|---------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| in_port       | unsigned 24-bit integer | The TCP or UDP port number on the internal side (the LAN side of the Exinda appliance) <sup>1</sup>                                                                                           |
| ex_port       | unsigned 24-bit integer | The TCP or UDP port number on the external side (the WAN side of the Exinda appliance) <sup>1</sup>                                                                                           |
| protocol      | unsigned 24-bit integer | The IANA assigned IP protocol number of the flow. See <a href="http://www.iana.org/assignments/protocol-numbers/">http://www.iana.org/assignments/protocol-numbers/</a> for more information. |
| app_id        | unsigned 24-bit integer | The internal Exinda Application ID assigned to this flow. This represents Exinda's classification of the flow. A zero value should be interpreted as unclassified.                            |
| bytes_in      | unsigned 64-bit integer | The number of inbound (WAN -> LAN) bytes recorded for this flow over the sample period.                                                                                                       |
| bytes_out     | unsigned 64-bit integer | The number of outbound (LAN -> WAN) bytes recorded for this flow over the sample period.                                                                                                      |
| packets_in    | unsigned 64-bit integer | The number of inbound (WAN -> LAN) packets recorded for this flow over the sample period.                                                                                                     |
| packets_out   | unsigned 64-bit integer | The number of outbound (LAN -> WAN) packets recorded for this flow over the sample period.                                                                                                    |
| intervals_in  | unsigned 24-bit integer | The number of 10 second intervals there was inbound (WAN -> LAN) traffic observed for this flow during the sample period.                                                                     |
| intervals_out | unsigned 24-bit integer | The number of 10 second intervals there was outbound (LAN -> WAN) traffic observed for this flow during the sample period.                                                                    |
| timestamp     | unsigned 32-bit integer | A UNIX timestamp (number of seconds since epoch - 1st Jan 1970) that represents the start of the sample period.                                                                               |
| max_tput_in   | unsigned 64-bit integer | The maximum inbound (WAN -> LAN) throughput observed for this flow during the sample period (bps).                                                                                            |
| max_tput_out  | unsigned 64-bit integer | The maximum outbound (LAN -> WAN) throughput observed for this flow during the sample period (bps).                                                                                           |

| Field          | Type                    | Description                                                                                                                       |
|----------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| rtt            | unsigned 32-bit integer | Round Trip Time in milliseconds. A measure of the time a packet takes to leave a device, cross a network and return. <sup>2</sup> |
| network_delay  | unsigned 32-bit integer | A normalized measure of the time taken for transaction data to traverse the network. <sup>2</sup>                                 |
| network_jitter | unsigned 32-bit integer | A normalized measure of the network_delay variability. <sup>2</sup>                                                               |
| server_delay   | unsigned 32-bit integer | A normalized measure of the time taken for a server to respond to a transaction request. <sup>2</sup>                             |
| bytes_lost_in  | unsigned 64-bit integer | The number of bytes lost due to retransmissions (WAN -> LAN). <sup>2</sup>                                                        |
| bytes_lost_out | unsigned 64-bit integer | The number of bytes lost due to retransmissions (LAN -> WAN). <sup>2</sup>                                                        |

» `in_port` and `ex_port` are only defined when the IP protocol is TCP (6) or UDP (17) and the Exinda was unable to classify the flow (so the `app_id` is 0).

» For more information, refer to [Using Application Performance reports](#) (page 332).

### summary\_hosts Table

The following table describes the schema of the `summary_hosts_in` and `summary_hosts_ex` SQL tables. The table fields are identical apart from the `ip` field - this field represent the IPv4 or IPv6 address of an internal host (`summary_hosts_in`) or an external host (`summary_hosts_ex`).

A host is internal if it is on the LAN side of the appliance and external when on the WAN side.

| Field         | Type                    | Description                                                                                                                      |
|---------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| ip            | binary string           | A string representation of the internal or external IPv4 or IPv6 address of the host.                                            |
| bytes_in      | unsigned 64-bit integer | The number of inbound (WAN -> LAN) bytes recorded for this flow over the sample period.                                          |
| bytes_out     | unsigned 64-bit integer | The number of outbound (LAN -> WAN) bytes recorded for this flow over the sample period.                                         |
| packets_in    | unsigned 64-bit integer | The number of inbound (WAN -> LAN) packets recorded for this flow over the sample period.                                        |
| packets_out   | unsigned 64-bit integer | The number of outbound (LAN -> WAN) packets recorded for this flow over the sample period.                                       |
| intervals_in  | unsigned 24-bit integer | The number of 10 second intervals there was inbound (WAN -> LAN) traffic observed for this flow during the sample period (bps).  |
| intervals_out | unsigned 24-bit integer | The number of 10 second intervals there was outbound (LAN -> WAN) traffic observed for this flow during the sample period (bps). |
| timestamp     | unsigned 32-bit integer | A UNIX timestamp (number of seconds since epoch - 1st Jan 1970) that represents the start of the sample period.                  |
| max_tput_in   | unsigned 64-bit integer | The maximum inbound (WAN -> LAN) throughput observed for this flow during the sample period.                                     |
| max_tput_out  | unsigned 64-bit integer | The maximum outbound (LAN -> WAN) throughput observed for this flow during the sample period.                                    |

| Field          | Type                    | Description                                                                                                           |
|----------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------|
| rtt            | unsigned 32-bit integer | Round Trip Time in milliseconds. A measure of the time a packet takes to leave a device, cross a network and return.1 |
| network_delay  | unsigned 32-bit integer | A normalized measure of the time taken for transaction data to traverse the network.1                                 |
| network_jitter | unsigned 32-bit integer | A normalized measure of the network_delay variability.1                                                               |
| server_delay   | unsigned 32-bit integer | A normalized measure of the time taken for a server to respond to a transaction request.1                             |
| bytes_lost_in  | unsigned 64-bit integer | The number of bytes lost due to retransmissions (WAN -> LAN).1                                                        |
| bytes_lost_out | unsigned 64-bit integer | The number of bytes lost due to retransmissions (LAN -> WAN).1                                                        |

For more information, refer to [Using Application Performance reports](#) (page 332).

## 0.19.5 Monitoring Configuration

You can configure details relevant to monitoring charts and the monitoring data that is collected. You can configure how the data is displayed, how the traffic is analyzed for monitoring purposes, which order of resolution methods are tried when resolving IP addresses to hostnames, whether data is collected, and whether collected data is deleted.

For configuring how data is to display, you can specify how many items are shown in the data tables, how many items are shown in the pie charts, and how many characters to show in the URLs.

For analyzing traffic, you can specify whether to recognize traffic according to layer 7 or layer 3 definitions, and how sensitive (or aggressive) to be when attempting to recognize BitTorrent, eDonkey, Skype, and flow detection.

For analyzing traffic for specific application types (Application Specific Analysis Modules (ASAM)), you can specify whether to extract data from Citrix, http, and SSL traffic, whether to identify anonymous proxies in the traffic, whether to analyze VoIP traffic, whether to calculate the performance and health of connections, whether to collect connection symmetry information, and whether to log every URL seen in the traffic.

For resolving IP addresses to hostnames, you can specify which methods are tried first, second and so on: network object, DSN, NetBios name lookup, and IP address.

For collection of monitoring data, you can specify whether to collect data for subnets and virtual circuits, and whether to collect detailed records for applications, hosts, URLs, users, conversations and subnets, and whether to collect data for traffic between internal network objects.

For deleting monitoring data, you can selectively delete various types of data collected by the appliance.

To configure monitoring charts display options

Go to **Configuration > System > Setup > Monitoring** tab - **Monitoring Options** form.

The following fields allow you to modify display options.

- » **Table Items** - Sets the maximum number of top items displayed in the monitoring tables. Acceptable values are 1-1000.
- » **Chart Items** - Sets the maximum number of top items to displayed in the chart and graphs. Acceptable values are 1-10. Note that this value will apply universally to ALL options on the Monitor menu.
- » **Maximum URL Size** - Sets the maximum length of URLs displayed on the Real Time report tables.
- » **Graph Display Options** - Specifies whether the graphs display in Flash or non-Flash format. The default is flash.

- » **Display for application details per subnet** - In a scheduled report, specifies whether the application chart within a subnet displays as a Time series chart (line chart), or as a Pie graph. When this option is selected, the Applications per subnet chart displays in the scheduled report as a line chart whereas all other charts continue to display as a pie graph. The default is Time series chart.
- » **Sort Subnets by Name** - Subnets are sorted by name within scheduled reports if the Enable checkbox is checked; otherwise the subnets are sorted by data volume.
- » **Dual-bridge bypass** - Specifies whether to show flows that touch multiple bridges as separate items per bridge in the real time monitor. If enabled, the same flow will be tracked separately on each bridge. This allows you to see the different policies and transfer rates that are being applied on each bridge, which may be desirable for accelerated back-hauled traffic. If disabled, the same flow through two different bridges will be shown as one merged flow. The real-time monitor will only show the policy of the last packet processed. This is preferred for load balancing, load failover, link bonding, or when seeing asymmetric routes (either locally or in a HA cluster). For more information, refer to [Dual Bridge Bypass](#) (page 570).

To configure how traffic is monitored

Go to **Configuration > System > Setup > Monitoring** tab - **Monitoring Options** form.

The following fields allow you to specify how sensitive the traffic classification analysis should be.

- » **Layer 7 Inspection** - Controls whether to analyze the application signatures within a packet to further classify the traffic within the reports. For example, when analyzing HTTP and FTP traffic and an MPEG file is detected within the packets, the application associated with the connection is changed to MPEG. When disabled, the Layer 7 signatures within packets are not analyzed and any application detection objects with Layer 7 rules are ignored.
- » **Monitor IPv6 Link Local Traffic** - Indicates whether to monitor IPv6 link local traffic, that is, non-routable traffic that is only valid on the single network segment. The default is to not monitor this traffic as it is not representative of your network user's traffic. It is mostly used for network discovery.
- » **OpenVPN Detection** - Indicates the sensitivity for detecting OpenVPN traffic. Setting this to 'aggressive' is the default, however, may result in some false positives. Setting this to 'safe' may result in false negatives.
- » **Bittorent Sensitivity** - Setting this to 'high' is recommended for most service provider environments. Setting it to 'low' is recommended in cases of high false positives.
- » **EDonkey Sensitivity** - Setting this to 'high' is recommended for most service provider environments. Setting it to 'low' is recommended in cases of high false positives.
- » **Skype Sensitivity** - Setting this to 'high' is recommended for most service provider environments.
- » **Reporting Sensitivity** - Controls the minimum number of packets needed to be seen on a flow before it is recorded in the database. Acceptable values are between 1 and 10, with 10 being the lowest sensitivity. Setting this to a low value is not recommended in high load environments. When the sensitivity is set to a low value such as 9, flows that contain less than nine packets over a five minute period are not stored in the database. This prevents port scans from loading hundreds of unnecessary rows of data into the database.

To enable or disable Application Specific Analysis Modules (ASAM)

Go to **Configuration > System > Setup > Monitoring** tab - **ASAM** form.

The Exinda appliance analyzes traffic and attempts to match it against criteria specific to the traffic type. The criteria for matching traffic is defined within Application Specific Analysis Modules (ASAM). Enable and disable the modules that are important for your network.

The following ASAM modules are available:

- » **Anonymous Proxy** - When enabled, the system attempts to anonymous proxies by matching the HTTP hostname and SSL common name against the list of anonymous proxy URLs downloaded by the appliance daily. Disable this

module if it appears that an applications is being misclassified as anonymous proxy.

- » **Citrix** - When enabled, the appliance attempts to extract user names and applications names from Citrix connections. Disable this module to stop the appliance in locations where privacy policy does not permit this type of user identification.
- » **DCE/RPC** - When enabled, this module categorizes client requests for Microsoft services such as MAPI and SMB. This should always be enabled.
- » **HTTP** - When enabled, this module attempts to further analyze connections identified as HTTP and attempts to extract information such as the host, URL, request type, and content type.
- » **Performance Metrics** - When enabled, this module calculates the network delay, server delay, round trip time (RTT), loss, efficiency, and TCP health for TCP connections. Disable this module if the RAM or CPU usage is increasing and affecting the performance of the appliance. For more information refer to [RAM Usage Report](#) and [CPU Usage Report](#).
- » **SSL** - When enabled, this module extracts public certificates from connections identified as SSL and decodes the information from those certificates (such as common name and organization unit).
- » **VoIP** - When enabled, this module extracts VoIP related information such as code type and call quality information (MoS and rFactor scoring) from connections identified as RTP.
- » **Asymmetric route** - When enabled, this module collects connection symmetry information. Disable this module if the network regularly has asymmetric routes, as it is unnecessary to alert administrators that asymmetrical connections are occurring.
- » **URL Logging** - When enabled, every URL seen by the appliance is logged to the database. Specify how long (in days) the data will be saved. This module is disabled by default.

To control the order of resolution methods tried when resolving IP addresses to hostnames

Go to **Configuration > System > Setup > Monitoring** tab - **Host Resolution Method** form.

There are multiple host resolution methods that can be used to resolve IP addresses to hostnames. The system will attempt to resolve the hostname using one of the methods. If that method fails it will try another method. You can determine the order of host resolution methods that the system will use by ranking the first method as 1, the next as 2, and so on.

The options for host resolution methods are the following:

- » **Network Object** - The IP addresses will be resolved according to the configured network objects.
- » **DNS** - The IP addresses will be resolved according to the DNS mappings.
- » **IP Address (no resolution)** - The IP addresses will NOT be resolved to hostnames.
- » **NetBIOS Name Lookup** - The IP addresses will be resolved to NetBIOS names.

To enable or disable collection of monitoring data

Go to **Configuration > System > Setup > Monitoring** tab

Various types of data is collected for traffic passing through the network. If the appliance is not performing as expected, data collection can be disabled to improve performance.

The following data collection can be disabled:

- » **Subnets**(shown in the **Statistics Collection** form) - If disabled, data is not collected for subnet reporting.
- » **Virtual Circuits**(and Applications) (shown in the **Statistics Collection** form) - If disabled, data is not collected for virtual circuit reporting. The collection of global application statistics also will not be collected since the global

application statistics are derived from the virtual circuit stats. Note that application reporting within a subnet is not affected by this setting. That is, if data collection is enabled for subnets and is disabled for virtual circuits, then the applications within a subnet will be reported, but the applications reported across the entire appliance or within a virtual circuit will not be reported.

» **Internal Hosts**(shown in the **Statistics Collection** form) - If disabled, data is not collected for internal hosts. You can disable this option to control the amount of data collected in situations where you have many hosts and want to ensure you do not run out of storage room. To view the amount of storage space allocated and how much is free, see [Allocate Disk Storage for System Services](#). Ensure you enable this option if you want to monitor or produce reports for internal host data or to display internal host data on the Application Performance screens in the [Solution Center](#).

» **External Hosts for Subnets**(shown in the **Statistics Collection form**) - Specify one or more network objects to collect external host data for specific network objects only. In cases where you have created a custom network object related to a specific set of IP Addresses, you can choose the network object to collect only the required data, rather than extraneous data from all objects.

#### NOTE

The amount of statistics collected increases for each network object you specify, which may also increase the amount of time necessary to generate reports that collect external host details. A large number of network objects selected may also increase the usage of the monitoring disk partition.

» **Detailed Record Retention**(shown in the **Monitoring Options** form) - Controls whether detailed monitoring records (Applications, Hosts, URLs, Users, Conversations and Subnets) are stored. If there are excessive traffic flows through the appliance, disabling this option will reduce CPU usage. However, the detailed records will no longer be collected and drill down information for Applications, Hosts, Conversations will no longer be available. Summary information, that is totals for the entire appliance, will be available for Applications, Hosts, and Conversations.

» **Ignore Internal-to-Internal**(shown in the **Monitoring Options** form) - Your network may have network objects on the WAN side of the appliance that have been configured as Internal objects, for example a router or firewall. Enabling the Ignore Internal-to-Internal option prevents traffic between internal network objects being included in the reports.

To delete collected monitoring data

Go to **Configuration > System > Setup > Monitoring** tab - **Clear Monitoring Records** form.

If the appliance is running out of disk space, you can delete collected data.

The following record types can be deleted:

- » **All Interface Records** - Deletes all data associated with the Interfaces charts - Interface Throughput and Interface Packets Per Second charts.
- » **All Network Summary Records** - Deletes all data associated with the Network Summary charts.
- » **All Control/Policy Records** - Deletes all data associated with the Control charts - Policies, Discard, and Prioritization Ratio charts.
- » **All Optimization Records** - Deletes all data associated with the Optimization charts - Reduction and Edge Cache charts.
- » **All SLA Records** - Deletes all data associated with Network Response (SLA) chart.
- » **All APS Records** - Deletes all data associated with Application Performance Score (APS) summary chart.
- » **All APM Records** - Deletes all data associated with Application Performance Metric (APM) charts, which are the detailed metric charts for the APS monitor.

- » **All Detailed Monitor Records** - Deletes all detailed data, that is, deletes all the drill down data for applications, hosts, URLs, users, conversations. Summary information, that is, the totals for the entire appliance will still be available.
- » **All Appliance Records** - Deletes all data associated with the system charts - Connections, Accelerated Connections, CPU Usage, CPU Temperature, RAM Usage, Disk IO, and Swap Usage charts.
- » **All Subnet Records** - Deletes all data associated with subnet charts.

All check boxes can be selected by clicking in the checkbox in the header area.

**CAUTION**

This will permanently delete the selected records from the monitoring database.

### 0.19.6 Netflow Configuration

Netflow allows the Exinda appliance to export flow records to 3rd party monitoring devices.

1. Use the form below to configure these Netflow targets.

**Add New Netflow Collector**

IP Address

Port

Version

| Property   | Description                                                                                                         |
|------------|---------------------------------------------------------------------------------------------------------------------|
| IP Address | Specify the IP Address of the Netflow target. The Exinda appliance will export Netflow data to this IP Address.     |
| Port       | Specify the Port number of the Netflow target. The Exinda appliance currently supports Netflow export on UDP ports. |
| Version    | Specify the Netflow version to export. Current supported versions are v1, v5 and v9.                                |

2. The form below allows customization of the flow records sent by Netflow.

| Common Options                          |                                            |
|-----------------------------------------|--------------------------------------------|
| Active flow timeout                     | <input type="text" value="1"/> minutes     |
| V9 Only Options                         |                                            |
| Use Long (64-bit) Byte Counters         | <input checked="" type="checkbox"/> Enable |
| Use Long (64-bit) Packet Counters       | <input type="checkbox"/> Enable            |
| Netflow Packet Payload Size             | <input type="text" value="1440"/> bytes    |
| Template Refresh Rate                   | <input type="text" value="100"/> packets   |
| Template Timeout Rate                   | <input type="text" value="600"/> seconds   |
| General Options Refresh Rate            | <input type="text" value="10000"/> packets |
| General Options Timeout Rate            | <input type="text" value="600"/> seconds   |
| Username Options Timeout Rate           | <input type="text" value="1440"/> minutes  |
| Inactive Username Expiry Rate           | <input type="text" value="168"/> hours     |
| V9 Optional Fields - General            |                                            |
| Export L7 Application ID                | <input checked="" type="checkbox"/> Enable |
| Export Policy ID                        | <input checked="" type="checkbox"/> Enable |
| Export Type of Service (TOS)            | <input checked="" type="checkbox"/> Enable |
| Export VLAN ID                          | <input checked="" type="checkbox"/> Enable |
| Export Min and Max Packet Sizes         | <input checked="" type="checkbox"/> Enable |
| Export Min and Max TTL                  | <input type="checkbox"/> Enable            |
| Export Flow Direction                   | <input checked="" type="checkbox"/> Enable |
| Export SNMP Input and Output Interfaces | <input checked="" type="checkbox"/> Enable |
| Export output byte and packet counters  | <input checked="" type="checkbox"/> Enable |
| Export username details                 | <input checked="" type="checkbox"/> Enable |
| Export VoIP MOS and rFactor             | <input checked="" type="checkbox"/> Enable |
| Export extra information (hostnames)    | <input checked="" type="checkbox"/> Enable |
| Export traffic class                    | <input type="checkbox"/> Enable            |
| V9 Optional Fields - Metrics            |                                            |
| Export RTT                              | <input checked="" type="checkbox"/> Enable |
| Export Network Delay                    | <input checked="" type="checkbox"/> Enable |
| Export Network Jitter                   | <input checked="" type="checkbox"/> Enable |
| Export Server Delay                     | <input checked="" type="checkbox"/> Enable |
| Export Bytes Lost                       | <input checked="" type="checkbox"/> Enable |
| Export APS Score                        | <input checked="" type="checkbox"/> Enable |

## Common Options:

| Option              | Description                                                                                                                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Active Flow Timeout | Specify how often long-term, persistent flows are exported. By default, flows are exported within 10 seconds of the flow terminating (this approach does not work well for long-term or persistent flows). This setting allows you to specify how often these long-term flows should be exported. |

## Netflow v9 Options:

| Option                        | Description                                                                |
|-------------------------------|----------------------------------------------------------------------------|
| Use Long Byte Counters        | Export byte counters as 64bit values instead of 32bit.                     |
| Use Long Packet Counters      | Export packet counters as 64bit values instead of 32bit.                   |
| Netflow Packet Payload Size   | Set maximum Netflow packet payload size.                                   |
| Template Refresh Rate         | Configure the maximum number of packets between exporting of templates.    |
| Template Timeout Rate         | Configure the maximum number of seconds between exporting of templates.    |
| Options Refresh Rate          | Configure the maximum number of packets between exporting of options.      |
| Options Timeout Rate          | Configure the maximum number of seconds between exporting of options.      |
| Username Options Timeout      | Configure maximum number of minutes between exporting of username options. |
| Inactive Username Expiry Rate | Configure the maximum time to remember inactive usernames.                 |

## Netflow v9 Optional Fields - General:

| Option                       | Description                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Export L7 Application ID     | Export Application identification information. The Application ID to Name mappings are exported as an options template. |
| Export Policy ID             | Export Optimizer Policy IDs and names.                                                                                  |
| Export Type of Service (TOS) | Export minimum and maximum Type of Service (TOS).                                                                       |
| Export VLAN ID               | Export VLAN identifier.                                                                                                 |
| Export Packet Sizes          | Export minimum and maximum packet sizes.                                                                                |
| Export Min and Max TTL       | Export minimum and maximum time-to-live (TTL).                                                                          |
| Export Flow Direction        | Export flow direction.                                                                                                  |
| Export SNMP Interfaces       | Export SNMP input and output interfaces.                                                                                |
| Export Output Counters       | Export output packet and byte counters, these can be compared to input byte and packet counters to calculate reduction. |
| Export Username Details      | Export AD usernames.                                                                                                    |
| Export VoIP MoS and rFactor  | Export MoS and rFactor values for VoIP calls.                                                                           |
| Export Extra Information     | Exports extra flow information, such as domain name for HTTP flows, published application name for Citrix.              |
| Export traffic class         | Export traffic class.                                                                                                   |

Netflow v9 Optional Fields - Metrics:

| Option                | Description                   |
|-----------------------|-------------------------------|
| Export RTT            | Export round trip time (RTT). |
| Export Network Delay  | Export network delay.         |
| Export Network Jitter | Export network jitter.        |
| Export Server Delay   | Export server delay.          |
| Export Bytes Lost     | Export lost bytes count.      |
| Export APS Score      | Export APS score.             |

### 0.19.7 Create a Scheduled Job

Cache pre-population, reboots, and firmware installations can be scheduled to run at a specific date and time, and at a set frequency.

**Add New Job**

|                                                                                    |                                                                                         |
|------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| <b>ID</b>                                                                          | <input style="width: 100%;" type="text" value="5"/>                                     |
| <b>Name</b>                                                                        | <input style="width: 100%;" type="text" value="Monthly Sales Collateral"/>              |
| <b>Comment</b>                                                                     | <input style="width: 100%;" type="text" value="Docs for sales team available 3rd day"/> |
| <b>Enable</b>                                                                      | <input type="button" value="Yes"/> <input type="button" value="↓"/>                     |
| <b>Fail-Continue</b>                                                               | <input type="button" value="Yes"/> <input type="button" value="↓"/>                     |
| <b>Schedule</b>                                                                    | <input type="button" value="Monthly"/> <input type="button" value="↓"/>                 |
| <b>Time</b>                                                                        | <input style="width: 100%;" type="text" value="3:00:00"/> (HH:MM:SS)                    |
| <b>Interval</b>                                                                    | <input style="width: 100%;" type="text" value="1"/> (months)                            |
| <b>Day-of-month</b>                                                                | <input style="width: 100%;" type="text" value="3"/> (-28 to -1 and 1 to 28)             |
| Please enter one or more commands and separate each command with <b>new line</b> . |                                                                                         |
| <b>Commands</b>                                                                    | <div style="border: 1px solid #ccc; height: 100px; width: 100%;"></div>                 |

Screenshot 279: Create the schedule

Where do I find this configuration?

Go to **Configuration > System > Setup > Scheduled Jobs**.

To schedule a job

1. In the **Add New Job** area, type a unique **ID** for the job.
2. Type a **Name** for the job.
3. [Optional] In the **Comment** field, type a description for the job.
4. To enable the job to execute upon next scheduled time, **Enable** the job.
5. If the job should be completed, even if one or more commands fail to execute, set **Fail-Continue** to **Yes**.
6. Set the schedule of the job. Jobs can be set to run Once, Daily, Weekly, Monthly, or Periodically.

- **Once:** Set the time and date when this job should be executed.
- **Daily:** Set the time that this job will execute every day.
- **Weekly:** Set the time and the day of the week that this job will execute.
- **Monthly:** Set the time of day, how frequently it recurs measured in intervals of months, and the day of the month. The day of the month is specified as 1 through 28 (E.g. March 23 would have a day of the month as 23), or the day of the month can be specified as -1 through -28, where it counts from the last day of the month (E.g. March 31 would have a day of the month as -1 and March 23 could be -9.)
- **Periodic:** Set the start time and date and how frequently it recurs as an interval. Start time is specified as HH : MM : SS, start date is entered as YYYY/MM/DD, interval is entered as 2h3m4s.

7. After selecting the schedule of the job, specify the parameters for the schedule. For example, set the time, date, interval, or day-of-the-week when the job runs.

8. In the **Commands** field, type the necessary commands for the job you want run. Each command must be on a new line. For scheduled pre-population jobs, leave the commands field blank. When creating the pre-population object, specify this scheduled job. The CLI for the pre-population object will automatically populate this commands field.

9. Click **Add Job**.

The job is added to the list, and is now available for selection in the Pre-population Object, if desired.

## 0.19.8 Alerts

Alerts will notify you when there are issues or potential issues with either the Exinda appliance system (such as CPU utilization and memory paging) or with your traffic (such as an application performance score dropped). The alerts can either be sent by email or by SNMP traps. Use the alerts to ensure the system and your network is operating the way you need it to.

### NOTE

To email alerts, valid SMTP and email settings are required. For more information, refer to [Email configuration](#) (page 544). Recipients of the email alerts are configured where SMTP is configured.

To send SNMP traps, valid SNMP settings are required. For more information, refer to [SNMP configuration](#) (page 547).

| Name                                 | Enable                              | Send Email                                 | Send SNMP Trap                             | Trigger Threshold | Clear Threshold |
|--------------------------------------|-------------------------------------|--------------------------------------------|--------------------------------------------|-------------------|-----------------|
| CPU Utilization                      | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable | 95 % Busy         | 80 % Busy       |
| Disk Usage                           | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable | 7 % Free          | 10 % Free       |
| Memory Paging                        | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| NIC Collisions                       | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable | 1 %               | 1 %             |
| NIC Link Negotiation                 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| NIC Dropped Packets                  | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| NIC Problems - RX                    | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| NIC Problems - TX                    | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Bridge Link                          | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Bridge Direction                     | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| System Startup                       | <input type="checkbox"/>            | <input checked="" type="checkbox"/> Enable | <input type="checkbox"/> Enable            |                   |                 |
| SMB Signed Connections               | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| SLA Latency                          |                                     | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| SLA Loss                             |                                     | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| APS                                  |                                     | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| APM                                  |                                     | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Redundant Power                      | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Redundant Storage                    | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Connection Limiting                  |                                     | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Max Accelerated Connections Exceeded | <input type="checkbox"/>            | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| Asymmetric Route Detection           | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |
| MAPI Encrypted Connections           | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Enable | <input checked="" type="checkbox"/> Enable |                   |                 |

Some alerts are enabled with no option to disable, but for all alerts you need to decide if you want email notifications and/or SNMP traps. For some alerts, you can specify operational thresholds to trigger or clear the alerts.

### Specified Thresholds Exceeded

- » **SLA Latency** – Alert raised when the specified latency for an SLA object is exceeded. For more information, refer to [Configuring service level agreement objects](#) (page 208).
- » **SLA Loss** – Alert raised when there is loss for a SLA.
- » **APS** – Alert raised when the defined threshold for an APS object is exceeded.
- » **APM** – Alert raised when the defined threshold for an APM object is exceeded.
- » **Connection Limiting** – Alert raised when one or more Virtual Circuits has connection limits enabled, and the threshold was reached.
- » **Max Accelerated Connections Exceeded** – Alert raised when the number of accelerated connections exceeds the licensed limit. Connections over the licensed limit pass through the appliance and are not accelerated.

## Particular Traffic Patterns Detected

- » **Asymmetric Route Detection** – Alert raised when traffic from a single connection comes in to the network through one interface or node, and goes out through another interface or node.
- » **SMB signed connections** – Alert raised when SMB signed connections are present.
- » **MAPI Encrypted Connections** – Alert raised when encrypted MAPI traffic to a Microsoft Exchange server is detected on an Exinda Appliance. Encrypted MAPI traffic cannot be accelerated.

## Appliance Issues

- » **CPU Utilization** – Alert raised when the CPU utilization threshold is reached. The defaults are 95% and 80% busy respectively.
- » **Disk Usage** – Alert raised when the used disk space threshold is reached. The defaults are 7% and 10% free respectively.
- » **Memory Paging** – Alert for memory use and paging.
- » **NIC Collisions** – Alert raised when collisions are present on the interfaces. The defaults are 20 and 1 per 30 sec respectively.
- » **NIC Link Negotiation** – Alert raised when the speed/duplex on an interface is set to auto, but it is negotiating at half duplex and/or 10Mbps.
- » **NIC Dropped packets** – Alert raised when dropped packets are present on the interfaces.
- » **NIC Problems - RX** – Alert raised when RX errors are present on the interfaces.
- » **NIC Problems - TX** – Alert raised when TX errors are present on the interfaces.
- » **System Startup** – Alert raised when the Exinda appliance boots up.
- » **Bridge Link** – Alert raised when one of the links on an enabled bridge is down.
- » **Bridge Direction** – Alert raised when the appliance cabling is incorrect. In most cases, it indicates the Exinda WAN interface has been incorrectly plugged into the LAN and vice versa.
- » **Redundant Power** – Alert raised when one of the power supplies fails (only available on platforms with power redundancy).
- » **Redundant Storage** – Alert raised when one of the hard disks fails (only available on platforms with storage redundancy).

## Enabling System Alerts

Use the following instructions to enable the system alerts.

Before you begin, read through [Alerts](#) for an understanding of what each of the alerts does.

1. Go to **Configuration > System > Setup > Alerts**.
2. For each of the listed alerts, decide upon which you need **Enabled**.
3. For each of the enabled alerts, select the types of notification to receive: **Send Email**, **Send SNMP Trap**, or both.
4. If selecting **CPU Utilization**, **Disk Usage**, or **NIC Collisions** alerts, specify the **Trigger Threshold** and **Clear Threshold** levels that cause the notifications to be sent.

#### NOTE

When the Trigger Threshold is reached, an alert notification is sent to the administrator. When the Clear Threshold values are reached, the notifications stop being sent.

5. Click **Apply Changes**.

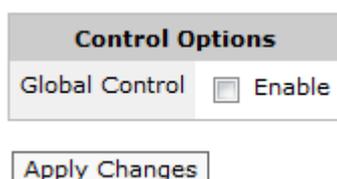
### 0.19.9 Control Configuration

There are two Optimizer modes that affect the behaviour of how Optimizer policies are treated in a multi-bridge deployments.

#### NOTE

To change Control Configuration, navigate to **Configuration > System > System Setup > Control Configuration** on the Web UI, advanced mode.

The form below is used to enable or disable Global Control:



The screenshot shows a web interface for 'Control Options'. It features a section for 'Global Control' with a checkbox that is currently unchecked, followed by the text 'Enable'. Below this section is a button labeled 'Apply Changes'.

**Independent Control (Global Control disabled) (Default):** Optimizer policies are applied to each bridge (LAN and WAN pair) independently. For example, if there were a single policy to restrict all traffic to 1Mbps, this would be applied independently to all bridges. So, the traffic through each bridge would not exceed 1Mbps.

**Global Control (enabled):** Optimizer policies are applied globally, to the entire system. For example, if there were a single policy to restrict all traffic to 1Mbps, this would be applied across all bridges. So, the sum of all traffic through all the bridges would not exceed 1Mbps. This is typically used when you are using multiple bridges and wish to QoS everything as one link.

#### NOTE

Global Control cannot be enabled if Dynamic Virtual Circuits are in use.

In Clustering/HA deployments, Optimizer policies are implemented globally, so this setting only affects how traffic through multiple bridges are treated. For example, a policy to restrict to 1Mbps on an Independent Control system would allow 1Mbps through on each bridge, shared across all appliances (so all the Bridge 0's would share 1Mbps, and all the Bridge 1's would share another 1Mbps, and so on for each bridge). A policy to restrict to 1Mbps on a Global Control system would allow 1Mbps through system wide across all bridges on all appliances.

### 0.19.10 Disk Storage Explained

The Exinda appliance can dynamically change the amount of storage allocated to system services. The Storage Configuration page allows you to see how much disk storage is currently allocated to each system service, as well as the amount currently in use. Users can [re-size](#) and reallocate disk space as required. The storage for particular services, such as cifs, edge-cache, logs, monitor, users, and wan-memory, can be [encrypted](#) if desired.

## CAUTION

- » Resize operation is not permitted on encrypted volumes.
- » Before changing the size of a partition, you must remove the encryption from the partition and put the appliance into Bypass mode. See, [CLI: Bypass](#) or [NIC Settings](#).

Disk Storage Map.



| Storage Configuration           |           |          |     |               |          |           |           |        |         |  |
|---------------------------------|-----------|----------|-----|---------------|----------|-----------|-----------|--------|---------|--|
| Service                         | Status    | Free     |     | Size          | Minimum  | Encrypted | Operation |        |         |  |
| cifs                            | available | 1898.23M | 97% | 1964.00M      | 1024.00M | ✘         | Resize    | Format | Encrypt |  |
| edge-cache                      | available | 569.83M  | 93% | 612.00M       | 1024.00M | ✘         | Resize    | Format | Encrypt |  |
| log                             | available | 974.64M  | 95% | 1024.00M      | 512.00M  | ✘         | Resize    | Format | Encrypt |  |
| monitor                         | available | 9401.04M | 92% | 10.00G        | 10.00G   | ✘         | Resize    | Format | Encrypt |  |
| users                           | available | 974.62M  | 95% | 1024.00M      | 512.00M  | ✘         | Resize    | Format | Encrypt |  |
| wan-memory                      | available | 2149.26M | 95% | 2252.00M      | 5120.00M | ✘         | Resize    | Format | Encrypt |  |
| unallocated storage             |           |          |     | 0.00          |          |           |           |        |         |  |
| <b>Total Available Storage:</b> |           |          |     | <b>18.71G</b> |          |           |           |        |         |  |

Screenshot 280: Disk Storage settings

The [disk storage map](#) shows which services are using disk storage and their current status. It also shows the amount of storage allocated to each service, along with their amount of free space and minimum storage requirements. Services have the capability to be encrypted. It also indicates whether the storage for those services is currently encrypted.

The Disk Configuration section shows a summary of storage by disk partition.

| Disk Configuration |        |           |           |
|--------------------|--------|-----------|-----------|
| Disk               | Status | Size      | Operation |
| sda9               | in-use | 914.22 GB |           |

[Refresh Disk Information](#)

## Related Topics

- » [Configure Storage with CLI](#)
- » [CLI: Storage](#)
- » [CLI: Bypass](#)
- » [NIC Configuration](#)

## The Disk Storage Map

- » **Service** – Services using disk storage
- » **Status** – the status of that storage; the disk storage may be in one of several states, depending on which operation has been selected:

- **available** – Storage is online and available to the service.
  - **growing** – Storage size was increased, and the file system is being reconfigured to use the newly created space.
  - **shrinking** – Storage size was decreased, and the file system is being reconfigured to use the decreased amount of storage available.
  - **formatting** – Storage is being formatted.
  - **checking** – Storage file system is being checked for consistency.
  - **error** – Storage is in an error state. Further information about the error will be displayed in a status message at the top of the form.
  - **unavailable** – Storage is not available.
- » **Free** – Displays the amount of free storage available, shown as the number of bytes as well as a percentage of available space.
  - » **Size** – Shows the total amount of storage allocated for this service.
  - » **Minimum** – Indicates the minimum amount of storage required for this service.
  - » **Encrypted** – identifies whether the storage for the service is currently encrypted or not.
  - » **Operation** – Lists options to perform operations on the storage (resize, format, encrypt).

#### Related Topics

- » [Resizing disk storage for a service](#)
- » [Deleting all data stored for a service](#)

#### Resizing disk storage for a service

The Storage configuration page allows you to see how much disk storage is currently allocated to each system service, as well as the amount currently in use. Users can re-size and reallocate disk space as required.

#### Important notes

- » Resize and format operations are not permitted on encrypted volumes.
- » Put the Exinda Appliance into bypass before changing the partition size of wan-memory. See, [CLI: Bypass](#) or [NIC Settings](#).
- » Storage size can be specified in terms of kilobytes (K), megabytes (M), gigabytes (G), or percentage (%). Use % when entering a storage size to indicate a storage amount as a percentage of free space available. This can be useful when re-allocating storage between services; entering 100% will increase the storage size by the currently unallocated space.
- » When decreasing the amount of storage available to a service, the service may be stopped until the storage operation has completed. If you are decreasing the amount of storage to less than that is currently being used, then the entire content of the storage for the specified service will be discarded.

#### Resizing procedure

Use the following instructions to resize the disk storage for a service. These instructions apply to each service.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.

4. Go to **Configuration > System > Setup > Storage**.
5. Find the entry for the service in the table.
6. In the **Size** column, edit the amount of storage available to a service.
7. On the same row, click **Resize**.

## Disk Storage Encryption

Encryption plays an essential role in securing the integrity of personal data since it ensures that information that could be intercepted by non-authorized third-parties is not accessible.

Exinda Appliance offers the administrators the capacity to encrypt and decrypt different services. Services that have encryption option available are:

| Service           | Description                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------|
| <b>CIFS</b>       | Service used by the Common Internet File System (CIFS).                                         |
| <b>Edge-Cache</b> | Service that catches files for the Edge-Cache feature.                                          |
| <b>WAN-memory</b> | Service used for the data deduplication module of Exinda's Application Acceleration Technology. |
| <b>Log</b>        | Service that enables the storage of debug log information.                                      |
| <b>Monitor</b>    | Service that monitors network object usage.                                                     |
| <b>Users</b>      | Service that monitors user objects' activities.                                                 |

### Important notes

- » In case you wish to switch back to a previous version of Exinda that does not support the encryption of log, monitor and users services, make sure that all these services are not encrypted before trying to boot the previous product version. Failure to do so may result in irreparable errors in the respective service, requiring a fresh installation.
- » Encryption results in more resources being utilized and hence, performance may be affected.
- » Do not interrupt the encryption/decryption operations to minimize the risk of data loss.
- » It is recommended to enable/disable encryption when the Exinda device is not in use to avoid service downtime.
- » During an encryption/decryption operation, the original data is moved to a temporary location on the device and then moved back to the original location. The administrator needs to ensure that the device has enough free space for this operation. The operation requires the size of the encrypted service plus ten percent to ensure that it completes successfully.
- » Resize and format operations are not permitted on encrypted volumes.
- » When Exinda is upgraded from a previous version, in most cases not enough space is allocated to the new log service and the Size column displays zero. The administrator is required to manually free up space on the device and [resize](#) the log service at least to the minimum required value before volume can be encrypted or formatted.
- » A reboot is required after an encrypt/decrypt operation for changes to take effect when encrypting the following services:
  - Monitor
  - Users
  - Log

#### Disk Storage Map.



| Storage Configuration           |           |              |               |          |           |           |        |            |  |
|---------------------------------|-----------|--------------|---------------|----------|-----------|-----------|--------|------------|--|
| Service                         | Status    | Free         | Size          | Minimum  | Encrypted | Operation |        |            |  |
| cifs                            | available | 1898.23M 97% | 1964.00M      | 1024.00M | ✓         | Resize    | Format | No Encrypt |  |
| edge-cache                      | available | 569.83M 93%  | 612.00M       | 1024.00M | ✓         | Resize    | Format | No Encrypt |  |
| log                             | available | 974.64M 95%  | 1024.00M      | 512.00M  | ✗         | Resize    | Format | Encrypt    |  |
| monitor                         | available | 9401.04M 92% | 10.00G        | 10.00G   | ✗         | Resize    | Format | Encrypt    |  |
| users                           | available | 974.62M 95%  | 1024.00M      | 512.00M  | ✓         | Resize    | Format | No Encrypt |  |
| wan-memory                      | available | 2149.26M 95% | 2252.00M      | 5120.00M | ✓         | Resize    | Format | No Encrypt |  |
| unallocated storage             |           |              | 0.00          |          |           |           |        |            |  |
| <b>Total Available Storage:</b> |           |              | <b>18.71G</b> |          |           |           |        |            |  |

Screenshot 281: Disk storage encryption options

### Enabling Encryption

To enable encryption:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Configuration > System > Setup > Storage**.
5. Click **Encrypt** next to the service you want to encrypt.

### Enabling Decryption

To enable decryption:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Configuration > System > Setup > Storage**.
5. Click **No Encrypt** next to the service you want to decrypt.

### Deleting all data stored for a service

Use the following instructions to delete all data from the disk storage for a service. These instructions apply to all services

#### CAUTION

Formatting a services storage will remove all associated application data and should not be necessary in most cases. Contact Exinda Support if you are unsure if this is necessary.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.

5. Click **Configuration > System > Setup > Storage**.
6. Find the entry for the service in the table.
7. On the same row, click **Format**.

## Configure Storage with CLI

The formula used to allocate a default storage size for each system service is shown in the table below, together with an example for the 6062 platforms.

### CAUTION

Before changing the size of a partition, you must remove the encryption on the partition and put the appliance into Bypass mode. See, [CLI: Bypass](#) or [NIC Settings](#).

| Formula                                                                                                                                                                                                                                                                                 | Example for 6062                                                                                                                                                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HDD size: M GB, X GiB<br>Base OS: 14GiB<br>Data Storage: X - 14<br>By default, the data storage is divided up as follows:<br>CIFS: 15%<br>Monitor: 15% or 10GiB, whichever is larger<br>User DB: 1GiB<br>Virt: 50GiB (not available on 2060)<br>WM (wan-memory): 55%<br>Edge Cache: 15% | HDD size: 1000GB (928GiB)<br>Base OS: 14GiB<br>Data Storage: 928 - 14 = 914GiB<br>CIFS: 15% = 58GiB<br>Monitor: 15% = 129GiB<br>User DB: 1GiB<br>Virt: 50GiB (supports virt)<br>WM (wan-memory): 55% = 474GiB<br>Edge Cache: 15% = 129GiB |

Fixed amounts, for example User DB and Virt, are allocated first, then the percentages are used to distribute the remainder.

### NOTE

To change the amount of storage allocated to system services, use the CLI `storage` command. For more information refer to [CLI: Storage](#).

### EXAMPLE

A 6062 is to be used for Control, Monitoring and Edge Cache only. Redistribute the default storage allocated for Virt, CIFS and WM (wan-memory) to the Monitor and Edge Cache services. To show the amount of storage allocated to each service, use the `show storage` CLI command:

```
(config) # show storage
Services: cifs: available - 127.45G free of 129.67G total
edge-cache: available - 127.45G free of 129.67G total
monitor: available 126.97G free of 129.45G total
users: available - 974.62M free of 1024M total
virt: available - 49.04G free of 50G total
wan-memory: available - 467.01G free of 474.65G total
Disks: sda9 (internal): in use - 914.22 GB
Total: 914.22G
Unallocated: 00
```

To redistribute the Logs, CIFS and wan-memory storage, first shrink the amount of storage allocated to these services to the minimum. The minimum size for each service is shown in the table below:

| Service           | Minimum Size |
|-------------------|--------------|
| <b>cifs</b>       | 1 GB         |
| <b>edge cache</b> | 1 GB         |

| Service         | Minimum Size                                  |
|-----------------|-----------------------------------------------|
| monitor         | 10 GiB or current usage, whichever is larger. |
| users           | 500 MiB or current usage, whichever is larger |
| logs            | 500 MiB or current usage, whichever is larger |
| WM (wan-memory) | 5 GB                                          |

To re-size a storage service, use the `storage service <service> size` command, and use the `show storage tasks` command to check the progress:

```
(config) # storage service wan-memory size 5G
(config) # show storage tasks
Storage tasks: Resize wan-memory to 5G: executing
(config) # show storage tasks
No pending tasks
```

Resize the storage for CIFS services:

```
(config) # storage service logs 512M
(config) # storage service cifs 1G
(config) # show storage
Services: cifs: available - 859.88M free of 1024M total
edge-cache: available - 127.23G free of 129.45G total
monitor: available - 126.97G free of 129.45G total
users: available - 974.62M free of 1024M total
logs: available - 363.91M free of 512M total
wan-memory: available - 4879.63M free of 5120M total
Disks: sda9 (internal): in use - 914.22 GB Total: 914.22G
Unallocated: 647.82G
```

There is now 647.82 GB of storage to be allocated to the edge-cache and monitor services. Increase the monitor space by 73.03 GB to a total of 200 GB.

```
(config) # storage service monitor size 200G
(config) # show storage service monitor
Service: monitor Status: growing Encrypted: no Free: 132.07G Size: 200G
```

#### NOTE

The status field is shown as `growing` whilst the resize operation is in progress. When the operation is complete, the status will change to `available`.

Now use the remainder of the disk (approximately 574.79 GB) to increase edge-cache to 706.71 GB (or 723671.04 MB).

```
(config) # storage service edge-cache size 723671.04M
(config) # show storage
Services: cifs: available - 859.88M free of 1024M total
edge-cache: available - 695.43G free of 706.71G total
monitor: available - 196.42G free of 200G total
users: available - 974.62M free of 1024M total
virt: available - 363.91M free of 512M total
wan-memory: available - 4879.63M free of 5120M total
Disks: sda9 (internal): in use - 914.22 GB Total: 219.02G
Unallocated: 12M
```

The status field can be referred while the resize operation is in progress. When the operation is complete, the status will change to `available`.

## NOTE

HDD manufacturers label storage capacity using a base 10 convention, where 1GB = 1,000,000,000 bytes. On the Exinda appliance storage sizes are represented in GiB, where 1 GiB = 1,073,741,824 bytes. So the actual storage of a hard disk, when represented in GiB, is less than what is labeled.

## 0.19.11 Virtualization

The Virtualization feature allows virtual machines to run on select Exinda appliances, allowing third-party operating systems and products to be installed. Once enabled, you can install any standard 32-bit or 64-bit (Intel-based) operating system onto an Exinda Virtualization Partition (EVP).

Most of the virtualization configuration is performed using the `virt` CLI command. In the Exinda Network Orchestrator Interface, you can see which virtual machines are configured and you can power on, power off, or power cycle them. You can also view the VGA console for each virtual machine.

## NOTES

- » Virtualization requires an additional license before this feature can be configured and used. Please contact Exinda TAC or your local Exinda representative if wish to use this feature.
- » The Virtualization menu does not appear if the Exinda Virtualization license has not been installed.

| Virtual Machines         |         |            |                                    |                    |
|--------------------------|---------|------------|------------------------------------|--------------------|
|                          | Name    | Comment    | Status                             | Actions            |
| <input type="checkbox"/> | Replify | Replify VA | Running - IP Address: 172.16.1.242 | Launch VGA Console |
| <input type="checkbox"/> | WinXP   | Windows XP | Running                            | Launch VGA Console |

Screenshot 282: List of Virtual Machines available

The VGA console uses a Java-based applet to launch a secure SSH-based protocol to encrypt the session. You need to have Java installed and enabled as well as direct SSH access to the Exinda appliance in order to use this feature. You also need to authenticate with your username and password.

### Where do I find these settings?

1. Launch the Exinda Network Orchestrator Console.
2. Go to **Configuration > System > Virtualization**.

The Virtualization menu does not appear if the Virtualization license has not been installed.

### What are the hardware requirements?

Only selected Exinda 60 series hardware is supported. Currently, the following hardware is available for Virtualization:

- » Exinda 4062, Exinda 4062 v2
- » Exinda 6060, Exinda 6062
- » Exinda 8060, Exinda 8062, Exinda 8063, Exinda 8064

- » Exinda 10060, Exinda 10062, Exinda 10063, Exinda 10064
- » Exinda 12063, Exinda 12064

An Exinda Virtualization license must be purchased for each Exinda appliance for which you want to enable virtualization.

#### NOTE

Virtualization requires 8 GB of RAM or higher. You may need to purchase a RAM kit.

### To configure and install virtualization

1. Ensure the Virtualization Enable checkbox is checked.
2. Configure virtualization through the CLI with the `virt vm` command.

These commands can be used to create virtual machines, configure CPU, storage, memory, install an operating system, etc. See [CLI: Virtualization](#).

### To manage your virtual machines

1. Select the desired virtual machine(s) by selecting the checkbox(es) beside the virtual machine(s) in the Virtual Machines list.
2. Click the desired action: **Power On**, **Power Off**, or **Power Cycle**.

### To launch the VGA console for a virtual machine

Click the Launch VGA Console button beside the desired virtual machine in the Virtual Machines list. This feature creates a secure SSH connection to the Exinda appliance so that the VGA console can be viewed securely. Therefore, direct SSH access to the Exinda appliance must be available.

Since there is not a trusted certificate as part of the secure SSH connection, a warning is presented that asks you to confirm that you trust the site.

### Troubleshooting: If the console does not appear

If the console does not appear, you may need to add the web address for the Exinda appliance to the Java Exception Site List. This is required since the Exinda appliance has a self-signed certificate for https access, which newer versions of Java are blocking.

### Troubleshooting: If you are unable to log in

If you are unable to log in, it may be due to an enhancement in JDK 7: "The JDK 7 release supports the Server Name Indication (SNI) extension in the JSSEclient. The SNI, described in RFC 4366, enables TLSclients to connect to virtual servers." (from the Oracle support page). If this is the case, the system pops up a java console window with an exception including:

```
Javax.net.ssl.SSLProtocolException: handshake alert: unrecognized_name
```

Options for working around the JDK issue:

- » Enable **Use SSL 2.0 compatible ClientHello format** in the Java console under **Advanced** OR
- » Set your hostname on the Exinda to the Fully Qualified Domain Name (FQDN) that will be used to access the UI and launch the VGA console.

## 0.20 Certificates

Learn managing the certificates and CA certificates that the Exinda Appliance(s) must use to decrypt and re-encrypt the traffic that is flowing into and out of your network.

---

|                                                        |     |
|--------------------------------------------------------|-----|
| 0.20.1 Managing Certificates and CA Certificates ..... | 642 |
| 0.20.2 View all certificates and private keys .....    | 645 |

---

### 0.20.1 Managing Certificates and CA Certificates

When accelerating encrypted traffic, the Exinda appliance transparently decrypts the traffic, performs the relevant application acceleration techniques, such as TCP Acceleration, WAN Memory, or Edge Cache caching, and then re-encrypts the traffic. The Exinda appliances must have access to the appropriate certificates or certificate authority (CA) certificates, and the public keys to decrypt and re-encrypt the traffic. You can import a certificate or a CA certificate, or generate a self-signed certificate or CA certificate.

On the Certificates tab, you can import normal certificates and you can generate untrusted self-signed certificates. Note that the normal certificates may be trusted Certificate Authority (CA)-signed certificates or self-signed certificates. In the **Certificates and Keys** table, you can see a list of all the (non-CA) certificates available on the appliance. You can show, delete, or export any of these certificates.

On the CA Certificates tab, you can import CA certificates and you can generate untrusted self-signed CA certificates. By importing CA certificates, the appliance can offer the entire chain of trust to clients when performing an SSL handshake. In the **CA Certificates and Keys** table, you can see a list of all the CA certificates available on the appliance. You can show, delete, or export any of these certificates.

#### NOTE

Certificates and keys are stored securely on the Exinda appliance. It is not possible to export or view the private key once it has been imported. If you lose the configuration or need to migrate the configuration to another appliance, you must manually load the private key again.

### Import Certificate and Key Details

|                        |                                                                       |            |
|------------------------|-----------------------------------------------------------------------|------------|
| Name                   |                                                                       | (optional) |
| Certificate/Key Format | <input checked="" type="radio"/> PKCS#12<br><input type="radio"/> PEM |            |
| Key Passphrase         |                                                                       | (optional) |
| Certificate File       | <input type="button" value="Browse..."/> No file selected.            |            |
| Private Key File       | <input type="button" value="Browse..."/> No file selected. (optional) |            |

**NOTE**

The interface for importing both Certificates and CA Certificates is the same.

### Generate Certificate and Key Details

|                                        |                 |
|----------------------------------------|-----------------|
| <b>Name</b>                            | localhost       |
| Key Size (bits)                        | 2048            |
| Days Valid                             | 365             |
| Organization Name (eg. company)        | Exinda Networks |
| Organizational Unit Name (eg. section) | Exinda ExOS     |
| <b>Common Name (eg. YOUR name)</b>     | localhost       |

Generate

**NOTE**

The interface for generating both Certificates and CA Certificates is the same.

| CA Certificates and Keys                                                                                                         |                                                |                                                |                                |             |      |               |
|----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|------------------------------------------------|--------------------------------|-------------|------|---------------|
| Name                                                                                                                             | Subject Name                                   | Issuer Name                                    | Validity                       | Private Key |      |               |
|  AAA_Certificate_Services                       | AAA Certificate Services                       | AAA Certificate Services                       | Dec 31 23:59:59<br>2028 GMT    | None        | Show | Delete Export |
|  ACEDICOM_Root                                  | ACEDICOM Root                                  | ACEDICOM Root                                  | Apr 13 16:24:22<br>2028 GMT    | None        | Show | Delete Export |
|  AC_Ra\C3\ADz_Cert\C3\Almara_S.A.               | AC Ra\C3\ADz Cert\C3\Almara S.A.               | AC Ra\C3\ADz Cert\C3\Almara S.A.               | Apr 2 21:42:02<br>2030 GMT     | None        | Show | Delete Export |
|  AOL_Time_Warner_Root_Certification_Authority_1 | AOL Time Warner Root Certification Authority 1 | AOL Time Warner Root Certification Authority 1 | Nov 20<br>15:03:00 2037<br>GMT | None        | Show | Delete Export |

Screenshot 283: List of CA Certificates and Keys (the list is similar on the Certificates tab for Certificates and Keys)

Where do I find these settings?

Go to **Configuration > System > Certificates**.

To learn more about SSL and Exinda's SSL Acceleration

Go to [How SSL Protocol Acceleration Works](#).

To import a certificate

In the **Import Certificate and Key Details** section:

1. Select the **Import Certificate** radio button.
2. (Optional) Type a **Name** for the certificate. If no name is specified, the filename of the certificate is used. Private keys are stored separately from certificates, and are automatically named the same as the certificate, with '\_key' appended to the end.
3. Select the **Certificate/Key Format**.
  - **PKCS#12**—Format used when the certificate and key are stored together, and usually have extensions such as `.pfx` and `.p12`.
  - **PEM**—Common format for certificates issued by Certificate Authorities. PEM certificates usually have extensions such as `.pem`, `.crt`, `.cer`, and `.key`. If PEM format is selected, an additional upload field is exposed so that the private key can be uploaded with the certificate.
4. If the key is password protected, in the **Key Password** field type the password.
5. In the **Certificate File** field, click **Choose File** and navigate to the file to be uploaded to the Exinda Appliance.
6. If the PEM format is selected, the private key must be uploaded. In the **Private Key File** field, click **Choose File** and navigate to the private key file.
7. Click **Import**.

The certificates are displayed in the Certificates and Keys table on the Certificates tab or CA Certificates and Keys table on the CA Certificates tab. From the tables the contents of a certificate can be viewed, or the certificate can be deleted or exported.

To generate a self-signed certificate

To encrypt SSL traffic that passes through the network without requiring the traffic to be signed, a self-signed certificate needs to be generated.

In the **Import Certificate and Key Details** section:

1. Select the **Generate Certificate** radio button.
2. Type a **Name** for the certificate.
3. In the **Key Size** field, specify the number of bits to use when encrypting the contents of the certificate.
4. Specify how many days the certificate is valid for.
5. Type the name of the organization and the name of the area that will be using this certificate.
6. In the **Common Name** field, type the name of the person issuing the certificate.
7. Click **Generate**.

After the certificate has been created, it appears in the list of certificates on the Certificates tab.

To display the contents of a certificate

View the contents of an SSL certificate to see the owner of the certificate, information on the issuer of the certificate, and the time period the certificate is valid.

1. In the **(CA) Certificates and Keys** table, locate the certificate in the list, and click **Show**.
2. To return to the list of certificates, click the **Back** button below the table.

To export a certificate

If an SSL certificate is only available on one appliance, export the certificate so it can be imported onto the other Exinda appliances on the network.

1. In the **(CA) Certificates and Keys** table, locate the certificate in the list, and click **Export**.
2. Select the format for the exported certificate.
3. Click **Save**.

The certificate is downloaded onto the computer accessing the Exinda Web UI.

To delete a certificate

Delete an SSL certificate from the Exinda Appliance when it expires, or becomes invalid.

1. In the **(CA) Certificates and Keys** table, locate the certificate in the list, and click **Delete**.
2. In the confirmation dialog, click **OK**.

The certificate is deleted.

## 0.20.2 View all certificates and private keys

The All Certificates tab displays a list of all Certificate Authority certificates, self-signed certificates, and all base certificates included on the Exinda Appliance. This is the combination of the certificate lists on the Certificates tab and the CA Certificates.

Where do I find these settings?

Go to **Configuration > System > Certificates > All Certificates**.

To display the contents of a certificate

View the contents of an SSL certificate to see the owner of the certificate, information on the issuer of the certificate, and the time period the certificate is valid.

1. In the **All Certificates and Keys** table, locate the certificate in the list, and click **Show**.
2. To return to the list of certificates, click the **Back** button below the table.

To export a certificate

If an SSL certificate is only available on one appliance, export the certificate so it can be imported onto the other Exinda appliances on the network.

1. In the **All Certificates and Keys** table, locate the certificate in the list, and click **Export**.
2. Select the format for the exported certificate.
3. Click **Save**.

The certificate is downloaded onto the computer accessing the Exinda Web UI.

To delete a certificate

Delete an SSL certificate from the Exinda Appliance when it expires, or becomes invalid.

1. In the **All Certificates and Keys** table, locate the certificate in the list, and click **Delete**.
2. In the confirmation dialog, click **OK**.

The certificate is deleted.

## 0.21 Optimization services

The Exinda optimization technology enables applications to run faster over the WAN. Latency in the network affects user productivity and satisfaction with their applications and network. Latency can be due to the sheer volume of data that must be returned for the given application as well as contention for the available bandwidth, the distance that the data must travel while the user is waiting for the data to be retrieved, including the number of back-and-forth communications of "chatty" applications, and failures of the data delivery requiring the data to be retransmitted .

The Network Orchestrator appliance uses a variety of techniques to address these issues. The appliance can reduce the amount of data transmitted over the WAN by using deduplication, compression, and caching techniques. The appliance can minimize delays associated with waiting for the data to be returned by reducing the chattiness of particular protocols and by anticipating requests for data and pre-fetching the data. The appliance can also reduce the frequency of data delivery failures so that data does not have to be retransmitted.

Learn about the optimization services available on the Exinda Appliance(s). There are several services that manage the various kinds of traffic on your network, each of which you can manage.

---

|                                                    |     |
|----------------------------------------------------|-----|
| 0.21.1 How Appliance Discovery Works .....         | 646 |
| 0.21.2 Configuring the Optimization Services ..... | 649 |
| 0.21.3 Universal Acceleration Service .....        | 650 |
| 0.21.4 Protocol-specific Acceleration .....        | 654 |
| 0.21.5 Data caching .....                          | 680 |

---

### 0.21.1 How Appliance Discovery Works

For the most part, acceleration requires two appliances: one to accelerate (such as compressing, or deduplicating data) and one to decelerate (such as recomposing the traffic from the compressed deduplicated traffic). Therefore each appliance must know of the other appliances with which it can accelerate. To find other appliances, the appliances have an auto-discovery process. It is used for two purposes:

- » The discovery of which connections can be accelerated.
- » The discovery of new Exinda appliances on the network.

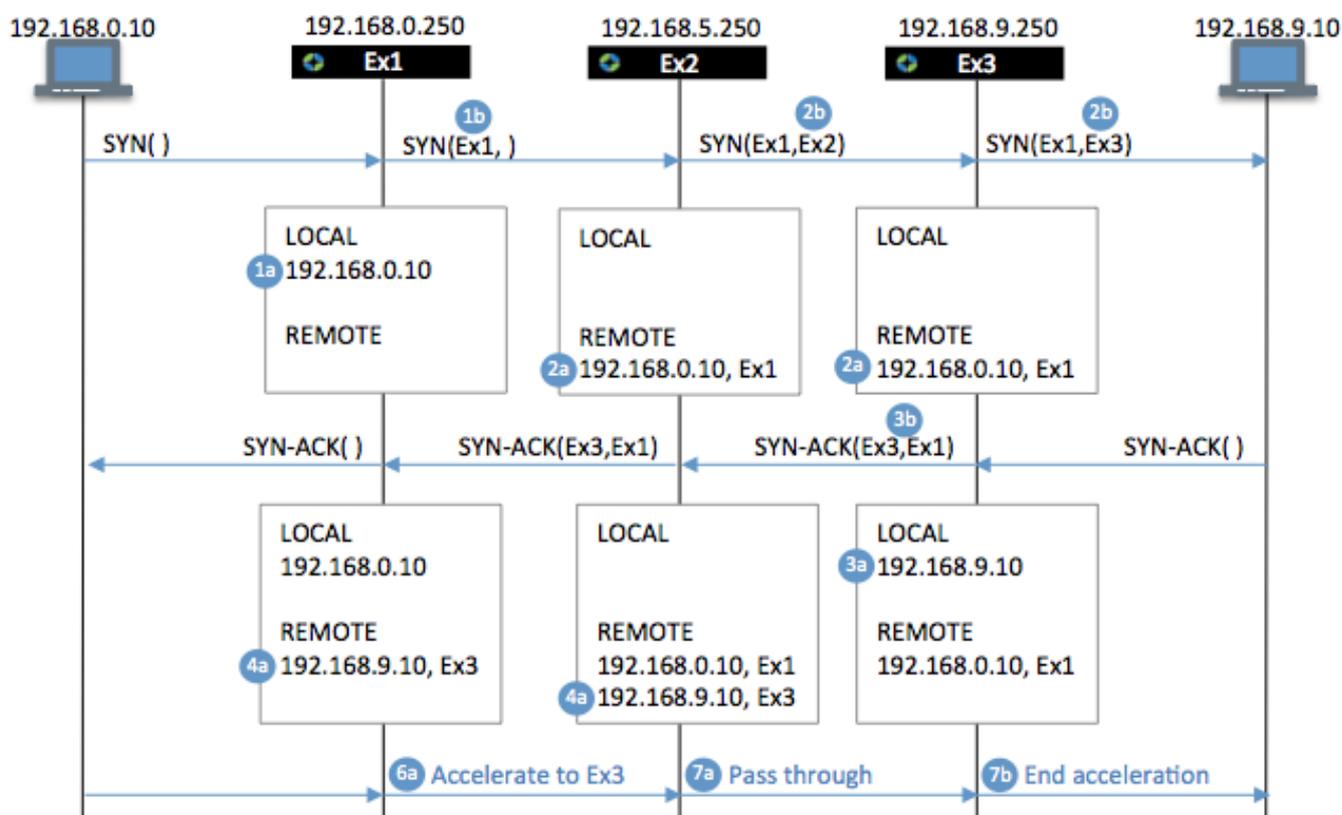
Appliances that can participate in accelerated communications are discovered by detecting extra information in the flow that is required during acceleration of a flow. The appliance adds the following information to the TCP option of SYN, SYN-ACK, and the first ACK packets of each new connection:

- » Source Appliance ID
- » Destination Appliance ID
- » Acceleration Module Map

Optionally an IP address corresponding to one of the appliances can be sent. In addition to this, each appliance must keep a list of the host IP/appliance ID pairs, which indicates which Exinda appliance terminates the acceleration for conversations with a given host IP.

## Discovery Process

The connection discovery process is as follows:



1. When an appliance receives a packet SYN from a client:

- It adds the client IP to its local list. E.g., adding 192.168.0.10 to the Ex1 local list.
- It adds the auto-discovery option to the packet, filling out the source details. E.g., adding Ex1 as the source of the SYN when leaving Ex1.
- If the server exists in the appliance remote list, then the destination field is filled out with the appliance details, otherwise the destination is left blank.

2. When an appliance receives a SYN packet containing the auto-discovery option:

- It records the client IP address and source appliance ID to its remote list. E.g., adding 192.168.0.10, Ex1 to the local lists of Ex2 and Ex3.

- It fills out the destination details and forwards the packet on. E.g., adding Ex2 as the destination of the SYN when leaving Ex2 and replacing the destination of the SYN with Ex3 when leaving Ex3.

3. When an appliance receives a SYN-ACK from the server without any auto-discovery option:

- It adds the server IP to its local list. E.g. adding 192.168.9.10 to the local list of Ex3.
- It adds an auto-discovery option with both the source and destination details filled out. E.g., adding Ex3 as the source and Ex1 as the destination in the SYN-ACK.

4. When an appliance receives the SYN-ACK containing the auto-discovery option, It adds the server IP address and source appliance ID to the remote list. For example, adding 192.168.9.10, Ex3 to the remote list of both Ex1 and Ex2.

5. After the SYN-ACK has passed through, both end appliances know which client or server that they are accelerating for and which other appliance they are accelerating with.

6. When an appliance receives a packet destined for a server, if it finds the source IP address of the packet in its local list and the destination IP address is in its remote list, then it performs acceleration techniques on the packet.

7. When an appliance receives a packet that has been accelerated:

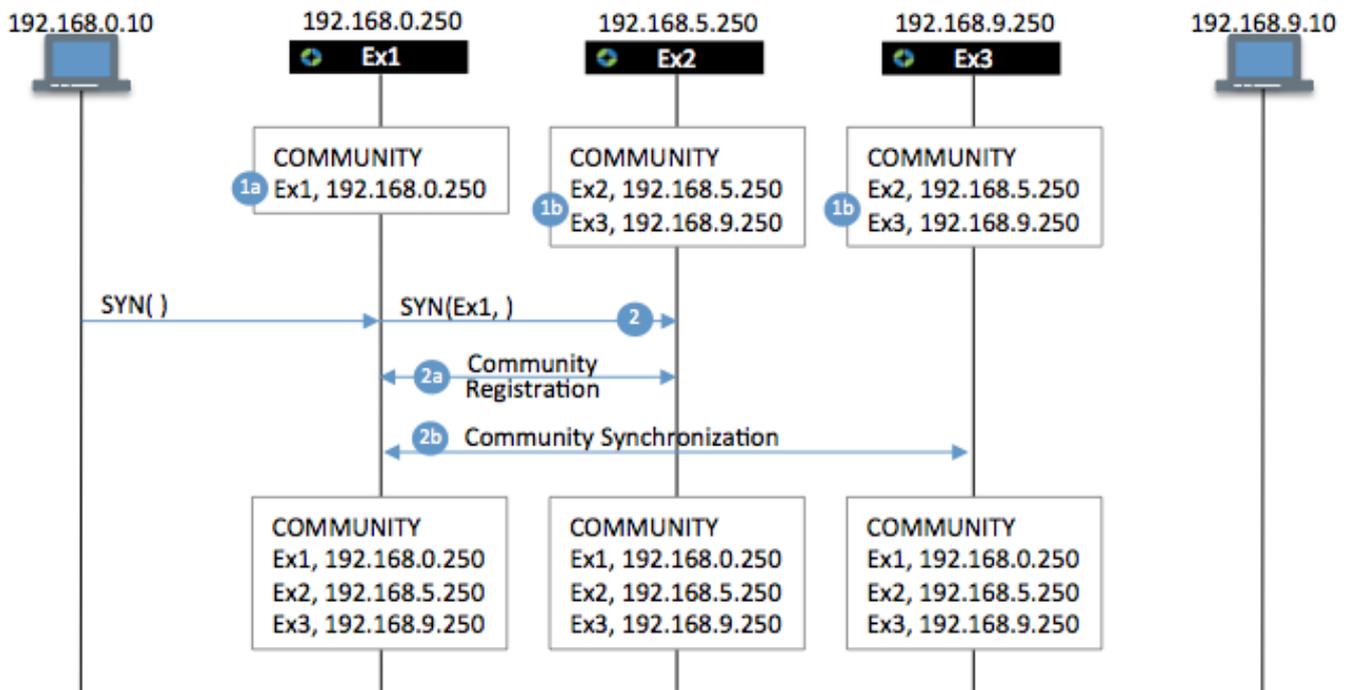
- If it finds that the destination does not refer to itself, then it will ignore all further packets that are part of that connection.
- If it finds that the destination refer to itself, then it will end the acceleration and forward the un-accelerated packets to the server.

## Exinda Community

A group of Exinda appliances in a network is referred to as a community. Exinda appliances that are part of the same community can accelerate to and from each other. The community is generally those Exinda appliances that were automatically discovered.

### NOTE

If an appliance is not automatically discovered, you can manually add the Exinda appliance to the community. To learn how to manually add an appliance to the community, read [Configure Exinda Appliance Community](#).



When an appliance receives an auto-discovery option from a source that the Exinda community does not know about, it can notify the community which will establish a connection to that appliance, and add it to the community.

1. The appliances may have established communities already.

- One appliance may not yet belong to the community.
- Other appliances may belong to the same community.

2. When an appliance receives an auto-discovery option from a source the Exinda community does not know about,

- It establishes a connection to that appliance and adds it to the community.
- It notifies other members of the community.

This may also cause two existing communities to join together.

The Auto Discovery process is very lightweight - it adds negligible latency or delay to packets as they pass through the Exinda appliance.

## 0.21.2 Configuring the Optimization Services

The Manage Optimization Services dialog allows you to start, stop, and disable the optimization services running on the Exinda appliance. Ensure that the service that you need is running. To find the Management Optimization Services controls, go to **Configuration > System > Optimization > Services**.

| Manage Optimization Services     |         |      |         |
|----------------------------------|---------|------|---------|
| NCP Acceleration: <b>Running</b> | Restart | Stop | Disable |
| SMB Acceleration: <b>Running</b> | Restart | Stop | Disable |
| TCP Acceleration: <b>Running</b> | Restart | Stop |         |
| WAN Memory: <b>Running</b>       | Restart | Stop | Disable |
| Exinda Community: <b>Running</b> | Restart | Stop |         |
| SSL Acceleration <b>Running</b>  | Restart | Stop | Disable |
| Edge Cache: <b>Running</b>       | Restart | Stop |         |

### CAUTION

If a service is disabled, any concurrently accelerated connections remain untouched, that is acceleration continues, whereas any new connections are not be able to use the service. When a service is stopped, all accelerated connections (new and concurrent) stop using the service immediately. Stopping services like SMB, TCP, and WAN Memory might cause a failure in currently accelerated connections requiring them to be re-established.

## Managing Optimization Compatibility

To enable compatibility with Exinda appliances that are running older firmware, you can manage the services that interact with these appliances. Go to **Configuration > System > Optimization > Services**. At the bottom of the page you can start, restart, and stop the Exinda Community service (for pre v6.4.0 appliances) and the SMB Acceleration service (for pre v6.3.0 appliances).

| Manage Optimization Compatibility Services    |         |      |
|-----------------------------------------------|---------|------|
| Exinda Community (pre v6.4.0): <b>Running</b> | Restart | Stop |
| SMB Acceleration (pre v6.3.0): <b>Running</b> | Restart | Stop |

### 0.21.3 Universal Acceleration Service

- » **Exinda Community** – Provides appliance auto-discovery and acceleration capability services between all Exinda appliances in the WAN. To learn more, read For more information, refer to [How Appliance Discovery Works](#) (page 646).
- » **WAN Memory** – Provides data reduction using deduplication and compression technology.

## Configure Exinda Appliance Community

A group of Exinda appliances in a network is referred to as a community. Exinda appliances that are part of the same community can accelerate to and from each other. Generally, Exinda appliances automatically discover each other when attempting application acceleration, however, if an appliance is not automatically discovered, you can manually add the Exinda appliance to the community. When the IP address of a manually added Exinda appliance changes, the community node must be updated as well.

| Community Peers |              |                             |                  |        |
|-----------------|--------------|-----------------------------|------------------|--------|
| Hostname        | Host ID      | IP Address(es)              | Firmware Version | Status |
| exinda7         | 00900b2695b4 | <a href="#">10.100.0.7</a>  | 7.0.0.2070       | ONLINE |
| exinda-tor-op   | b8ac6f863261 | <a href="#">10.20.0.205</a> | 7.0.0.2014       | ONLINE |
| ex-beta1-wan    | b8ac6f874f7c | <a href="#">10.10.10.11</a> | 7.0.0.2069       | ONLINE |
| ex-beta1-lan    | b8ac6f879c8f | <a href="#">10.10.10.10</a> | 7.0.0.2069       | ONLINE |
| exinda-tor-op2  | bc305bd6c2ea | <a href="#">10.20.0.206</a> | 6.4.3.2784       | ONLINE |
| weber-exinda    | d4ae528e15a5 | <a href="#">10.0.0.10</a>   | 6.4.3.2784       | ONLINE |

Screenshot 284: List of automatically discovered Exinda appliances

#### NOTE

The Community service uses port 8017 to communicate between Exinda Appliances. Please ensure this port is open for proper functionality.

#### VERSION INFO

- » In a pre-6.4 version, by default, the community was larger than it needed to be, which caused some inefficiencies. In this case, user-defined community groups allow you to create multiple separate smaller Exinda Communities in the same network.
- » In 6.4 and later versions, appliances automatically join the community of appliances with which they are accelerating. If you want your 6.4 or later appliance to belong to a community of pre-6.4 appliances, you need to configure the community settings to match your pre-6.4 appliances.

### Exinda Communities: adding, editing, and removing appliances

Use the following sets of instructions to edit the listing of Exinda community members. For appliances with firmware versions v6.4.0 or later, you do not need to manage communities unless you want such an appliance to join a community created for a pre-6.4.0 appliance group.

#### NOTE

An Exinda appliance can belong to multiple community groups. By default, all appliances belong to the community group with Group ID 0. As a security measure, the **Community Group ID** can be used like a PIN to restrict access to any other Exinda appliance from joining your community.

#### Manually adding an Exinda appliance to the community

1. Go to **Configuration > System > Optimization > Community**.
2. In the **Manually Add New Community Node** area, type a **Name** and the **IP Address** for the Exinda appliance.

| Manually Add New Community Node |                                              |
|---------------------------------|----------------------------------------------|
| Name:                           | <input type="text" value="Waterloo Exinda"/> |
| IP Address                      | <input type="text" value="10.0.0.10"/>       |

3. Click **Apply Changes**.

The appliance is added to the list of manually added community nodes.

| Manually Added Community Nodes |            |                                     |                                       |
|--------------------------------|------------|-------------------------------------|---------------------------------------|
| Name                           | IP Address | Edit                                | Delete                                |
| Waterloo Exinda                | 10.0.0.10  | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |

#### Editing manually added communities

1. Go to **Configuration > System > Optimization > Community**.
2. On the **Manually Added Community Nodes** panel, for the particular appliance, click **Edit**.

| Manually Added Community Nodes |              |                                     |                                       |
|--------------------------------|--------------|-------------------------------------|---------------------------------------|
| Name                           | IP Address   | Edit                                | Delete                                |
| A                              | 192.168.5.20 | <input type="button" value="Edit"/> | <input type="button" value="Delete"/> |

Screenshot 285: The edit screen opens.

Define the manually added community nodes that exist on this host.

| Manually Add New Community Node                                                    |                                           |
|------------------------------------------------------------------------------------|-------------------------------------------|
| Name:                                                                              | <input type="text" value="A"/>            |
| IP Address                                                                         | <input type="text" value="192.168.5.20"/> |
| <input type="button" value="Apply Changes"/> <input type="button" value="Cancel"/> |                                           |

3. Modify the name or IP address of the appliance.

4. Click **Apply Changes**.

#### Removing manually added Exinda appliances from the community

1. Go to **Configuration > System > Optimization > Community**.
2. To remove individual appliances, on the **Manually Added Community Nodes** panel, find the appliance and click

the **Delete** button next to its entry.

3. To remove all appliances from the community, click **Remove all community peers from system**.

## WAN Memory

WAN Memory is the data deduplication module of Exinda's Application Acceleration Technology. It is a bi-directional and universal byte-level cache that stores repetitive patterns on the Exinda appliances's hard disk drive and uses these patterns to compress accelerated traffic between two or more Exinda appliances.

In addition to data deduplication, the appliance can be configured to compress accelerated traffic with a standard LZ compression algorithm. The WAN memory cache can be configured to be persistent such that it survives a system reboot. When using the appliance in cluster mode, WAN memory can be configured to synchronize the WAN memory caches to all members of the cluster.

### VERSION INFO

The LZ compression algorithm changed in version 6.4.3. For appliances that are accelerating traffic between each other, this traffic must be accelerated using the same compression algorithm. Version 6.4.3 contains both the old and the new compression algorithm. Therefore, 6.4.3 and earlier appliances will successfully accelerate and 6.4.3 and later appliances will successfully accelerate. If you have an environment where you have pre-6.4.3 and post-6.4.3 appliances, it is recommended that you install 6.4.3 on your core appliances so that regardless of the versions that the branch appliances are using, they will successfully accelerate with your core appliances.

The WAN memory cache can be cleared immediately or upon the next restart of the WAN memory process.

Configure WAN Memory options.

| WAN Memory Options |                                                               |
|--------------------|---------------------------------------------------------------|
| LZ Compression     | <input checked="" type="checkbox"/>                           |
| Persistent cache   | <input checked="" type="checkbox"/>                           |
| HA cache sync      | <input checked="" type="checkbox"/> When in cluster mode only |

Clear WAN Memory Cache

| WAN Cache Options     |                                                                    |
|-----------------------|--------------------------------------------------------------------|
| Force Data Expiration | <input type="button" value="Expire"/>                              |
| Reset Persistent Data | <input type="button" value="Reset"/> Requires a Wan Memory Restart |

### NOTE

Each Exinda appliance running WAN Memory will connect to each other in order to maintain cache synchronization. This communication happens over TCP port 8013, so this port must be open and available between all Exinda appliances. For security purposes, data sent across these WAN Memory synchronization connections is obfuscated.

Where do I find these settings?

Go to **Configuration > System > Optimization > WAN Memory**.

To configure WAN memory cache

1. Select **LZ Compression** to compress accelerated traffic with a standard LZ-based compression algorithm. This is in addition to data deduplication operations.
2. Select **Persistent Cache** to ensure that WAN memory patterns will be stored on the Exinda appliance's hard disk and will survive a system reboot.
3. Select **HA Cache Sync** when in cluster mode so that WAN memory caches will be mirrored to the WAN memory caches on the other appliances in the cluster.
4. Click **Apply Changes**.

To clear WAN memory cache

Do one of the following:

- » **Force Data Expiration** by clicking **Expire**. This will remove any patterns stored on the Exinda appliance's hard disk drive. Note that this may take several minutes depending on the amount of data.
- » **Reset Persistent Data** by clicking **Reset**. The next time the WAN memory process restarts, it will load any persistent data from the hard disk drive. Using this function and restarting the WAN Memory service is a quick way to clear the WAN Memory

## 0.21.4 Protocol-specific Acceleration

- » **TCP Acceleration** – Provides layer 4 (TCP) protocol optimization.

The TCP protocol can be optimized by establishing a protocol tunnel to avoid subsequent 3-way TCP handshake chattiness and ensuring that the tunnel is kept alive. TCP Acceleration also allows the administrator to set the TCP receive window size to optimize the amount of data in flight given the environment characteristics and to set congestion control algorithms to best match the environment. TCP Acceleration will also reduce chattiness and the amount of data on the wire by acknowledging the receipt of packets in batches instead of acknowledging each packet individually. TCP Acceleration also notifies ECN-aware (Explicit Congestion Notification) routers without dropping packets.

- » **SSL Acceleration** – Provides acceleration for SSL encrypted connections.

SSL Acceleration provides acceleration of SSL encrypted TCP sessions by intercepting SSL connections to configured servers and decrypting them, performing acceleration techniques, then re-encrypting them again. Only traffic to servers that are explicitly configured is SSL accelerated. Any SSL traffic that the Exinda appliance sees that does not belong to a configured server is ignored.

- » **SMB Acceleration** – Provides layer 7 SMB1 and SMB2 (Windows File Sharing) protocol optimization.

SMB (Server Message Block), operates as an application-layer network protocol used for providing shared access to files, printers, serial ports, and miscellaneous communications between nodes on a network. May also be known as CIFS (Common Internet File System), where CIFS is a dialect of the SMB protocol.

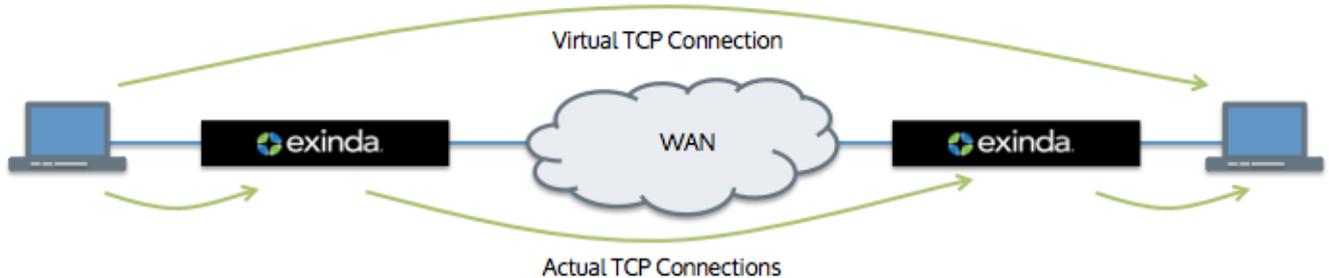
SMB Acceleration is the file transfer specific component of the Exinda Application Acceleration Technology. To deal with inefficiencies in the SMB protocol, the Exinda Appliance has several optimizations to improve the performance of applications using this protocol, including reading ahead of the data stream, writing behind the data stream, and caching meta data on files and folders.

- » **NCP Acceleration** – Provides layer 7 NCP (NetWare Core Protocol over TCP port 524) protocol optimization.

NCP is used in some products from Novell. NCP is used to access file, print, directory, clock synchronization, messaging, remove command execution, and other network service functions in these Novell products.

## Configure TCP Acceleration

TCP Acceleration is the heart of Exinda's application acceleration technology. All accelerated connections pass through TCP Acceleration. In order to accelerate traffic over the WAN, Exinda transparently proxies TCP connections at each end. Both the client and server think they have established a connection with each other; however, they have connected with their local Exinda devices.



In addition to facilitating other acceleration technologies like WAN Memory and SMB acceleration, TCP acceleration also provides performance improvements over and above regular TCP, while being fully compliant with TCP.

You can configure various settings for TCP acceleration.

- » Acceleration occurs between two Exinda appliances. The appliances can discover other appliances on the network that they can accelerate with. You can configure whether your appliance will attempt to discover other appliances and which IP address your appliance will use when other appliances discover your appliance.
- » TCP acceleration and auto-discovery require the use of custom TCP options. If other equipment is stripping the required TCP option from the TCP headers, you can specify that the TCP accelerated traffic must tunnel to avoid the options being stripped. Once a connection is setup via the protocol tunnel, subsequent connections can use the protocol connection and avoid the 3-way TCP handshake, which will reduce the number of TCP connections traversing the WAN and reduces the TCP connection setup time.
- » Also TCP option 30 which historically has been used to indicate Exinda acceleration, has been assigned to indicate multi-path TCP. Exinda now uses both option 30 and option 230 to indicate Exinda acceleration. You can specify which option code should be used in acceleration. Your choice will depend on the Exinda appliance version you are using and whether you are seeing multi-path TCP traffic in your network. Also you can indicate not to accelerate traffic (bypass acceleration) when traffic is using multi-path TCP.
- » Throughput is limited by two windows: the TCP receive window size and the congestion control window. Setting the appropriate sizes for these windows can ensure efficient use of the available bandwidth:
  - Windows scaling increases the TCP window size, which allows more data to be in-flight before TCP requires acknowledgments. This means higher throughput can be achieved on WAN links with higher levels of latency.
  - Selectable congestion control algorithms can be chosen to match the WAN environment. For instance, for high speed high latency links, High Speed TCP should be used. For Satellite links, or other high-latency links, TCP Hybla should be used. This allows for better TCP performance over different WAN technologies.
  - Adaptive Initial Congestion Window allows automatic adjustment of the initial window size depending on the connectivity properties of the WAN link between the Exinda appliances.
  - Slow Start with Congestion Avoidance is used to reset the send window size temporarily to avoid congestion.
- » TCP keep-alive signals prevent the link between accelerated appliances from being broken. You can set whether to use keep-alives and how frequently to send the keep-alive signals.
- » When accelerating traffic in a backhauled setting, the dual-bridge bypass setting will ensure that acceleration processing happens on only one bridge so that the traffic is not re-accelerated. This option, by default, is disabled.

» When you know that a particular Exinda is always at the end of an acceleration chain, you can indicate that it is the end and therefore should not pass through option 30 packets. This is useful when the traffic is transported to a server or firewall that does not know how to handle option 30 packets or when the traffic is forwarded to the Internet. Note that you should not use this setting if the Exinda appliance is also an Exinda in the middle in some scenarios.

Furthermore, TCP acceleration also provides performance improvements without requiring configuration:

- » Delayed and Selective Acknowledgments are used to acknowledge the receipt of packets in batches, instead of acknowledging every single packet. This reduces the amount of return data on the wire.
- » Explicit Congestion Notification (ECN) allows end-to-end (between the Exinda appliances) notification of network congestion without dropping packets. Traditionally, TCP/IP networks signal congestion by dropping packets. When ECN is enabled, an ECN-aware router may set a mark in the IP header instead of dropping a packet in order to signal impending congestion.

Use the form below to configure various TCP Acceleration settings.

### TCP Acceleration Options

|                                               |                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Appliance Auto-Discovery                      | <input checked="" type="checkbox"/>                                                                                                                                                                                                                                                                               |
| Appliance Auto-Discovery IP Address           | <input type="text" value="auto"/>                                                                                                                                                                                                                                                                                 |
| Transport Type                                | <input type="text" value="Transparent TCP"/>                                                                                                                                                                                                                                                                      |
| Window Scaling Factor                         | <input type="text" value="5 / 2M"/>                                                                                                                                                                                                                                                                               |
| Congestion Control                            | <input checked="" type="radio"/> [cubic] General Purpose<br><input type="radio"/> [hybla] Satellite (High speed, high round-trip-time)<br><input type="radio"/> [highspeed] High speed<br><input type="radio"/> [veno] Wireless (Loss handling)<br><input type="radio"/> Other: <input type="text" value="reno"/> |
| TCP Keep-Alive Enabled                        | <input checked="" type="checkbox"/>                                                                                                                                                                                                                                                                               |
| TCP Keep-Alive Timeout                        | <input type="text" value="3600"/> seconds                                                                                                                                                                                                                                                                         |
| Dual-bridge bypass                            | <input checked="" type="checkbox"/>                                                                                                                                                                                                                                                                               |
| Acceleration TCP Options Mode                 | <input type="text" value="30+230: Send 230 where possible, 30 when unsure, accept both (default)"/>                                                                                                                                                                                                               |
| Multi-Path TCP (MPTCP) Acceleration Bypass    | <input checked="" type="checkbox"/>                                                                                                                                                                                                                                                                               |
| End Acceleration (no acceleration on the LAN) | <input type="checkbox"/>                                                                                                                                                                                                                                                                                          |

Where do I find these settings?

Go to **Configuration > System > Optimization > TCP**.

How to use the settings?

- » **Appliance Auto-Discovery** – Allows the appliance to discover other Exinda appliances in its community.
  - When enabled, the Exinda appliance attempts to automatically discover other Exinda appliances on the network when making acceleration. Exinda appliances do this by injecting TCP option 30 into any TCP-SYN packets that the Exinda appliance is attempting to accelerate. If unknown TCP options are removed or blocked by other equipment in your network (e.g., VPN terminators, firewalls, IPS/IDS systems, etc.) then auto-discover may

not work or traffic may be blocked.

- When disabled, or if TCP option 30 is stripped or blocked by other equipment on your network, you will need to manually specify the location of another Exinda appliance in your network on the [Configuration > System > Optimization > Community](#) page.

» **Appliance Auto-Discovery IP Address** – Allows you to set the IP address that identifies this appliance when other appliances are trying to discover their appliance community. Usually this is the management IP address or the IP address to which the default route is associated.

- **auto** – The appliance auto-detects its own IP address.
- **<IP address>** - Select an IP address from the list.

» **Transport Type**– Allows you to specify whether acceleration should be fully transparent or tunneled. Exinda acceleration makes use of TCP option 30 and/or TCP option 230. This setting makes the use of these options transparent or hidden.

- **Transparent TCP** – Ensures Exinda's Application Acceleration is fully transparent. Source and destination IP addresses and port numbers are maintained on all accelerated connections, so any equipment in between two accelerating Exinda appliances can still see correct IP and TCP headers.
- **Protocol Tunnel**– When equipment in between two accelerating Exinda appliances will strip the TCP options, you will need to use the Protocol Tunnel transport type so that other equipment cannot strip these required option codes.

» **Window Scaling Factor**– Specifies how large the TCP receive window can be, according to the formula: TCP Window Size kB = 64 kB x 2 ^ Window Scaling Factor. The scaling factor ranges between 0 and 14, meaning that the receive window can be between 64 kB and 1 GB. The TCP window scaling factor is needed for efficient transfer of data when the bandwidth-delay product is greater than 64kB. The default scaling factor is 5 resulting in a window size of 2MB. Larger window sizes result in more potential memory usage. However, the window size may need to be increased in high-bandwidth, high-latency environments to ensure the bandwidth is fully utilized.

» **Congestion Control**– Indicates which congestion control algorithm should be used. The most common congestion control algorithms are listed together with their intended usage. Set this according to the type of WAN the Exinda appliances are deployed into. This setting only affects outbound traffic to the WAN, so the same setting should be applied to all Exinda appliances on the WAN.

» **TCP Keep-Alive Enabled** – Enables the sending of keep-alive packets on the WAN.

» **TCP Keep-Alive Timeout** – Specifies the amount of time, in seconds, that a connection may be idle before sending keep-alive packets is enabled. Keep-alive packets are sent once per minute until either a response is received, or 5 minutes passes. If five minutes passes without a response the connection is terminated.

» **Dual Bridge Bypass**– Specifies whether acceleration should be handled on a single bridge or multiple bridges when traffic is passing through an Exinda appliance multiple times.

- **Enabled**– All acceleration processing is performed on one bridge only. This is desirable for accelerated back-hauled traffic.
- **Disabled**– Acceleration processing is performed on whatever bridge it arrives on. This is desirable for asymmetric routing or link aggregation.

To learn more about this feature, see [Dual Bridge Bypass](#).

» **Acceleration TCP Options Mode** – Specifies which option code to use in tagging accelerated packets. Historically, Exinda appliances used TCP option 30, however, TCP option 30 has been assigned to indicate multi-path TCP. Exinda started using option 230 in v7.0.

- **30+230: Send 230 where possible, 30 when unsure, accept both (default)** – The system will respond with whatever the remote used. If both options 230 and 30 are available, option 230 will be used. This is the best choice when some of your appliances are v7.0 and later and some are v6.x and earlier. This is a compatibility mode.
  - **230: Send and accept options 230 only** – Only option 230 will be used. Packets with option 30 will not be looked at. Use this when all your appliances are v7.0 and later.
  - **30: Send and accept options 30 only** – Only option 30 will be used. Packets with option 230 will not be looked at. Use this when all your appliances are earlier than v7.0.
  - **230-compat: Send options 230, accept incoming options 30** – Always sends using option 230 and uses compatibility mode for receiving packets, that is, it will handle receiving both option 230 and option 30.
  - **30-compat: Send options 30, accept incoming options 230** – Always sends using option 30 and uses compatibility mode for receiving packets, that is, it will handle receiving both option 230 and option 30.
- » **Multi-Path TCP (MPTCP) Acceleration Bypass** – Specifies whether to attempt acceleration if the traffic is identified as multi-path TCP and falls into an acceleration policy.
- When enabled and the traffic falls into an acceleration policy, multi-path TCP flows are not accelerated. The multi-path TCP options are not stripped and the flows will continue to work in a multi-path TCP fashion.
  - When disabled and the traffic falls into an acceleration policy, the multi-path TCP options will be stripped and acceleration will be attempted.
  - When the traffic does not fall into an acceleration policy (regardless of this setting), the multi-path TCP options will not be stripped and the flows will work in a multi-path TCP fashion.
- » **End Acceleration (no acceleration on the LAN)** – Forces acceleration to end on this appliance.
- When enabled, any incoming acceleration connections on the WAN will be terminated at this appliance and no attempt will be made to find another appliance on the LAN interface. This has no effect on accelerated connections arriving on the LAN.

Consider traffic passing from the client to the server through two accelerated Exinda appliances:

```
Client -> (LAN-side) Exinda (WAN-side) -> WAN -> (WAN-side) Exinda (LAN-side) -> Server
```

Normally, the server side Exinda would send out an option 30 packet to the server. However, if the server does not know how to handle with an option 30, it will return a SYN/ACK without an option 30. Enabling this setting allows the server-side Exinda to know that it is the last appliance in the chain and so it will not send out a SYN with option 30 and it terminates the acceleration connection.

In addition to stopping this appliance from sending option 30 packets to servers that are known to not handle them, it also reduces the timeouts that happen with protocol 139 when attempting to accelerate past the last appliance. It allows servers/firewalls that refuse options to work. It prevents sending random options out to the Internet, which is the case in an accelerated backhauled traffic environment with only a single pair of Exinda appliances. If you have a hub-and-spoke topology then you will not want to enable this setting.

## Configure SSL Acceleration Servers

SSL Acceleration provides acceleration of SSL encrypted TCP sessions by intercepting SSL connections to configured servers by decrypting these sessions, performing acceleration techniques, and then re-encrypting them. Only traffic to the configured servers is SSL accelerated. Any SSL traffic that the Exinda appliance sees that does not belong to a configured server is ignored.

By configuring the SSL Acceleration Server, you are specifying:

- » The location of the server (IPv4 address and port)
- » The SNI (Server Name Indication) which is the hostname of a virtual host when multiple secure websites are hosted on a single host where you want each website to use its own certificate
- » Which certificate is used to re-encrypt the traffic
- » Which certificate is used to authenticate the traffic and what type of validation to perform using that certificate
- » If any CA validation is chosen, then you can choose whether to check if that CA certificate is still valid or whether it has been revoked.

**NOTE**

If the revocation check cannot be done or the certificate has been revoked, then the SSL Acceleration Server is disabled. If the OCSP Responder is offline, the server is disabled. The appliance periodically tests the connection and re-enables the server when it is back up. If the OCSP response verification fails or if the certificate has been revoked, then the connection is reset and the server is disabled.

**NOTE**

If there are any problems with the certificate or key associated with a configured SSL server (e.g., a missing key, or an expired certificate), then SSL Acceleration ignores that traffic until the issue is resolved. The traffic may still be accelerated, just not SSL-accelerated.

**NOTE**

The SSL Acceleration service uses port 8018 to communicate between Exinda Appliances. Please ensure this port is open for proper functionality

| Add SSL Acceleration Server |                                          |
|-----------------------------|------------------------------------------|
| Name                        | <input type="text"/>                     |
| IPv4 Address                | <input type="text"/>                     |
| Port                        | <input type="text" value="443"/>         |
| SNI                         | <input type="text"/>                     |
| Certificate                 | <input type="text"/>                     |
| Client Auth Certificate     | <input type="text"/>                     |
| Validation                  | <input type="text" value="Any CA"/>      |
| Cert Revoked Check          | <input type="text" value="OCSP-Server"/> |
| OCSP Server URI             | <input type="text"/>                     |

**IMPORTANT**

Before a server with an SNI extension can be added to the Exinda Appliance, the server must be added to the appliance without the SNI extension. A server without an SNI extension can be used as a fallback in event that the client is unable to process the SSL certificate with SNI. A server with the same IP address and port number can be added to the appliance by specifying a unique SNI extension for each server.

**IMPORTANT**

A server cannot be deleted if another server with the same IP address and port number, and an SNI extension has been configured on the Exinda Appliance. Servers with SNI extensions must be deleted before the server can be deleted.

Where do I find these settings?

Go to **Configuration > System > Optimization > SSL**.

To configure an SSL Acceleration server

1. In the **Add SSL Acceleration Server** area, type a name for the server or application you wish to enable for SSL Acceleration.
2. Type the **IPv4 Address** of the server running the SSL enabled application.
3. Type the **Port** number running the SSL enabled application on the server.

4. If the server has multiple SSL certificates with a Server Name Indication (**SNI**) specified, type the SNI extension in the field. The server (without an SNI) must be added before the server with the same IP and port number and an SNI can be added.

5. Select the **Certificate** to use for re-encryption of the SSL session. The certificates available here are those that are configured in the Certificate and Key page.

6. Select the **Client Auth Certificate** to authenticate sessions on the SSL server.

7. Select the type of validation to apply to the server's certificate.

- **None**— SSL Acceleration accepts and processes the connection even if the server's SSL certificate is invalid or expired.
- **Reject**— SSL Acceleration does not process the connection under any circumstances. The connection is still accelerated, but is not SSL accelerated.
- **Certificate**— SSL Acceleration accepts and processes the connection only if the server's certificate matches the specific certificate named in the Client Auth Certificate field. Otherwise, the connection is not processed.
- **Any CA**— SSL Acceleration accepts and processes the connection if the server's certificate matches any CA certificate that is loaded on to the Exinda appliance.
- **Any**— SSL Acceleration accepts and processes the connection if the server's certificate matches any certificate (CA or non-CA) that is loaded on to the Exinda appliance.

8. If **Certificate** is selected as the **Validation** type, select the certificate to validate against.

9. If **Any CA** or **Any** is selected as the **Validation** type, select the **Cert Revoked Check** type.

- **None**— No check is performed. The client auth certificate is used regardless of whether the certificate is revoked or not.
- **OCSP-AIA**— The Online Certificate Status Protocol (OCSP) Authority Information Access (AIA) check is performed. The method uses the location of the authority embedded in the certificate to check for the certificate's revocation status. Note that if the AIA location is not specified in the certificate when this option is chosen, then the certification revoke check will not happen.
- **OCSP-Server**— The Online Certificate Status Protocol (OCSP) check is performed. This method presents an **OCSP Server URI** field where you can type the location of the authority to check for the certificate's revocation status.

10. Click **Add SSL Server**.

The servers are displayed at the top of the page, where they can be edited or deleted.

To edit an SSL Accelerated server

1. Locate the server in the **SSL Acceleration Servers** list, and click **Edit**.
2. Modify the settings for the server, and click **Apply Changes**. The settings for the server are changed.

To delete an SSL Accelerated server

1. Locate the server in the **SSL Acceleration Servers** list, and click **Delete**. Servers with SNI extensions must be deleted before the server with the same IP and port number (but without an SNI) can be deleted.
2. In the confirmation dialog, click **OK**. The server is deleted.

To troubleshoot a disabled SSL Acceleration Server

If the server is disabled, check the status message in the SSL Acceleration Servers list or Remote SSL Acceleration Servers list. The list will provide feedback on why the server is disabled. Perhaps the certificate validation failed or the OCSP

validation failed.

To fix the problem, you can try relaxing the certificate validation a step at a time. For example, turn off OCSP validation and see what happens. Then turn off or broaden the certification validation, such as using ANY, or ANY-CA and see what happens. You can also use the openssl client to check the SSL handshake:

```
openssl s_client -state -msg -connection <ip:port> -ssl3 -showcerts  
openssl s_client -connect <ip:host> -tls1 -showcerts -servername <server-name>
```

## How SSL Protocol Acceleration Works

### How SSL works

SSL is the standard protocol for establishing a secure encrypted link between a remote application server and the client Web browser on the local user computer. The SSL protocol secures each session link by automatically establishing connections on-demand using standards-based protocols, encryption techniques, and certificate exchange.

SSL encryption requires a certificate on the server to authenticate the identity of a server. A certificate is an electronic confirmation that you, as the owner of a public key, are who you claim to be and that you hold the private key corresponding to the public key in the certificate.

You create this certificate by generating a certificate and sending a certificate signing request to a Certificate Authority (CA) using your public key. The CA checks with a registration authority to verify your identity and then signs and returns the certificate. You then upload the signed certificate and public key onto the server.

When a client browser visits a website hosted on the server over HTTPS, the server offers the signed certificate and public key. The client browser verifies that the certificate is valid for the site that is being visited and that it has not expired. Then it verifies the chain of trust by looking at who has signed the certificate:

- » If the certificate is a root-certificate, it will compare it to the ones shipped with the operating system or browser.
- » If it is a non-root-certificate, it will follow the chain of trust up each level until reaching a root-certificate.

Now the server has the private key and the client has the public key, a private encrypted tunnel is created to allow them to appropriately communicate by encrypting and decrypting the traffic between them. When the session is over, the connection is automatically terminated.

### Other certificate signing options

#### Self-signed certificate

You can create a self-signed certificate, where the certificate has signed itself and therefore there is no chain of trust. The browser will issue a warning, telling the user that the site certificate cannot be verified. To continue, the user will have to confirm that they trust the site. When the browser visits this site again later, the warning will not be presented again since the user has already confirmed their trust of the site. An alternate use case, is the company that created the self-signed certificate can provide the certificate to the client users and tell them to load the certificate into their browsers. This is equivalent to confirming trust when the warning is shown. Using self-signed certificates is reasonable in situations where you want encryption but you do not need the third party verification, such as an internal system where you want your internal users to have password protection, however as the clients, and the server are behind a firewall, you do not need the third party verification.

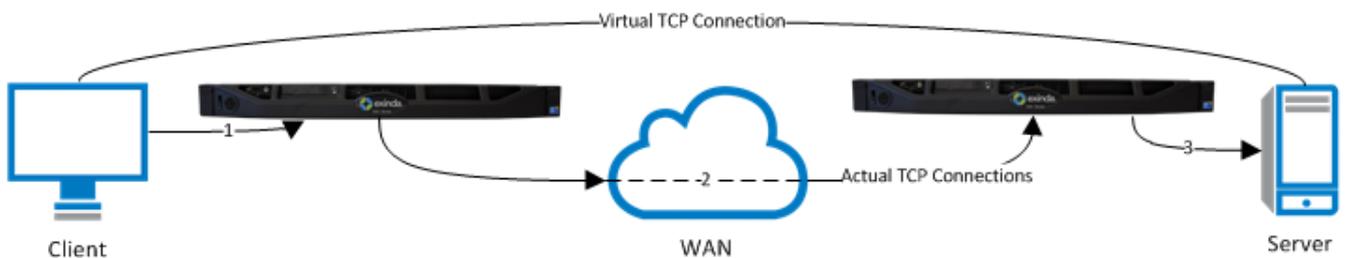
#### Self-signed Certificate Authority (CA)

You can create your own self-signed CA certificate for signing other certificates. In this case, the certificates that your self-signed CA certificate signs will have no chain of trust. Similar to self-signed certificates, using your own self-signed CA certificate to sign other certificates is reasonable in situations where you want encryption but you do not need the third-party verification. The difference is that a warning is shown to the client for each server when using self-signed certificates, whereas when using a self-signed CA certificate to sign multiple other certificates, the warning will only be

shown once for all the certificates that were signed by the CA certificate, that is, once the client trusts a certificate that is signed by the self-signed CA certificate, the client automatically trusts all other certificates signed by that self-signed CA certificate.

In the case where there are multiple virtual hosts on a single server, a Server Name Indication (SNI) is used to indicate what virtual hostname the client is attempting to connect to during the handshaking process. This allows a single server to present individual certificates for each of its multiple secure websites without requiring all of the sites to use the same certificate.

### How Exinda accelerates the SSL protocol



For SSL acceleration, a server-side Exinda appliance and a client-side appliance is put in line for this SSL traffic. The traffic between these appliances are accelerated. The benefits that can be gained by generic application acceleration on encrypted data are limited. For example, the Exinda WAN Memory technology achieves higher reduction on clear text rather than encrypted data. However, the SSL acceleration feature is designed to overcome these limitations by transparently decrypting accelerated traffic, performing the relevant application acceleration techniques such as TCP Acceleration and WAN Memory, then re-encrypting the traffic again. This means Exinda can apply all application acceleration technologies to the traffic as if it were clear text, while still maintaining SSL connections.

The server-side appliance acts on behalf of the client in the communication between the appliance and the server and the client-side appliance acts on behalf of the server for communication between the client and the appliance. To decrypt and re-encrypt the traffic, the Exinda appliances must have access to the appropriate certificate and public key for each server that clients communicate with over SSL. Furthermore, the Exinda appliances must be configured to know which servers can receive traffic that is SSL accelerated. These servers are defined by IP address and port, certificate, and other details. Only traffic to servers that are explicitly configured in this way is SSL accelerated. If the server is hosting multiple virtual hosts, when defining the server, you can define an acceleration server for each of the virtual hosts by specifying the SNI virtual host so that the virtual host name is presented during the handshake process with the appliance.

#### NOTE

If you upload the appropriate certificates and configure SSL Acceleration Servers on the server-side appliance, the appliance will use the Exinda acceleration community feature to push these certificates and server configurations to the other appliances in the community. Configurations that have been pushed to the remote appliances will appear in the **Remote SSL Acceleration Servers** list on the **Optimization > SSL** page.

By default, the Exinda appliances are pre-loaded with several root CA certificates. The site-specific certificates will be loaded onto the appliances by the user (or distributed using the community feature). When the client attempts to access the website, during the handshake, the appliance sends to the client all of the certificates in the chain of trust.

#### NOTE

If you are concerned about any decrypted data on the Exinda appliance, then you can choose to use storage disk encryption.

To configure SSL acceleration

1. [Configure SSL certificates and private keys](#) (or configure SSL CA certificates and private keys) to use for SSL acceleration.
2. Configure the servers to use for SSL acceleration. For more information, refer to [Configure SSL Acceleration Servers](#) (page 658).
3. Create Optimizer policies that allow SSL traffic to be accelerated.
4. If you're concerned about decrypted data on the appliance, then enable storage disk encryption.

How to tell if SSL Acceleration is working?

1. Ensure that the SSL Acceleration Server is ok by seeing its status on the SSL server configuration page. There is a green checkmark next to each SSL Acceleration Server that has a good state.
2. On the real-time conversations page, turn on Show Policies and ensure the SSL traffic that you are interested is accelerated. If the traffic shows in a gold band, then it is processed by an accelerated policy. If the traffic has a lock icon, then it is being SSL accelerated.

If your traffic is not accelerating and you need to troubleshoot, try the following steps:

1. Check that the policies seem to be configured properly and that they are in the proper order. You want to ensure that another policy earlier in the tree is not capturing your desired traffic.
2. Check the SSL Acceleration Server details. Ensure you are using correct spelling, etc. More troubleshooting help for disabled SSL Acceleration Servers is offered in the [Configure SSL Acceleration Servers](#) section.
3. Check that the Exinda community feature has distributed the certificates and SSL Acceleration Server configuration properly to your appliance. They will appear as **remote**.

### [Create policies to accelerate SSL traffic](#)

The default policies that are created as a result of running the policy configuration wizard captures SSL traffic in a QoS only policy, meaning no attempt is made to accelerate any SSL traffic by default. To accelerate SSL traffic, you need to create an acceleration policy for the SSL application server you want to accelerate. Any SSL traffic that matches an acceleration policy is passed to SSL Acceleration. If a valid certificate and key are configured for that SSL traffic, then SSL acceleration occurs.

To create a policy for accelerating an SSL application

1. Go to the Optimizer and create a new policy in the appropriate circuit & virtual circuit by clicking **Create new Policy...**
2. Specify the **Action** as **Optimize**.
3. Check the **Acceleration** checkbox.
4. Add a filter for the specified host and the specified SSL application.
5. Add the policy on all Exinda appliances.
6. Once the desired policies are in place on all Exinda appliances, restart the Optimizer.

## EXAMPLE

To Accelerate an SSL application, this policy needs to be placed above any other policy that generically captures SSL traffic in the policy tree.

**Edit Policy**

|                                                     |                                                                                             |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------|
| Policy Name: <input type="text" value="SSL Accel"/> | <input checked="" type="checkbox"/> Guaranteed Bandwidth: <input type="text" value="5"/> %  |
| Schedule: <input type="text" value="ALWAYS"/>       | Burst (Max) Bandwidth: <input type="text" value="100"/> %                                   |
| Action: <input type="text" value="Optimize"/>       | Burst Priority: <input type="text" value="4"/>                                              |
| Policy Enabled: <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> Acceleration: <input type="text" value="Acceleration"/> |
|                                                     | WM Reduction Type: <input type="text" value="Disk"/>                                        |
|                                                     | <input type="checkbox"/> ToS/DSCP Mark: <input type="text"/>                                |

| VLAN                             | Host                             | Direction                                | Host                                    | ToS/DSCP                         | Application                        |
|----------------------------------|----------------------------------|------------------------------------------|-----------------------------------------|----------------------------------|------------------------------------|
| <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="&lt; - &gt;"/> | <input type="text" value="SSL Server"/> | <input type="text" value="ALL"/> | <input type="text" value="HTTPS"/> |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="&lt; - &gt;"/> | <input type="text"/>                    | <input type="text"/>             | <input type="text"/>               |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="&lt; - &gt;"/> | <input type="text"/>                    | <input type="text"/>             | <input type="text"/>               |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="&lt; - &gt;"/> | <input type="text"/>                    | <input type="text"/>             | <input type="text"/>               |
| <input type="text"/>             | <input type="text"/>             | <input type="text" value="&lt; - &gt;"/> | <input type="text"/>                    | <input type="text"/>             | <input type="text"/>               |

## Encrypt Disk Storage

SSL acceleration requires the SSL traffic to be decrypted and cached so that various acceleration techniques can be applied to the data. If you are concerned about this, then you can encrypt storage for WAN memory.

If the storage for WAN memory is encrypted, a green check-mark is shown in the Encrypted column.

### Disk Storage Map.



| Storage Configuration           |           |              |  |                                       |          |                                     |                                       |                                       |                                           |
|---------------------------------|-----------|--------------|--|---------------------------------------|----------|-------------------------------------|---------------------------------------|---------------------------------------|-------------------------------------------|
| Service                         | Status    | Free         |  | Size                                  | Minimum  | Encrypted                           | Operation                             |                                       |                                           |
| cifs                            | available | 1898.23M 97% |  | <input type="text" value="1964.00M"/> | 1024.00M | <input checked="" type="checkbox"/> | <input type="button" value="Resize"/> | <input type="button" value="Format"/> | <input type="button" value="No Encrypt"/> |
| edge-cache                      | available | 569.83M 93%  |  | <input type="text" value="612.00M"/>  | 1024.00M | <input checked="" type="checkbox"/> | <input type="button" value="Resize"/> | <input type="button" value="Format"/> | <input type="button" value="No Encrypt"/> |
| log                             | available | 974.64M 95%  |  | <input type="text" value="1024.00M"/> | 512.00M  | <input checked="" type="checkbox"/> | <input type="button" value="Resize"/> | <input type="button" value="Format"/> | <input type="button" value="Encrypt"/>    |
| monitor                         | available | 9401.04M 92% |  | <input type="text" value="10.00G"/>   | 10.00G   | <input checked="" type="checkbox"/> | <input type="button" value="Resize"/> | <input type="button" value="Format"/> | <input type="button" value="Encrypt"/>    |
| users                           | available | 974.62M 95%  |  | <input type="text" value="1024.00M"/> | 512.00M  | <input checked="" type="checkbox"/> | <input type="button" value="Resize"/> | <input type="button" value="Format"/> | <input type="button" value="No Encrypt"/> |
| wan-memory                      | available | 2149.26M 95% |  | <input type="text" value="2252.00M"/> | 5120.00M | <input checked="" type="checkbox"/> | <input type="button" value="Resize"/> | <input type="button" value="Format"/> | <input type="button" value="No Encrypt"/> |
| unallocated storage             |           |              |  | 0.00                                  |          |                                     |                                       |                                       |                                           |
| <b>Total Available Storage:</b> |           |              |  | <b>18.71G</b>                         |          |                                     |                                       |                                       |                                           |

Screenshot 286: Disk Storage settings

To enable encryption for wan-memory:

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Go to **Configuration > System > Setup > Storage**.
5. To encrypt WAN memory storage, click **Encrypt** button for wan-memory.

## Ciphers supported in SSL acceleration

SSL Acceleration supports the following ciphers (encryption/decryption algorithms).

| Protocol Key | Length   | Cipher Name  |
|--------------|----------|--------------|
| SSLv3        | 256 bits | AES256-SHA   |
| SSLv3        | 128 bits | AES128-SHA   |
| SSLv3        | 168 bits | DES-CBC3-SHA |
| SSLv3        | 128 bits | RC4-SHA      |
| SSLv3        | 128 bits | RC4-MD5      |
| TLSv1        | 256 bits | AES256-SHA   |
| TLSv1        | 128 bits | AES128-SHA   |
| TLSv1        | 168 bits | DES-CBC3-SHA |
| TLSv1        | 128 bits | RC4-SHA      |
| TLSv1        | 128 bits | RC4-MD5      |

- » If the client does not support any of these ciphers, the SSL connection is rejected.
- » If the server does not support any of these ciphers, it is automatically removed.

## Host multiple secure websites on Windows Server 2012

On a corporate network, it may be necessary to have multiple secure websites being served from a single Windows server, on a single IP address. Previously, attempting to host multiple secure sites on a single IP address would cause certificate requests to be perceived as man-in-the-middle attacks, and the connections would be refused. IIS 8.0, available only on Windows Server 2012, introduces the Server Name Indication (SNI) extension which allows a hostname or domain name to be included in SSL certificate requests. With SNI, multiple secure websites can be served from a single IP address as the certificates requests for the sites include the SNI extension, allowing the correct certificate to be presented to the client browser.

To host multiple secure websites on Windows Server 2012, configure the websites to include the SNI extension in the connection requests.

1. [Install IIS 8.0 on Windows Server 2012](#)
2. [Add sites to the web server](#) .
3. Ensure the certificates required for the sites are available on the server. Depending on how your organization manages SSL certificates, this may involve generating a self-signed certificate or importing a certificate from a Certificate Authority. For instructions managing the certificates on the Windows Server, refer to the Microsoft help.
4. (Optional) If the site requires Server Name Indication (SNI), create a self-signed certificate that identifies the ID of the site. See, [Create self-signed certificates for each site requiring Server Name Indication](#).
5. [Identify the certificate to be used by each website](#)
6. [Export SSL certificates from Windows Server 2012](#)
7. [Managing Certificates and CA Certificates](#)
8. [Configure SSL Acceleration Servers](#)

## Install IIS 8.0 on Windows Server 2012

IIS 8.0 must be installed on the Windows server before certificates with Server Name Indicators (SNI) can be configured.

1. Open the **Server Manager**.
2. Select **Manage > Add Roles and Features**.
3. Select **Role-based or Feature-based Installation**. Click **Next**.
4. Select the appropriate server and click **Next**.
5. From the list of Server Roles, select **Web Server (IIS)**.
6. In the Add Roles and Features Wizard dialog, click **Add Features**.
7. Click **Next**.
8. Do not select any additional features, and click **Next**.
9. On the Web Server Role (IIS) information screen, click **Next**.
10. Accept the default role services, and click **Next**.
11. Review the selections, and click **Install**. When the IIS installation completes, the wizard reflects the installation status.
12. To exit the wizard click **Close**.

### Add sites to the web server

Add sites that require SSL certificates with Server Name Indicators (SNI) to the IIS Manager to manage what certificates are used by each site.

1. In the **Server Manager**, and click **IIS**.
2. Right-click the server name, and select **Internet Information Services (IIS) Manager**.
3. Double-click the server name.
4. Right-click **Sites** and select **Add Website**.
5. Add the parameters for the website.
6. In the Binding area, ensure you type the host name of the server.
7. Click **OK**.
8. Repeat these steps for each secure website that will be available on the server.

### Create self-signed certificates for each site requiring Server Name Indication

The SelfSSL tool is installed with IIS, and allows you to create self-signed certificates that include the ID of the site within the certificate.

1. In the **Internet Information Services (IIS) Manager**, click Sites and make note of the ID of each website using the self-signed certificate that has Requires Server Name Indication selected.
2. Open a command prompt and navigate to **C:\Program File (x86)\IIS Resources\SelfSSL**.
3. At the prompt type the parameters for the certificate, ensuring you specify the site ID for the site requiring Server Name Indication. For example:

## NOTE

In the command, /V is the number of days the certificate is valid, /S is the ID of the site. Use the values that correspond to your site in the command.

```
selfssl.exe /N:CN=TEST.SITE.3 /K:1024 /V:<days-valid> /S:<site-ID> /P:443
```

The certificate is created.

4. When prompted to replace the SSL settings for the site, type Y.

5. Modify the site to use the new certificate in the bindings. See For more information, refer to [Identify the certificate to be used by each website](#) (page 668)..

### Identify the certificate to be used by each website

Specify the certificate that the secure website uses when receiving requests.

1. In the **Internet Information Services (IIS) Manager**, locate the site created in [Add sites to the web server](#).
2. In the Actions list, select **Bindings**.
3. In the Type list, select **https** and click **Edit**.
4. Type a host name.
5. Select the appropriate SSL certificate.
6. If this site uses the same IP address as another secure site, select **Require Server Name Indication**.
7. To add the binding, click **OK**.
8. Click **Close**. The binding is added for the site.
9. Repeat this task for each site configured on the server.

### Export SSL certificates from Windows Server 2012

Export the certificates from the Windows server so they can be imported onto the Exinda Appliance.

1. In the **Server Manager**, and click **IIS**.
2. Right-click the server name, and select **Internet Information Services (IIS) Manager**.
3. Double-click **Server Certificates**.
4. Right-click the certificate, and select **Export**.
5. Specify the location where the exported certificate should be saved, and type a name for the certificate. Click **Open**.
6. Type and confirm the password required to use the certificate.
7. Click **OK**. The certificate is exported to the specified location.

### Host multiple secure websites on Apache

On a corporate network, it may be necessary to have multiple secure websites being served from a single Apache server, on a single IP address. Previously, attempting to host multiple secure sites on a single IP address would cause certificate requests to be perceived as man-in-the-middle attacks, and the connections would be refused.

Configure the websites served up from Apache to include the SNI extension in the connection requests.

## NOTE

Only Apache 2.2.12 and later and OpenSSL 0.9.8j and later have support for Server Name Indication (SNI). SNI is not supported on Internet Explorer running on Windows XP.

1. Create all the secured sites on the Apache server.
2. Copy the certificate files for the secure sites onto the Apache server. Put the certificate files in the same location as the other certificates on the server. The certificates should be readable by the web server process only.
3. [Enable SSL on Apache](#)
4. [Specify the ports referenced by the virtual hosts](#)
5. [Add a <VirtualHost> block for each secure site on the server](#)
6. [Verify the secure server configuration](#)
7. [Managing Certificates and CA Certificates](#)
8. [Configure SSL Acceleration Servers](#)

### Enable SSL on Apache

To use SSL on Apache, the `mod_ssl` module must be enabled.

To enable the `mod_ssl` module, type the following command: `sudo a2enmod ssl`

### Specify the ports referenced by the virtual hosts

A SSL web server must run on a different port than an unencrypted web server. The standard port for HTTPS traffic is 443, but any port number can be used. Apache will not accept incoming connections to any ports if they are not specified with a `Listen <port_number>` directive in the active configuration set.

1. Navigate to `/etc/apache2/conf.d` and open the `ports.conf` file in an editor.
2. Locate the `<IfModule mod_ssl.c>` block.
3. Ensure `Listen 443` is included in the block.
4. Add `NameVirtualHost *:443` to the block.
5. Save the configuration file.

### Add a <VirtualHost> block for each secure site on the server

For each domain name or domain subset we want to support SSL for, a `VirtualHost` block must be declared. This block identifies the domain name to support connections for, and what Certificate or Key files to use for it.

1. Navigate to `/etc/apache2/sites-enabled` and open the folder for the secure site.
2. Open the `<site_name>.conf` file in an editor.
3. Add the `<virtualhost>` block for the secure server.

The block will look similar to this:

```
<VirtualHost *:443>
ServerName "secure2.example.com"
ServerAdmin webmaster@example.com
DocumentRoot /home/demo/public_html/secure1.example.com/public
ErrorLog /home/demo/public_html/secure2.example.com/log/error.log
```

```

LogLevel warn
CustomLog /home/demo/public_html/secure2.example.com/log/access.log combined
<Directory /home/demo/public_html/secure2.example.com/public>
Options Indexes FollowSymLinks MultiViews
AllowOverride None
Order allow,deny
allow from all
</Directory>
SSLEngine On
SSLCertificateFile /var/www/certs/secure2.pem
SSLCertificateKeyFile /var/www/keys/secure2.key
</VirtualHost>

```

Update the sample block to reflect the file locations on your Apache server, and ensure each block references the correct secure site.

4. Save the configuration file.

#### Verify the secure server configuration

It is always best to check your Apache config files for any errors before restarting, because Apache will not start again if your config files have syntax errors.

1. Check the Apache config files for errors, run the following command:

```
sudo apachectl configtest
```

2. Display the secure sites in a browser that supports SNI, and verify no errors are displayed.

3. Restart the Apache server to commit the configuration changes.

```
sudo apachectl stop
```

```
sudo apachectl start
```

4. After the server has restarted, run the command `sudo netstat -tnlp | grep 443` and verify that the server is listening on port 443.

### Configure file acceleration

SMB Acceleration is the file transfer specific component of the Exinda Application Acceleration Technology. To deal with inefficiencies in the SMB protocol, the Exinda Appliance has several optimizations to improve the performance of applications using this protocol, including reading ahead of the data stream, writing behind the data stream, and caching meta data on files and folders.

SMB acceleration makes the following scenarios more efficient:

- » File Download (Read) – The SMB client is reading a file from an SMB server. The server-side Exinda proactively requests future read events and passes the read information to the client-side Exinda so that it is available locally and immediately to the SMB client.
- » File Upload (Write) – Similar to the read scenario, the Exinda appliance proactively transfers write data to the other Exinda. The client-side Exinda responds locally to write requests from the SMB client and passes the data to the server-side Exinda at WAN link speed to complete the write operation.

- » Remote Access of Microsoft Office Files – Microsoft office files (Word, PowerPoint, Excel, etc.) which reside on a remote SMB server are often opened from a SMB client. The Exinda SMB Acceleration addresses slow downloads by pre-fetching the file data and populating it on the client side Exinda. Consequently, all SMB client requests for the file data are served from the client side Exinda at LAN speeds.
- » Directory Browsing – When browsing a remote file system using Windows Explorer, the SMB protocol transfers various bits of information about the files you are browsing. This metadata is transferred in special SMB instructions called transactions. The Exinda appliance also caches these transactions such that they can be served locally, from the client-side Exinda appliance. This significantly improves the performance of directory browsing using the SMB protocol.

#### Related Topics

- » [Configuring SMB File Acceleration](#)
- » [How SMB/CIFS \(File Transfer\) Acceleration Works](#)
- » [Diagnostics Files](#)
- » [Manage optimization services](#)

#### Configuring SMB File Acceleration

Use the following instructions to configure the Exinda Appliance to use the SMB optimization features. The SMB file acceleration configuration has several options for managing transfers under a variety of situations.

1. Go to **Configuration > System > Optimization**.
2. On the right, select the **SMB** tab.
3. Select, as needed, the **Enabled** check boxes for SMB1 and SMB2.

#### NOTE

Under the SMB1 section there are several options that you can use. The "Read Ahead" and "Read Behind" options, however, relate only to SMB1 traffic. For the "Data to Prefetch" option, this again relates only to SMB1.

4. In the **Acceleration Options** section complete the following, as needed:
  - a. To prefetch data from the SMB server in anticipation of subsequent client requests, select **Read Ahead**.
  - b. To update the SMB cache, aggregating requests to the SMB server, select **Write Behind**.
  - c. To enable caching of SMB file attributes such as file access times and size, select **Meta-Data Caching**.
  - d. In the **Data to Prefetch** field, specify the amount of SMB data to prefetch when performing read-ahead or write-behind.

#### NOTE

The default value is 1024 kb. If you experience network latency with file transfers, for example if latency is greater than 500 milliseconds, increase the amount of data being prefetched for the SMB1 cache.

5. To accelerate digitally signed connections, in the **SMB Signing** area select **Enabled**.

**SMB Signing**  
 Enabled

**Domain**  
 Domain:

| Credentials |                      |                      |                          |        |
|-------------|----------------------|----------------------|--------------------------|--------|
| Domain      | Username             | Password             | Enabled                  | Modify |
| (default)   | <input type="text"/> | <input type="text"/> | <input type="checkbox"/> |        |

6. In the **Domain** area, type the NetBIOS name, *not* the fully qualified domain name, of a domain or computer where shared files are located, and click **Add**. The domain appears in the SMB signing Credentials section below. If necessary, add any other domains as well.

7. In the **Credentials** area, type the **Username** and **Password** to use when generating the signing key.

**IMPORTANT**

The account used for signing the key must be able to authenticate against the specified server, but should be a highly restricted account that does not have permissions to access the files being requested by the client computers, or administrator access to the domain.

8. Select whether the authentication credentials are enabled or disabled.

**NOTE**

The Exinda Appliance will use the recorded credentials for the requested domain. If the appliance is unable to connect to that domain because the server is unavailable or because of incorrect credentials, the status of the signed connection is reported as `Bypassed` or `Unhandled`.

If the request is for a domain that is not registered on the Exinda Appliance, the credentials for the `(default)` domain are used. If the credentials are incorrect for the requested domain, the status of the signed connection is reported as `Bypassed` or `Unhandled`.

9. In the **Compatibility Options** area, if on your network there are Exinda appliances that are running a firmware version prior to v6.3, select **Support versions (pre v6.3.0) Enabled**.

10. Click **Apply Changes**.

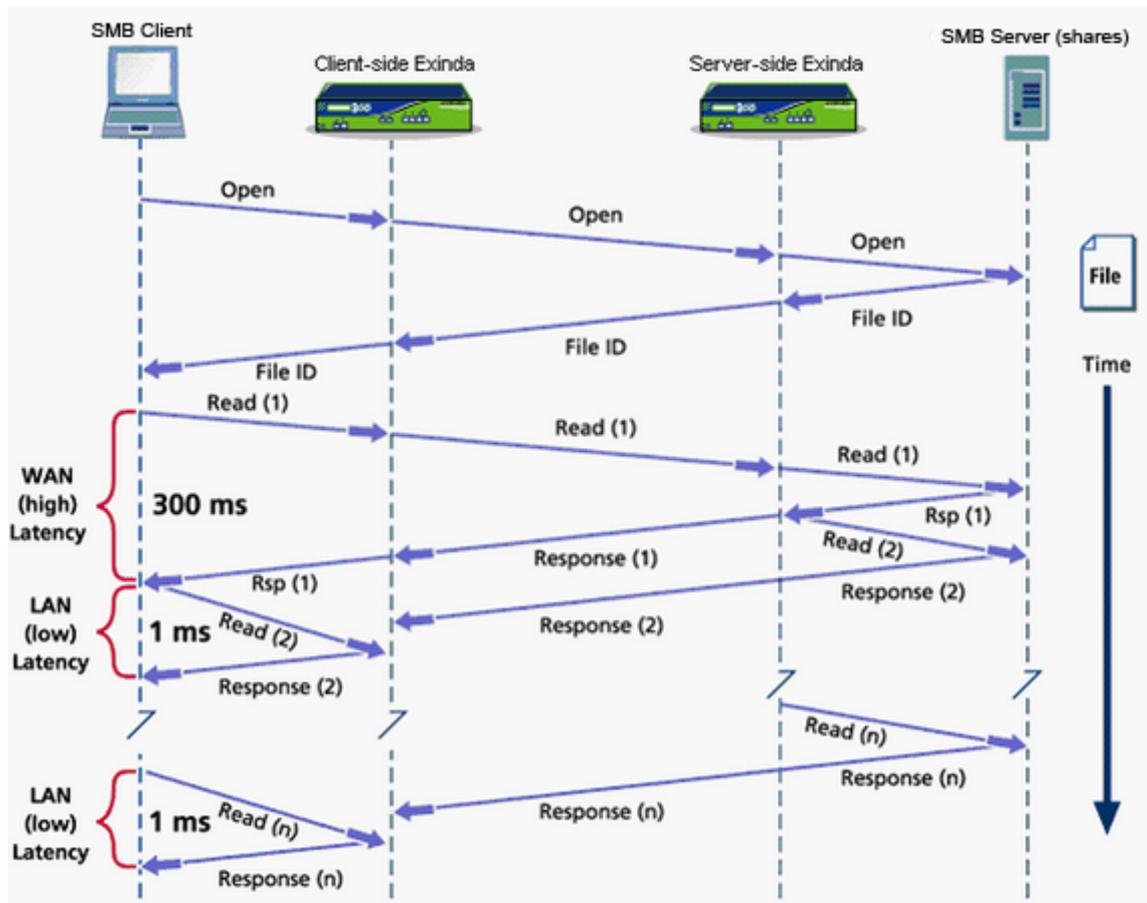
**How SMB/CIFS (File Transfer) Acceleration Works**

SMB1 and SMB2 are remote file access protocols that form the basis for Windows file sharing. Each time you browse or access files on a Windows server using Windows Explorer, Windows uses the SMB protocols to transport information (files or directory information) back and forth between your computer and the server.

In addition to file sharing, SMB is also used as a transport protocol for various higher level Microsoft communications protocols, as well as for network printing, resource location services, remote management/administration, network authentication (secure establishment services), and RPC (Remote Procedure Calls). SMB operates very poorly over a high latency WAN link because by design SMB sends a large number of back and forth transactions to complete a request.

The primary goal of Exinda SMB acceleration is to reduce the overall accumulated latency introduced by the "chattiness" of the SMB protocols. Each Exinda appliance can act on behalf of an SMB client and server to make the interaction

between the two computers more efficient. Exinda maintains a state machine and database of SMB behaviors that it relies upon to optimize future SMB related transactions. When Exinda determines that a certain SMB transaction is likely to occur, it pre-fetches data and temporarily stores it in the appliance memory for future reference. Once the pre-fetched data is referenced, the data is deleted from the memory. See the following figure.



SMB acceleration works seamlessly for TCP Acceleration, WAN Memory, and Compression, and benefits from the ability of the WAN memory to reduce data traversing the WAN, just as with other applications such as FTP, HTTP, or email.

To deal with the inefficiencies of the SMB protocol Exinda has developed several optimizations to improve the performance of applications using this protocol. Each version of SMB handles file transfer optimizations in different ways, and may include reading ahead of the data stream, writing behind the data stream, and caching meta data about files and folders. In addition to this, the Exinda Appliance ensures that data is efficiently compressed and deduplicated.

### SMB1

With SMB1 there are several opportunities to provide optimizations: object caching, read ahead and write behind of data, and meta-data caching.

#### » Object Cache

This refers to the saving of files to an internal file storage area on the appliance. When a client reads a file, it is cached on both the client side and server side appliances. This significantly improves response time for successive reads of the same file as it occurs at LAN speed instead of WAN speed. When a client writes a file, the object cache is updated which allows successive reads of the file to be served as efficiently as possible without having to use the WAN.

#### » Read Ahead

Reading ahead of the data stream is an optimization by which the appliance pre-fetches the contents of a file ahead of the client that is attempting to read it. When the Exinda Appliance detects a client attempting to perform a sequential bulk read of a file, the appliance fabricates read requests to the server on behalf of the client. The end result is that the

appliance is effectively sending the reads to the server and pre-populating both the client side and server side cache. Since SMB1 clients perform reads serially, this dramatically improves cold pass read performance and helps to populate the object cache quickly.

» Write Behind

Writing behind the data stream is an optimization by which the Exinda Appliance immediately responds to the client when it is trying to write a file. When the appliance detects a client attempting to perform a bulk write to a file, it immediately responds to the client from the client side appliance. The end result is that the Exinda Appliance is effectively sending the write requests to the server so the conversation between the client and client side appliance is occurring at LAN speed. Since SMB1 clients perform writes serially, the immediate response by the appliance allows the client write requests to fill the connection, making it appear to be asynchronous and significantly improving write performance.

» Meta-data Caching

Meta-data caching is an optimization by which the Exinda Appliance caches the properties related to files and folders on both the client side and server side appliances. When a client queries the properties of a file or folder, it is served from cache which eliminates the need to go across the WAN. This occurs quite frequently when browsing a file share location that has a larger number of file and folder entries. Similar to the object cache, change notifications are registered to ensure that the meta-data cache does not serve stale information.

## **SMB2**

With the addition of SMB2, most of the optimizations that were implemented for SMB1 no longer apply. Below is a rationale for each of these and why they are no longer needed.

» Read Ahead and Write Behind

In SMB2, read ahead and write behind requests are built in to the client, effectively stacking the requests one on top of the other in an asynchronous manner without any gaps between them. As a result, there is no accumulation of latency and therefore no need for the appliance to attempt to perform any sort of read prefetching or immediate write response.

» Meta-data Caching

In SMB2, meta-data caching is performed by the client. This eliminates the need for the appliance to do any caching in the middle as the client very quickly caches its own copy of the file and folder meta-data locally and uses that for the duration of the session.

## **Compression and Deduplication**

Aside from the protocol specific optimizations that are provided by the appliance, the Exinda SMB acceleration framework also provides some significant downstream optimization benefits, primarily in the areas of compression and deduplication. The SMB acceleration framework is reconstructing the SMB messages in their entirety before processing them. This means that for large data centric operations like reading and writing a file, the appliance is actually operating on large blocks of data as opposed to individual packets of fragmented data. In doing so, Exinda passes off these large blocks of data to our WAN memory framework. This allows the WAN memory framework to heavily optimize for compression and deduplication.

## **Accelerating Exchange and Outlook Traffic (MAPI)**

---

---

### **Accelerate Exchange and Microsoft Outlook traffic**

Exinda x800 appliances have built-in support for accelerating MAPI traffic. To gain the most benefits from Exinda's acceleration, compression and WAN Memory technologies, it is recommended that any native encryption be disabled in

Exchange and Outlook.

- » For more information, refer to [Enable MAPI Acceleration on the Exinda Appliances](#) (page 677). Enable MAPI Acceleration on the Exinda Appliances For more information, refer to [Enable MAPI Acceleration on the Exinda Appliances](#) (page 677).
- » To disable encryption of all MAPI traffic, you must turn off encryption on both Microsoft Outlook and the Exchange server. See [Turn off MAPI encryption in Microsoft Outlook](#) and [Disable encryption on the Exchange server](#).
- » Verify MAPI traffic is being accelerated.
  - For more information, refer to [View MAPI Acceleration Results](#) (page 679). View MAPI Acceleration Results For more information, refer to [View MAPI Acceleration Results](#) (page 679).
  - For more information, refer to [Monitoring conversations in real time](#) (page 246). Conversations in Real Time For more information, refer to [Monitoring conversations in real time](#) (page 246)..

#### NOTE

If you encounter any issues, see [Troubleshoot problems with MAPI acceleration](#).

### Turn off MAPI encryption in Microsoft Outlook

Exinda recommends that encryption of contents and attachment for output messages should be disabled to maximize reduction, as each user will encrypt files with a different key. MAPI Encryption is a client side configuration parameter in Outlook. Therefore, to disable MAPI encryption you need to make the change on each Outlook client.

You must disable encryption on all Microsoft Outlook clients as well as the Exchange server. See For more information, refer to [Disable encryption on the Exchange server](#) (page 677)..

### Turn off MAPI encryption in Outlook 2003

1. Open Microsoft Outlook.
2. Configure the Trust Center.
  - a. On the **Tools** menu select **Options**.
  - b. Switch to the **Security** tab.
  - c. Ensure the **Encrypt contents and attachments for outgoing messages** checkbox is not selected.
  - d. To close the dialog and save the settings, click **OK**.
  - e. To close the Options dialog, click **OK**.
3. Configure the Account Settings.
  - a. On the **Tools** menu select **Email Accounts**.
  - b. On the **E-mail** tab, select the email account and click **Change**.
  - c. On the Server Settings page, click **More Settings**.
  - d. Switch to the **Security** tab.
  - e. Ensure the **Encrypt data between Microsoft Outlook and Microsoft Exchange** checkbox is not selected.
  - f. To close the dialog and save the settings, click **OK**.
  - g. To close the Server Settings dialog, click **Next** and **Finish**.

## Turn off MAPI encryption in Outlook 2007

1. Open Microsoft Outlook.
2. Configure the Trust Center.
  - a. On the **Tools** menu select **Trust Center**.
  - b. Click **Trust Center Settings**. The Trust Center dialog opens.
  - c. Click **E-mail Security**.
  - d. Ensure the **Encrypt contents and attachments for outgoing messages** checkbox is not selected.
  - e. To close the Trust Center dialog and save the settings, click **OK**.
  - f. To close the Options dialog, click **OK**.
3. Configure the Account Settings.
  - a. Outlook 2007: On the **Tools** menu select **Account Settings**.
  - b. On the **E-mail** tab, select the email account and click **Change**.
  - c. On the Server Settings page, click **More Settings**.
  - d. Switch to the **Security** tab.
  - e. Ensure the **Encrypt data between Microsoft Outlook and Microsoft Exchange** checkbox is not selected.
  - f. To close the dialog and save the settings, click **OK**.
  - g. To close the Server Settings dialog, click **Next** and **Finish**.

## Turn off MAPI encryption in Outlook 2010 and 2013

1. Open Microsoft Outlook.
2. Configure the Trust Center.
  - a. On the **File** menu select **Options**.
  - b. Click **Trust Center > Trust Center Settings**. The Trust Center dialog opens.
  - c. Click **E-mail Security**.
  - d. Ensure the **Encrypt contents and attachments for outgoing messages** checkbox is not selected.
  - e. To close the Trust Center dialog and save the settings, click **OK**.
  - f. To close the Options dialog, click **OK**.
3. Configure the Account Settings.
  - a. On the **File** menu select **Info > Account Settings**.
  - b. On the **E-mail** tab, select the email account and click **Change**.
  - c. On the Server Settings page, click **More Settings**.
  - d. Switch to the **Security** tab.
  - e. Ensure the **Encrypt data between Microsoft Outlook and Microsoft Exchange** checkbox is not selected.

f. To close the dialog and save the settings, click **OK**.

g. To close the Server Settings dialog, click **Next** and **Finish**.

#### NOTE

These parameters are configurable through a global change so that each client does not need to be individually changed.

### Enable MAPI Acceleration on the Exinda Appliances

Turn Acceleration ON for the policy that captures MAPI traffic. By default, MAPI falls into the Mail policy.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > Optimizer**.
6. At the end of the Virtual Circuit policy list, type a priority number in the **Order** field, and select **Mail - Guarantee Low 5%-100% - Accelerate**.
7. Click **Add to WAN outbound**. The policy is added to the active policies for the virtual circuit.

| Virtual Circuit 15 - WAN outbound (1024000 kbps to ALL) |                      |                                                                                                                 | --Actions--           |
|---------------------------------------------------------|----------------------|-----------------------------------------------------------------------------------------------------------------|-----------------------|
| <input checked="" type="checkbox"/>                     | 10                   | P2P - Choke 1%-3% (Optimize, 1%-3%, Priority 10)                                                                | --Actions--           |
| <input checked="" type="checkbox"/>                     | 20                   | Recreational - Limit Low 2%-10% (Optimize, 2%-10%, Priority 10)                                                 | --Actions--           |
| <input checked="" type="checkbox"/>                     | 30                   | Software Updates - Guarantee Low 5%-100% - Accelerate (Optimize, 5%-100%, Priority 6, Application Acceleration) | --Actions--           |
| <input checked="" type="checkbox"/>                     | 40                   | Voice - Guarantee Critical 15%-100% (Optimize, 15%-100%, Priority 1)                                            | --Actions--           |
| <input checked="" type="checkbox"/>                     | 50                   | Interactive and Secure - Guarantee High 10%-100% (Optimize, 10%-100%, Priority 3)                               | --Actions--           |
| <input checked="" type="checkbox"/>                     | 60                   | Thin Client - Guarantee High 10%-100% (Optimize, 10%-100%, Priority 3)                                          | --Actions--           |
| <input checked="" type="checkbox"/>                     | 70                   | Files - Guarantee Med 8%-100% - Accelerate (Optimize, 8%-100%, Priority 4, Application Acceleration)            | --Actions--           |
| <input checked="" type="checkbox"/>                     | 80                   | Web - Guarantee Med 8%-100% - Accelerate (Optimize, 8%-100%, Priority 4, Application Acceleration)              | --Actions--           |
| <input checked="" type="checkbox"/>                     | 90                   | Mail - Guarantee Low 5%-100% - Accelerate (Optimize, 5%-100%, Priority 6, Application Acceleration)             | --Actions--           |
| <input checked="" type="checkbox"/>                     | 100                  | Database - Guarantee Med 8%-100% - Accelerate (Optimize, 8%-100%, Priority 4, Application Acceleration)         | --Actions--           |
| <input checked="" type="checkbox"/>                     | 200                  | ALL - Guarantee Low 5%-100% (Optimize, 5%-100%, Priority 7)                                                     | --Actions--           |
| <b>Order:</b>                                           | <input type="text"/> | <b>Policy:</b> ALL - Accelerate                                                                                 | Add To 'WAN outbound' |

8. To restart the Optimizer, in the system toolbar click **Restart**.

#### NOTE

This can be done by following the Optimizer Wizard in the Basic User Interface. Select 'Yes' when asked if you would like to accelerate.

### Disable encryption on the Exchange server

For 2007, 2010, and 2013 Exchange servers, Exinda recommends that encryption of the MAPI protocol should be disabled to maximize reduction, as each user will encrypt files with a different key. You must disable encryption on all Microsoft Outlook clients as well as the Exchange server to maximize the benefit. See For more information, refer to [Turn off MAPI encryption in Microsoft Outlook](#) (page 675).

## NOTE

Encryption cannot be disabled on Exchange 2003 servers, but it will not enforce a policy requiring encrypted communications between Exchange and Outlook.

Turn off encryption on Exchange 2007 servers

1. Open the **Exchange Management Shell**.
2. At the command prompt, type the following command: `Set-MailboxServer <ExchangeServerName> -MAPIEncryptionRequired:$false`
3. To verify the change to the encryption status, type the following command:  
`Get-MailboxServer <ExchangeServerName>`

Turn off encryption on Exchange 2010 and 2013 servers

1. Open the **Exchange Management Shell**.
2. At the command prompt, type the following command: `Set-RpcClientAccess -Server <ExchangeServerName> -EncryptionRequired $false`
3. To verify the change to the encryption status, type the following command: `Get-RpcClientAccess -Server <ExchangeServerName>`

## Verify MAPI traffic is being accelerated

Confirm that the MAPI traffic from Microsoft Outlook and Exchange are being accelerated.

### [View MAPI Acceleration Results](#)

### [Conversations in Real Time](#)

## Compare compression of MAPI traffic

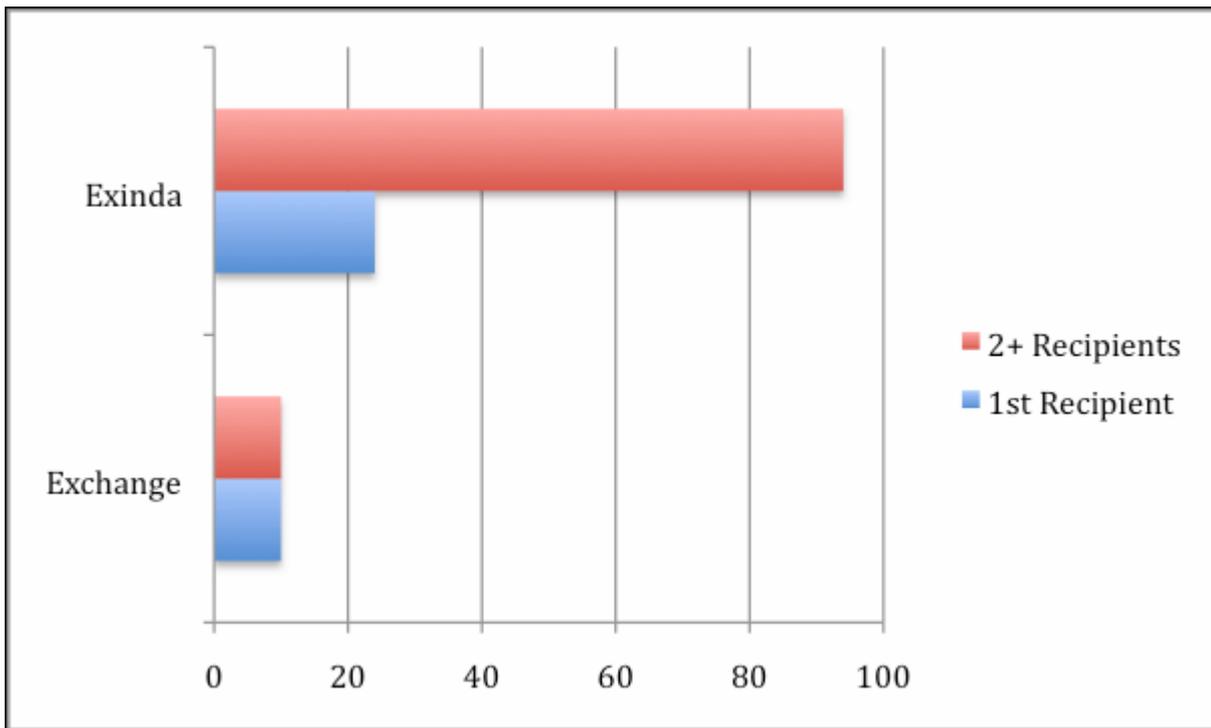
The following two tests compare the performance of Exinda's compression and WAN Memory technologies to Microsoft Outlook and Exchange compression.

### Exchange Results

- 1 recipient: 3.3MB was compressed to 3MB – 10% reduction
- 2nd and subsequent recipients: 3.3MB was compressed to 3MB – 10% reduction

### Exinda Results

- 1 recipient: 3.3MB was compressed to 2.5MB – 24% reduction
- 2nd and subsequent recipients: 3.3MB was compressed to 0.2MB – 94% reduction



Screenshot 287: Compression percentage achieved by Exinda and Exchange.

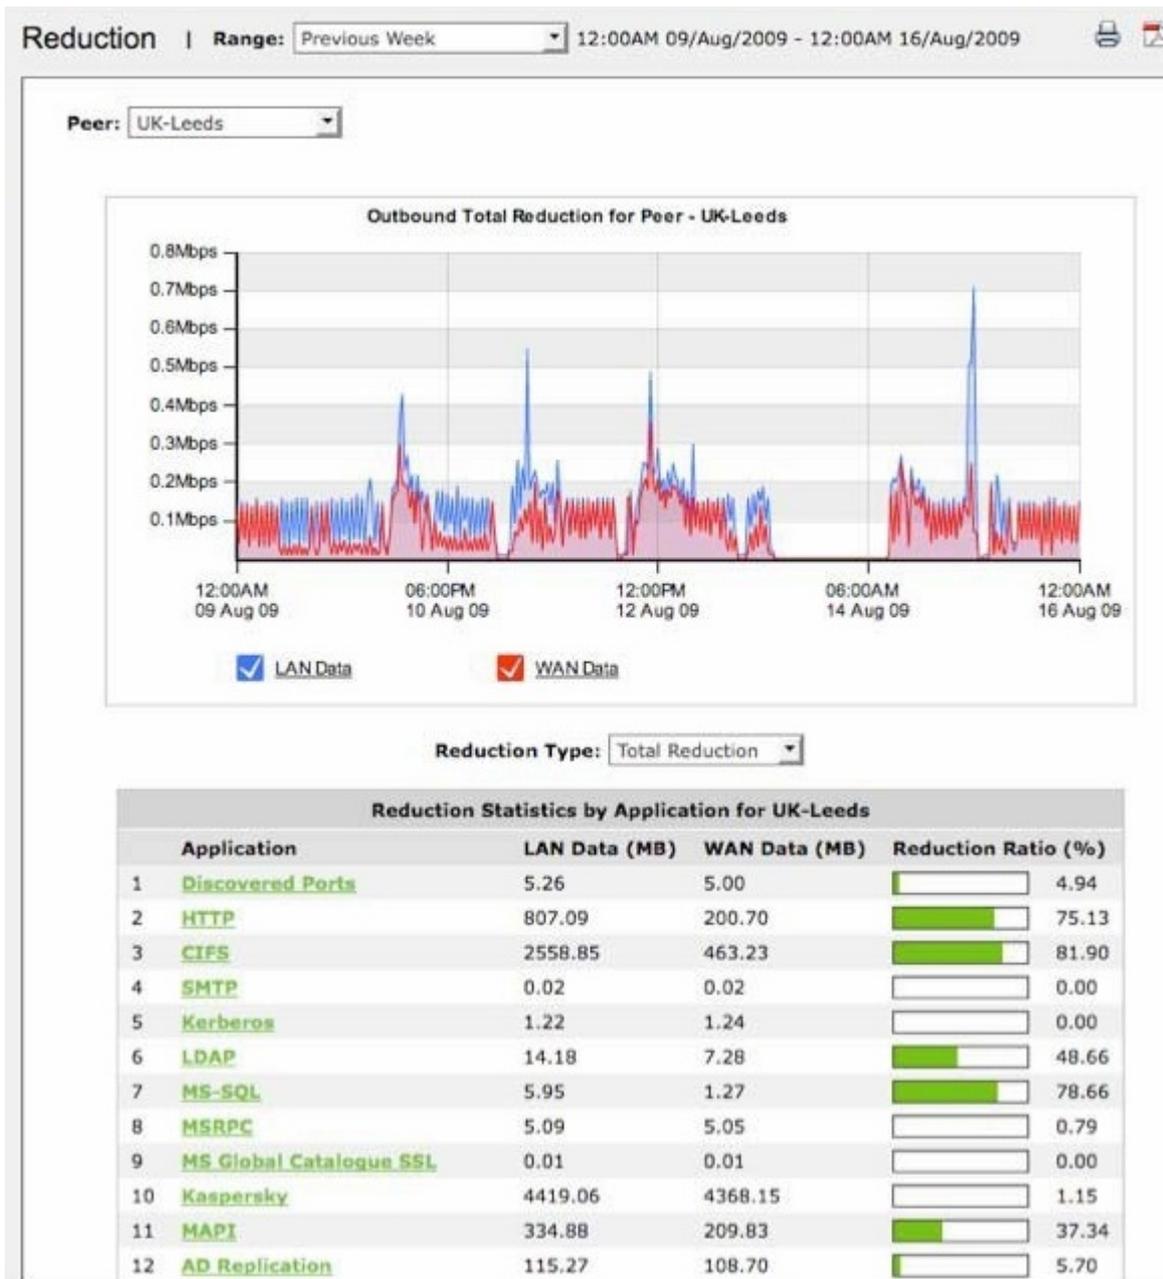
#### Test parameters

- » File size: 3.3MB PDF file
- » Link speed: 2Mbps, with 30ms RTT
- » Email: Exchange 2007 server/Outlook 2003 client
- » Exinda devices: 2 x Exinda 4800 appliances

#### View MAPI Acceleration Results

View the reduction in MAPI traffic on the network.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Monitor > Optimization** and switch to the **Reduction** tab. The report displays the reduction in MAPI traffic.



## 0.21.5 Data caching

**Edge Cache** – Provides acceleration of static web content such as HTML, GIF, JPEG, ZIP, RAR, ISO as well as dynamic content including YouTube, Google Video, Vimeo.

### NOTE

The Acceleration feature (including universal & protocol-specific acceleration) is licensed separately. Edge Cache is also licensed separately. Please contact your local Exinda representative if you wish to enable either of these features.

Related topics

[How Edge Cache Works](#)

[Edge Cache Configuration](#)

[Edge Cache Report](#)

## How Edge Cache Works

Edge Cache enables single-sided caching of Internet-based content, including web objects, videos and software updates. Edge Cache requires only one Exinda appliance.

When web objects are downloaded from the Internet or across WAN links, Edge Cache stores them at the edge of the network. When subsequent requests come for the same material, the content is quickly delivered from Edge Cache, without the need to download the data again over the WAN. The result is the ability to experience LAN speeds of WAN objects, and provide users with a better network experience.

Edge Cache also supports HTTPS sites allowing the appliance to be a forward proxy and decrypt content for caching. This is important as more and more applications and services are moving to the cloud. These SaaS-based applications are typically delivered over HTTPS and so to be effective, Edge Cache must support caching this HTTPS traffic.

Edge Cache also offers cache statistics, which provide insight into the amount of repetitive data being off-loaded from the WAN link, how cacheable the network data is, and how frequently the cache is being accessed.

## Caching Internet-based Content

To cache web traffic, a client-side Exinda appliance is put in line with the traffic. When a network user visits a URL with cacheable content, Edge Cache first determines if the content is available in its cache. If not, Edge Cache retrieves the content from the URL. Upon retrieving the content, it is stored in the cache with its expiry date as specified on the source website. This assumes that the content is cacheable and falls within the Edge Cache setting parameters, such as size of object and whether or not the URL is blacklisted. The next time a network user visits the same URL, Edge Cache determines that the content is available in the cache and that the content is not stale by looking at the object expiry date. The content is then served to the client from the cache, rather than retrieving from the URL over the WAN.

Edge Cache uses a least recently used (LRU) algorithm for expiring cached data to make room for new objects. This means the most popular and most used content is stored the longest. You also have the ability to manually clear the entire cache if desired.

Edge Cache operates as a transparent proxy since it is running on an inline device. As a result, your browsers do not need to be configured with an explicit proxy configuration.

## Caching Encrypted Internet-based Content

### VERSION INFO

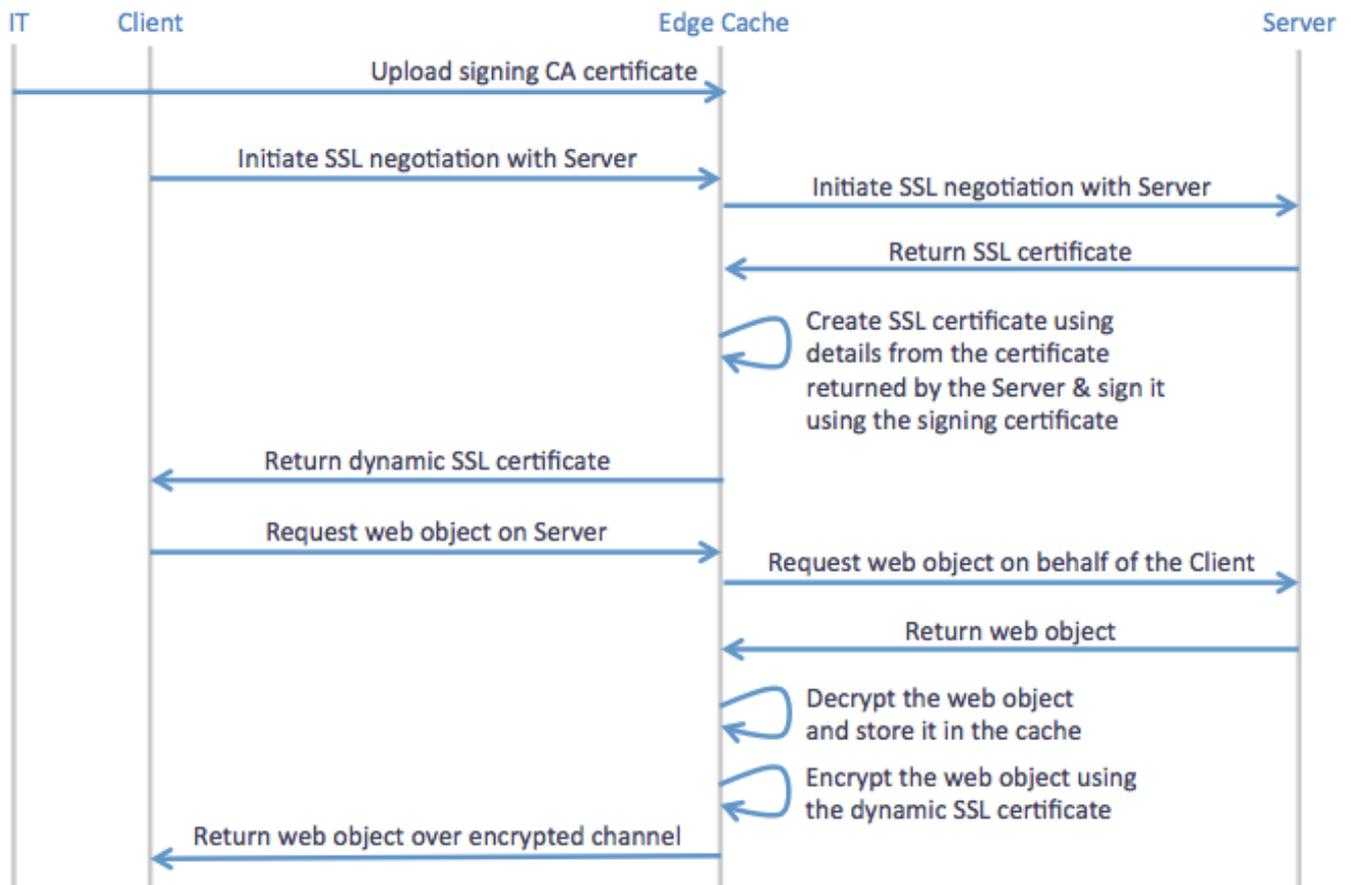
As of version 7.0.2, Edge Cache can cache HTTPS content, as well as HTTP content.

When the network user visits an HTTPS URL, if HTTPS caching is not enabled, Edge Cache is unable to determine what is being requested because the traffic is encrypted, including the URL being requested. Even if it could cache the encrypted data, the next request for the same HTTPS URL would not contain the same cached data because the encryption pattern would be different. By enabling HTTPS caching, Edge Cache is able to act as a forward proxy, and retrieve the content from the server, decrypt it, and provide it to the client over an encrypted communication channel. Later requests can then be served from the cache.

To support this feature, you need to upload a trusted certificate to the appliance, which is then used by Edge Cache to sign all dynamically generated site certificates. All client devices must trust this certificate as a signing authority.

To cache encrypted web traffic, the client tries to communicate with the HTTPS web server. The Exinda appliance intercepts, keeping the communication open with the client. Then Edge Cache tries to establish a conversation with the server. Upon receiving the certificate details from the server, Edge Cache extracts the certificate details, constructs a new certificate and signs it using the signing certificate that was loaded and specified in the Edge Cache settings. Edge Cache then presents this new certificate to the client. The client trusts this certificate because the details match its HTTPS URL request as the client has previously been told to trust anything signed by this signing certificate (see below). The communication negotiation between the client and Edge Cache is now complete. The client then requests the specific

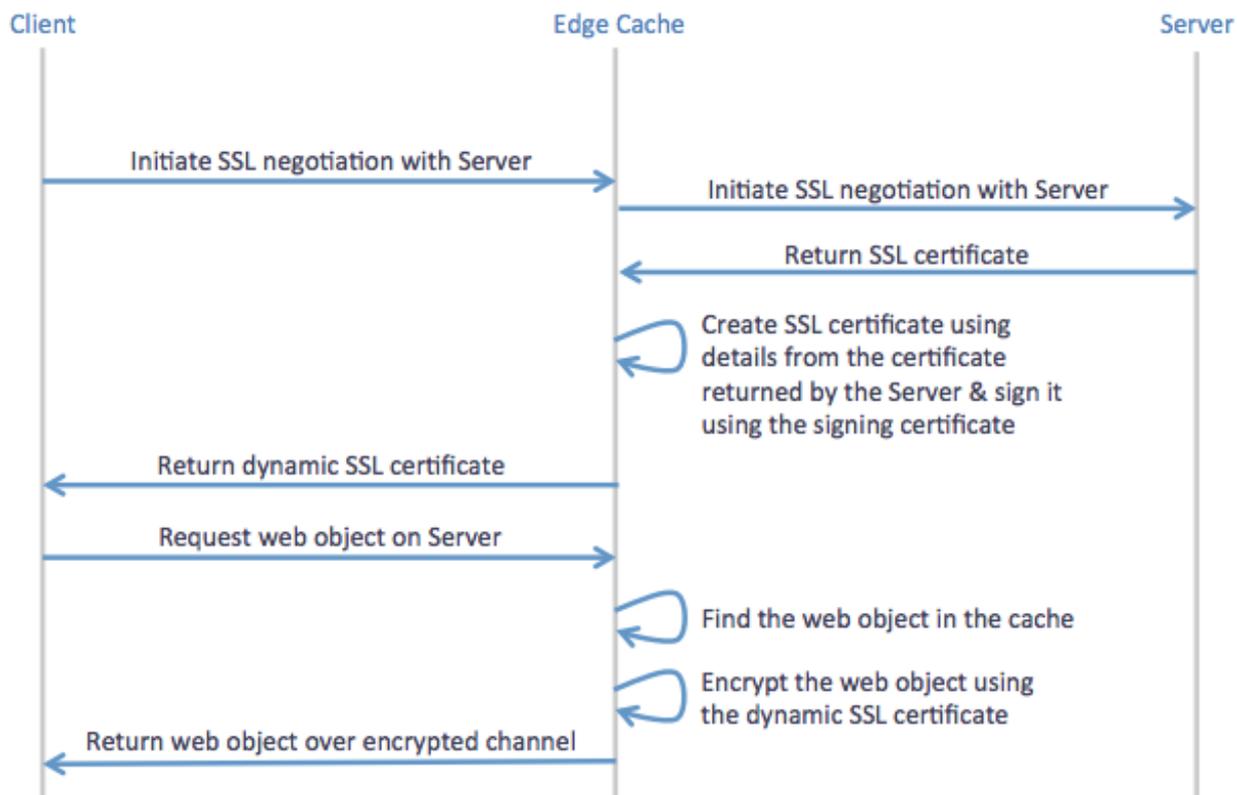
web object from Edge Cache as if it were the server. Edge Cache then requests the web object from the server over its previously established trusted connection. The server provides the content to Edge Cache, which then stores the content, if applicable. The connection with the server can be closed. Edge Cache then serves the content to the client and the connection with the client can then be closed.



Screenshot 288: Sequence of operations for caching encrypted traffic using Edge Cache

When determining whether the content should be stored in the cache, Edge Cache evaluates whether it is the appropriate size and whether it is white- or black-listed. The whitelist and blacklist can consist of source IP, destination IP, source domain, and destination domain. Note that the domains are resolved using DNS, so the resulting IP addresses are reverse mapped to determine the domain that is used to configure the Edge Cache engine.

The next time a client requests the same content, the same negotiation happens where the client requests a secure communication channel with the server, the Exinda appliance intercepts and forms a secure communication channel with the server, forges the certificate and establishes a secure communication channel with the client (on behalf of the server). The client then requests the specific web content. Edge Cache determines that the requested content is available in cache and serves it to the client. Edge Cache then closes the communication channels with both the server and the client.



Screenshot 289: Sequence of operations for use of previously cached objects from encrypted sites using Edge Cache

## Licensing

The Edge Cache Acceleration feature is a separately licensed component. To see if you are licensed for Edge Cache, go to **Configuration > System > Setup > License**. You are licensed for Edge Cache, if **Max Edge Cache Connections** is listed and is greater than 0. For more information, refer to [Licensing information](#) (page 145).

Please contact your local Exinda representative if you wish to enable this feature.

## Overview - Configuration & Usage

To use Edge Cache, you'll need to ensure all the required configuration is set.

» Configure the Edge Cache settings, including:

- indicating what size of objects you want to cache
- specifying how long you are willing to let Edge Cache wait for a response from the WAN when fetching objects
- specifying the signing certificate and private key if you plan to cache content from HTTPS sites
- specifying blacklisted sites to not cache, or specifying to only cache whitelisted sites (for HTTPS sites only)
- clearing out the cache, if desired.

### NOTE

Before version 7.0.2, Edge Cache requires you to restart the Edge Cache process after making any modifications to the Edge Cache configuration. With 7.0.2 and later, the Edge Cache process automatically restarts when needed.

For more information, refer to [Edge Cache Configuration](#) (page 684).

- » You also need to ensure that the DNS server configuration information is set. For more information, refer to [Configuring DNS](#) (page 694).
- » For HTTPS caching, you need to ensure that the certificate used for HTTPS caching is trusted by your clients. For more information, refer to [Preparing & Trusting a Certificate for Encrypted Traffic](#) (page 691).
- » If you have an upstream proxy in your environment, you can configure it as a proxy peer to ensure that Edge Cache can fetch content from the Internet.
- » You can add one or more Edge Cache policies to a virtual circuit in the Optimizer. For more information, refer to [Creating an Edge Cache Policy in the Optimizer](#) (page 695). Then you can start the Edge Cache process.
- » Since the appliance has dynamic disk partitioning, if needed, you can increase the cache storage capacity dynamically. For more information, refer to [Disk Storage Explained](#) (page 633).

Once Edge Cache is configured and started, you can monitor the caching performance.

- » You can look at the Edge Cache monitor to determine the reduction ratio or throughput comparing the LAN (cached and non-cached) data to the WAN (non-cached) data. The Edge Cache monitor will also report the requests per second versus the hits per second, where hits are the number of requests that could be satisfied by cached content. These charts can look at long term historical caching or can be as little as the last 5 minutes with 10 second samples. For more information, refer to [Monitoring edge cache reduction](#) (page 310).
- » You can also look at the Real Time page to see which conversations are passing through Edge Cache. Conversations with a blue background indicate that the flow passed through Edge Cache, however, it does not necessarily mean that any of the requests were satisfied by cached content. For more information, refer to [Monitoring edge cache traffic in real time](#) (page 249).

## Related topics

[Edge Cache Configuration](#)

[Edge Cache Report](#)

## Edge Cache Configuration

When web objects are downloaded from the Internet or across WAN links, Edge Cache stores them at the edge of the network. When subsequent requests come for the same material, the content is quickly delivered from the Edge Cache without the need to download the data again over the WAN. Edge Cache can cache web objects, videos, software updates, and other content on the WAN. You can control whether to exclude particular sources, or sizes of objects from being cached. The cache objects can also be shared among peer appliances if desired, meaning that if the content cannot be found in the cache of the appliance through which the traffic is passing, then the appliance can request the content from peer appliances. You are also able to clear the cache, if desired.

To learn the details of how edge cache works, see [How Edge Cache Works](#) or to learn all the details of how to configure Edge Cache including configuring DNS and HTTPS certificates and blacklists and whitelists, see the [Edge Cache](#).

You may want to limit the size of objects that can be cached. Since the cache storage size is limited, you may want to decide whether to allow a few large objects or lots of little objects. You should ensure that the size is aligned with the types of objects that you want to be cached. For example, iOS updates tend to be approximately 1 GB and Mac updates tend to be 6 GB and so if you want Edge Cache to help with caching of these updates, you'll want to ensure that the minimum and maximum allowed objects size accommodates these sizes.

You may want to blacklist particular sites if Edge Cache is not working properly and is preventing the network user from getting access to the site. Also, you may want to blacklist secure sites because you do not want to cache sensitive data such as financial or banking sites. Alternatively, if you want to only cache particular secure sites once they've been identified as important to your network, you can whitelist them such that only those sites listed will be cached. Whitelisting is only available for HTTPS caching. You can specify your whitelist and blacklist as source IP, destination IP, source domain, or destination domain. Domains are resolved using the DNS.

To use HTTPS caching, you will need to specify which signing certificate will be used by Edge Cache to generate a certificate to negotiate with the client on behalf of the server.

#### VERSION INFO

For versions prior to 7.0.2, Edge Cache requires you to restart the Edge Cache process after making any modifications to the Edge Cache configuration. With 7.0.2 and later, the Edge Cache process automatically restarts when needed.

#### NOTE

Objects in the traffic that are matched with an Edge Cache policy but are excluded from storage in Edge Cache due to these settings, will still pass through Edge Cache unprocessed and will be highlighted on the Real Time conversations screen in blue (indicating that they passed through and were evaluated by Edge Cache).

Where do I find this functionality?

Go to **Configuration > System > Optimization > Edge Cache**.

To set the range of object size that you would like to cache

| Memory Object Options |                                      |         |
|-----------------------|--------------------------------------|---------|
| Min Object Size       | <input type="text" value="0"/>       | kB      |
| Max Object Size       | <input type="text" value="8190000"/> | kB      |
| Connection Timeout    | <input type="text" value="20"/>      | seconds |

**Apply Changes**

Screenshot 290: Setting the size of the objects that can be cached

1. In the **Memory Object Options** area, type the minimum and maximum size of the objects to be cached. Only objects that are within this size range will be stored in Edge Cache.
2. Click the **Apply Changes** button.

#### VERSION INFO

For versions prior to 7.0.2, you will need to restart Edge Cache for the changes to take effect. Go to **Configuration > System > Optimization > Services**, and click the Edge Cache **Restart** button.

To set the fetch connection timeout

| Memory Object Options |                                         |
|-----------------------|-----------------------------------------|
| Min Object Size       | <input type="text" value="0"/> kB       |
| Max Object Size       | <input type="text" value="8190000"/> kB |
| Connection Timeout    | <input type="text" value="20"/> seconds |

Screenshot 291: Setting the connection timeout

1. In the **Memory Object Options** area and in the **Connection Timeout** field, type the maximum time in seconds that the Edge Cache will wait for a response from the WAN when fetching objects. You may need to increase this if connection timeouts are occurring regularly. Browsers typically return a message similar to the following when this occurs: (110) Connection timed out
2. Click the **Apply Changes** button.

#### VERSION INFO

For versions prior to 7.0.2, you need to restart Edge Cache for the changes to take effect. Go to **Configuration > System > Optimization > Services**, and click the Edge Cache **Restart** button.

To blacklist certain HTTP URLs to never cache

| URL             | Delete                                |
|-----------------|---------------------------------------|
| docs.exinda.com | <input type="button" value="Delete"/> |

| Add URL/Domain |                      |
|----------------|----------------------|
| URL            | <input type="text"/> |

Screenshot 292: HTTP caching with blacklisted sites

1. In the **Add URL/Domain** area, type the HTTP URL or domain that will be excluded from the Edge Cache.
2. Click the **Add URL** button.
3. Repeat until you have your desired blacklist.
4. Remove an HTTP URL or domain from the list by clicking the **Delete** button for the specified URL/domain.

## VERSION INFO

For versions prior to 7.0.2, you need to restart Edge Cache for the changes to take effect. Go to **Configuration > System > Optimization > Services**, and click the Edge Cache **Restart** button.

To blacklist certain (encrypted) HTTPS URLs to never cache

All https traffic can be cached according to the policy except those sites listed in the blacklist.

### HTTPS Caching

Enable caching of HTTPS Contents

Signing Certificate

HTTPS Site Exceptions  Only allow specified (whitelisted) HTTPS sites to be cached  
 Attempt to cache all HTTPS sites except for blacklisted exceptions

Apply Changes

Configure which HTTPS URL/Domains to not cache

| Type               | Value              | Delete |
|--------------------|--------------------|--------|
| Destination Domain | www1.royalbank.com | Delete |
| Destination Domain | easyweb.td.com     | Delete |
| Source IP          | 10.6.2.251/32      | Delete |

### Add IP or Domain

Type  Value

Add To List

Screenshot 293: Enabling HTTPS caching with blacklisted sites

1. In the **HTTPS Caching** area, select the **Enable caching of HTTPS content** checkbox.

**NOTE**

You cannot enable caching if DNS is not configured.

2. Select the **signing certificate** to use to re-sign the traffic.
3. Select **Attempt to cache all HTTPS sites except for blacklisted exceptions** from the drop-down list.
4. Click the **Apply Changes** button.
5. In the **Add IP or Domain** area, specify the type of traffic you want to blacklist and the **Value**(IP or domain name) for that type of traffic.

You can specify **Source IP**, **Destination IP**, **Source Domain**, or **Destination IP**.

Domains are resolved using the DNS. Ensure the domains are in the format that are required by DNS (i.e. without https://).

The specified domain name is resolved to an IP address, then the IP address is reverse mapped to the actual domain name that corresponds to that IP address. Note that the domain name may resolve to multiple IP addresses and multiple other domain names.

6. Click the **Add URL** button.
7. Repeat until you have your desired blacklist.

**NOTE**

The blacklist takes effect immediately as Edge Cache automatically restarts.

8. Remove an HTTPS URL or domain from the list by clicking the **Delete** button for the specified URL/domain.

To only allow specified whitelisted (encrypted) HTTPS URLs

Only https traffic specified in the policy and in this whitelist will be cached.

**HTTPS Caching**

|                                  |                                                                                                                                                                                          |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable caching of HTTPS Contents | <input checked="" type="checkbox"/>                                                                                                                                                      |
| Signing Certificate              | edge-cache-cert ▾                                                                                                                                                                        |
| HTTPS Site Exceptions            | <input checked="" type="radio"/> Only allow specified (whitelisted) HTTPS sites to be cached<br><input type="radio"/> Attempt to cache all HTTPS sites except for blacklisted exceptions |

Apply Changes

Configure which HTTPS URL/Domains can be cached

| Type               | Value             | Delete                                                               |
|--------------------|-------------------|----------------------------------------------------------------------|
| Destination IP     | 106.21.113.215/32 | <div style="border: 1px solid #ccc; padding: 2px 10px;">Delete</div> |
| Destination Domain | www.youtube.com   | <div style="border: 1px solid #ccc; padding: 2px 10px;">Delete</div> |

**Add IP or Domain**

|      |             |       |                                                                         |
|------|-------------|-------|-------------------------------------------------------------------------|
| Type | Source IP ▾ | Value | <div style="border: 1px solid #ccc; width: 150px; height: 20px;"></div> |
|------|-------------|-------|-------------------------------------------------------------------------|

Add To List

*Screenshot 294: Enabling HTTPS caching with whitelisted sites*

1. In the **HTTPS Caching** area, select the **Enable caching of HTTPS content** checkbox.

**NOTE**

You will not be able to enable caching if DNS is not configured.

2. Select the **signing certificate** to use to re-sign the traffic.
3. Select **Only allow specified whitelisted HTTPS sites to be cached** from the drop-down list.
4. Click the **Apply Changes** button.
5. In the **Add IP or Domain** area, specify the type of traffic you want on the whitelist and the **Value** (IP or domain name) for that type of traffic.

You can specify **Source IP**, **Destination IP**, **Source Domain**, or **Destination IP**.

Domains are resolved using the DSN. Ensure the domains are in the format that are required by DNS (i.e. without https://).

The specified domain name is resolved to an IP address, then the IP address is reverse mapped to the actual domain name that corresponds to that IP address. Note that the domain name may resolve to multiple IP address and multiple other domain names.

6. Click the **Add URL** button.

7. Repeat until you have your desired whitelist.

**NOTE**

The whitelist takes effect immediately as Edge Cache automatically restarts

8. Remove an HTTPS URL or domain from the list by clicking the **Delete** button for the specified URL/domain.

To not blacklist or whitelist any URLs

Set the blacklist mode, but do not add any URLs.

This allows all sites to be cached except those listed. Since the list is empty, all sites are allowed.

To clear the Edge Cache cache

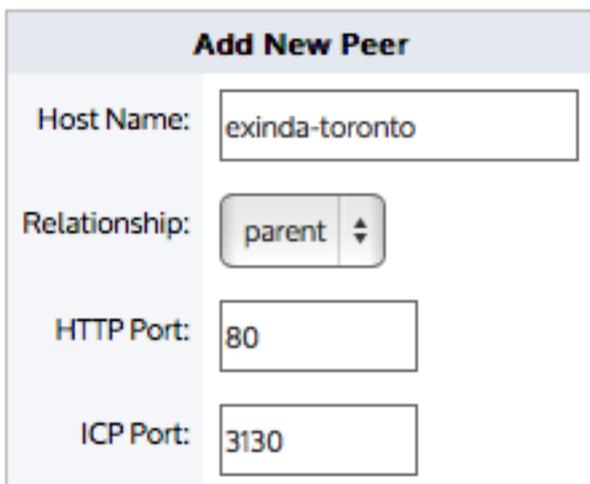
**The clear action removes all objects from the memory.**



Click the **Clear** button. This removes all objects from storage.

To manage with which appliance peers to share Edge Cache content

When Edge Cache appliance peers are specified, if the requested content is not available in the appliance's cache, Edge Cache can request the content from it's appliance peers.

A form titled "Add New Peer" with four input fields: "Host Name" containing "exinda-toronto", "Relationship" with a dropdown menu set to "parent", "HTTP Port" containing "80", and "ICP Port" containing "3130".

Screenshot 295: To add a new peer appliance for Edge Cache

Add New Peer

| Host Name      | Type   | HTTP Port | ICP Port | Edit | Delete |
|----------------|--------|-----------|----------|------|--------|
| exinda-toronto | parent | 80        | 3130     | Edit | Delete |

Screenshot 296: Peer appliances for Edge Cache

1. Click the **Add New Peer** button. An **Add New Peer** screen appears.
2. Type the **Host Name** of the peer.
3. Select the **Relationship** with the peer.
4. Type the **HTTP Port** of the new peer.
5. Type the **ICP Port** of the new peer.
6. Click the **Add New Peer** button. The peer appears in the list of configured peers.
7. Peers can be edited or deleted by clicking the **Edit** or **Delete** button for the specified peer.

### Preparing & Trusting a Certificate for Encrypted Traffic

When caching encrypted content, you need to specify a certificate that Edge Cache can use to create and sign a dynamically created certificate on behalf of the server. You need to ensure that this certificate is trusted by all the computers on your network that have traffic passing through Edge Cache. It is recommended that you create a self-signed CA certificate (as opposed to a self-signed certificate without the CA designation) to simplify the loading and trusting of the certificate by the computers in your network.

#### How to create a self-signed CA certificate for Edge Cache to use

When using Edge Cache for encrypted traffic, you must create and import a signing certificate in the **Certificates and Keys** store. This certificate and its corresponding key are used by Edge Cache to generate and sign dynamic SSL certificates for proxied sites. For all practical purposes, this certificate becomes a root certificate and you become a Root CA.

To create a CA certificate and private key using OpenSSL

Use the following OpenSSL commands:

```
openssl genrsa -out myCompanyCA.key 2048
openssl req -x509 -new -key myCompanyCA.key -out myCompanyCA.cer -days 1000
-subj /CN="myCompany CA"
```

This will generate two files: a key file and a .cer file that can be uploaded in the CA certificates UI.

To import your CA certificate and private key to the appliance's Certificates and Keys store

Screenshot 297: Import certificate in CA Certificates store

1. Go to **Configuration > System > Certificates > CA Certificates**.
2. Select the **Import Certificate** radio button.
3. Optionally, Type a **Name** for the certificate. If no name is specified, the filename of the certificate is used. Private keys are stored separately from certificates, and are automatically named the same as the certificate, with '\_key' appended to the end.
4. Select the **Certificate/Key Format**.
  - **PKCS#12**—Format used when the certificate and key are stored together, and usually have extensions such as `.pfx` and `.p12`.
  - **PEM**—Common format for certificates issued by Certificate Authorities. PEM certificates usually have extensions such as `.pem`, `.crt`, `.cer`, and `.key`. If PEM format is selected, an additional upload field is exposed so that the private key can be uploaded with the certificate.
5. If the key is password protected, in the **Key Passphrase** field type the password.
6. In the **Certificate File** field, click **Choose File** and navigate to the file to be uploaded to the appliance.
7. If the PEM format is selected, the private key must be uploaded. In the **Private Key File** field, click **Choose File** and navigate to the private key file.
8. Click **Import**.

The certificates are displayed in the **Certificates and Keys** table on the CA Certificates and Keys table on the CA Certificates tab. From the tables the contents of a certificate can be viewed, or the certificate can be deleted or exported.

#### How to export the certificate for use on client computers

If the browsers in your network do not trust the certificate, you may get a warning or the sites may fail to load. In this case, each computer needs to import the certificate so that the certificate will be trusted when negotiating with Edge Cache over SSL.

You will need to export the certificate from the appliance and import it to the desired computers.

To export the certificate from the appliance:

1. Go to **Configuration > System > Certificates > CA Certificates** and find your desired certificate in the list.
2. Export the certificate by clicking the **Export** button.
3. Ensure that the Export **Certificate Format** is set to **PEM**. The PEM format encodes the certificate and private key. It may include an entire certificate chain including public key, private key, and root certificates.
4. Click the **Save** button.

#### How to deliver and install the certificate on machines across your network

The method used to install the trusted certificate on client machines depends upon both the browser and the operating system:

##### Chrome and Explorer on Windows machines

Both Chrome and Internet Explorer on Windows Machines use the Windows certificate store to trust the Exinda Edge Cache SSL certificate. Two methods are available to distribute the certificates: using a domain controller or using a manual method.

To use a Domain Controller:

In this case, it is recommended that you follow the instructions provided by Microsoft to use a domain controller to distribute the certificate:

<http://technet.microsoft.com/en-us/library/cc772491.aspx>

#### NOTE

These instructions assume that you are using a domain controller or a workstation running the domain admin MMC snapins while logged into a domain as a domain admin. Some of the elements that are referred to in the instructions will not exist if you are using Windows Server 2008 R2.

To use the manual method:

Follow this method to add the certificates to the Trusted Root Certification Authorities store on each local computer.

1. In the Windows Search field, type **mmc**, and then press **ENTER** to launch the Console screen.
2. Click **File > Add/Remove Snap-in**.
3. Under Available snap-ins, click **Certificates**, and then click **Add** to move the Certificates option to the Selected snap-ins list.
4. In the pop-up window, select the **Computer Account** option under 'This snap-in will always manage certificates for' and then click **Next**.
5. Click **Local computer**, and click **Finish**.
6. In the console tree, double-click **Certificates**.
7. Right-click the **Trusted Root Certification Authorities store**.
8. Click **Import** to import the certificates and follow the steps in the Certificate Import Wizard.

##### Chrome on Linux machines

To install the trusted certificate on Linux machines for Chrome, you must use the NSS command line tools. To import a personal certificate and private key stored in a PKCS12 file, use the command below, substituting the details between the <> with the certificate file name.

```
pk12util -d sql:$HOME/.pki/nssdb -i <PKCS12_file.p12>
```

If the certificate was generated as a root CA certificate, use the following command, substituting the details between the <> with the certificate nickname and file name.

```
certutil -d sql:$HOME/.pki/nssdb -A -t "C,," -n <certificate nickname> -i <certificate filename>
```

## Firefox on Windows machines

Mozilla Firefox uses its own certificate store and requires a unique process to trust the certificate.

1. Launch the Firefox browser.
2. Go to **Options > Advanced > Certificates**.
3. Click **View Certificates**.
4. Click **Import**.
5. Navigate to the certificate you generated and exported from the Exinda appliance and import it.

## MAC OS

If installing the certificate on a MAC, you must use the **Keychain Access** program. To start the **Keychain Access** program, double-click certificate file.

If you are importing a CA certificate:

1. Double-click the exported PEM file for the CA certificate to start the Keychain Access program.
2. When prompted, type your computer's admin password.
3. The **Keychain Access** window appears. The certificate has automatically been installed with no additional steps.

If you are importing a non-CA certificate:

1. Double-click the exported PEM file for the non-CA certificate to start the Keychain Access program.
2. When prompted, type your computer's admin password.
3. In the **Keychain Access** window, select the **System** keychain to install for all users on the machine, or **Login** keychain to install only for the current user account.
4. Find the desired certificate in the list and right-click and select **Get Info**.
5. In the **Trust** section, select **Always Trust** for the **When using this certificate** drop-down list.

## Configuring DNS

If DNS is not configured, Edge Cache will not work properly. That is the DNS server(s) must either be specified on **Configuration > System > Network > DNS** or by selecting **DHCP** on **Configuration > System > Network > IP Address**.

The user interface will not allow you to get into a state where DNS is not configured when using HTTPS Caching. For instance, if DNS is not configured, then a warning will be shown next to the **HTTPS Caching** section and the enable checkbox will be disabled. Also if HTTPS Caching and DNS are configured, and you attempt to remove DNS configuration, then the system will warn you and prevent you from making the change.

### Caution

The system does not prevent using non-encrypted caching when DNS is not configured. In this case, the system will not warn you and Edge Cache will fail. The logs will report the issue.

To learn more about configuring DNS. See DNS and Domain Names Configuration in the main user guide.

## Creating an Edge Cache Policy in the Optimizer

Edge Cache works on outbound, HTTP based conversations (and HTTPS based conversations, for version 7.0.2 and later). To enable Edge Cache, create a policy that will capture the HTTP application traffic that you wish to cache and add the policy to the appropriate virtual circuit in the Optimizer. You can specify a subset of the network to use Edge Cache by specifying the source network object.

### CAUTION

You can create an Edge Cache policy on anything except a L7 signature. That is, you can set the policy filter using VLAN, source and/or destination network objects, ToS/DSCP, or applications that are based on protocol, ports, network objects, or DSCP.

### NOTE

It is possible to shape traffic that is being cached on the first pass. However, there is one extra consideration. Traffic hitting the Edge Cache engine is only shaped by the policy configured for the “Web” application-group or the HTTP/HTTPS application objects. For example, if a flow being classified as “Software Updates” is falling under the “Software Updates” policy, it is actually shaped as using the “Web” policy settings.

To learn more about configuring policies, see *Policy* in the main user guide.

| Edit Policy                                                       |                                                                                           |                                   |                                  |                                  |                                    |
|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------|-----------------------------------|----------------------------------|----------------------------------|------------------------------------|
| Policy Name: <input type="text" value="Cache Web for Marketing"/> | <input type="checkbox"/> Guaranteed Bandwidth: <input type="text" value="128"/> kbps      |                                   |                                  |                                  |                                    |
| Schedule: <input type="text" value="ALWAYS"/>                     | Burst (Max) Bandwidth: <input type="text" value="2048"/> kbps                             |                                   |                                  |                                  |                                    |
| Action: <input type="text" value="Optimize"/>                     | Burst Priority: <input type="text" value="1 (High)"/>                                     |                                   |                                  |                                  |                                    |
| Policy Enabled: <input checked="" type="checkbox"/>               | <input checked="" type="checkbox"/> Acceleration: <input type="text" value="Edge Cache"/> |                                   |                                  |                                  |                                    |
|                                                                   | <input type="checkbox"/> Packet Marking                                                   |                                   |                                  |                                  |                                    |
| Filter Rules:                                                     |                                                                                           |                                   |                                  |                                  |                                    |
| VLAN                                                              | Source                                                                                    | Direction                         | Destination                      | ToS/DSCP                         | Application                        |
| <input type="text" value="ALL"/>                                  | <input type="text" value="Marketing"/>                                                    | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="HTTP"/>  |
| <input type="text" value="ALL"/>                                  | <input type="text" value="Marketing"/>                                                    | <input type="text" value="Both"/> | <input type="text" value="ALL"/> | <input type="text" value="ALL"/> | <input type="text" value="HTTPS"/> |
| <input type="text" value=""/>                                     | <input type="text" value=""/>                                                             | <input type="text" value="Both"/> | <input type="text" value=""/>    | <input type="text" value=""/>    | <input type="text" value=""/>      |
| <input type="text" value=""/>                                     | <input type="text" value=""/>                                                             | <input type="text" value="Both"/> | <input type="text" value=""/>    | <input type="text" value=""/>    | <input type="text" value=""/>      |
| <input type="text" value=""/>                                     | <input type="text" value=""/>                                                             | <input type="text" value="Both"/> | <input type="text" value=""/>    | <input type="text" value=""/>    | <input type="text" value=""/>      |
| <input type="text" value=""/>                                     | <input type="text" value=""/>                                                             | <input type="text" value="Both"/> | <input type="text" value=""/>    | <input type="text" value=""/>    | <input type="text" value=""/>      |

To create an Edge Cache policy

1. Click **Configuration > Optimizer > Policies > Create New Policy** or click **Create New Policy** in the appropriate virtual circuit.
2. Type a name for the policy.
3. Select the **Acceleration** checkbox and select **Edge Cache** from the Acceleration list.
4. Create the filter rules for the policy, ensuring that HTTP (or HTTPS) or an application based on protocol, port, network object, or dscp is selected from the application list.
5. Click **Add New Policy**.

To apply Edge Cache to a subset of the network

1. When creating the policy, specify a network object as the source in the policy filter.
2. Click **Add New Policy**.

Only the traffic in that source network object will use Edge Cache.

### Pre-populate the cache

The pre-population feature allows you to warm the Exinda SMB, Edge, and WAN memory caches in advance of end users accessing the selected data. With a pre-populated cache, first time access of the data is served up locally with LAN-like performance.

The pre-population service uses the management IP address on the Exinda as the “Client” IP for this process. The IP address must either be configured on the inline bridge itself or on a dedicated Management Interface residing on the “inside network” of the Exinda.

Create and schedule a pre-population job to ensure content is added to the SMB or HTTP object cache before users request it.

#### Prerequisites for Pre-population

Before configuring pre-population, you must have the following:

- » Exinda 7.x or later
- » One of the following licenses:
  - Any license with acceleration capabilities
  - An Edge Cache license

#### NOTE

In an acceleration environment you need an active Community and Wan Memory for SMB prepopulation to work

#### Create a pre-population job in the Exinda Web UI

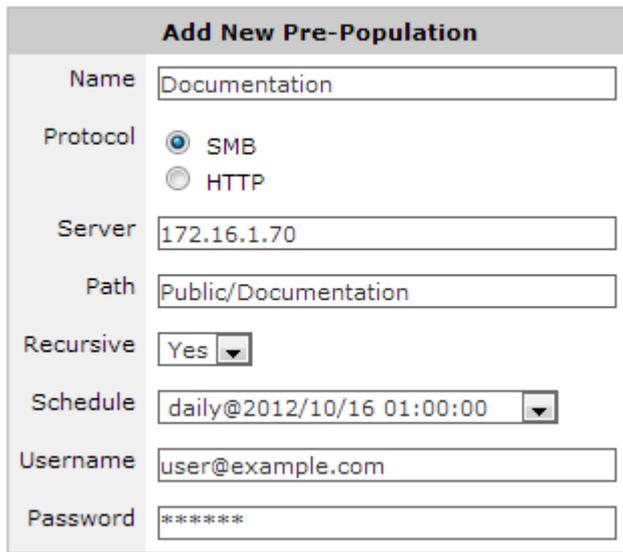
Add and configure pre-population jobs in the Exinda Web UI. SMB traffic can be stored in SMB object cache and WAN memory cache, and HTTP can be stored in either WAN memory cache or Edge cache depending on the configuration of your policies.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Configuration > System > Optimization** and click the **Pre Population** tab.
5. In the Add New Pre-Population area, type a name for the pre-population job.
6. Select the type of cache to pre-populate.
7. Type the hostname or IP address of the server containing the cache.
8. Type the path to the file or cache directory on the server.
9. To recursively fetch all the files in the specified directory, as well as those in sub-directories, select **Yes**.
10. Select the appropriate schedule job for how frequently the cache should be updated. If no jobs are available, [create a job schedule](#).

11. If required, type the username and password for the server.
12. Click **Add New**. The job appears in the list of configured Pre-population jobs.
13. To begin pre-populating the selected cache, in the list of configured pre-population jobs click **Start**. The content from the specified location begins to populate the cache.
14. To view the communication between the server and the Exinda appliance, click **Monitor > Real Time > Conversations**.

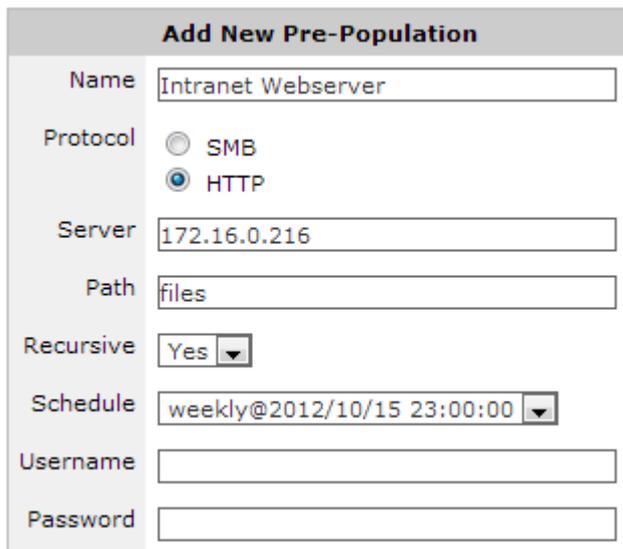
Examples -

» **Intranet Windows File Share**—all files and folders from the "Public\Documentation" folder. Show image...



| Add New Pre-Population |                                                                    |
|------------------------|--------------------------------------------------------------------|
| Name                   | Documentation                                                      |
| Protocol               | <input checked="" type="radio"/> SMB<br><input type="radio"/> HTTP |
| Server                 | 172.16.1.70                                                        |
| Path                   | Public/Documentation                                               |
| Recursive              | Yes ▼                                                              |
| Schedule               | daily@2012/10/16 01:00:00 ▼                                        |
| Username               | user@example.com                                                   |
| Password               | *****                                                              |

» **Intranet HTTP**—all files and folders from a directory called "files" on an intranet web server. Show image...



| Add New Pre-Population |                                                                    |
|------------------------|--------------------------------------------------------------------|
| Name                   | Intranet Webserver                                                 |
| Protocol               | <input type="radio"/> SMB<br><input checked="" type="radio"/> HTTP |
| Server                 | 172.16.0.216                                                       |
| Path                   | files                                                              |
| Recursive              | Yes ▼                                                              |
| Schedule               | weekly@2012/10/15 23:00:00 ▼                                       |
| Username               |                                                                    |
| Password               |                                                                    |

### Create a Pre-population Job in the CLI

Add and configure pre-population jobs using CLI commands.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.

3. Click **Login**.

5. Click **Configuration > System > Tools** and click **Console**.

6. Log into the console. The prompt `exinda-appliance >` is displayed.

7. To enter configuration mode, type `enable` then `configure terminal`. The prompt `exinda-appliance (config) #` is displayed.

8. To create the pre-population job, at the prompt type the following command with the appropriate parameters:

```
acceleration prepopulate <name> {location|username|password|recursive|start|stop}
```

where:

- `<name>` is the name of the pre-population job.
- `location cifs <server> <path>` configures an SMB server and path.
- `location http <url>` configures an HTTP path.
- `username | password` is the authentication information for the server.
- `recursive` fetches all the files in the specified directory, as well as those in sub-directories.
- `start | stop` controls the transfer.

9. To view the pre-population rules and the status of those rules, run the following command:

```
show acceleration prepopulate
```

#### EXAMPLE

In the following examples, the pre-populate job is called `videos`.

```
» All files and folders from the "Public\Documentation" folder on a Windows file share: exinda-appliance (config) # acceleration prepopulate videos location cifs MYSERVER public\documentation
```

```
» Including a username and password for authenticating to the server: Exinda-appliance (config) # acceleration prepopulate videos username myname password mypassword
```

```
» Download the entire path specified, including sub-folders: Exinda-appliance (config) # acceleration prepopulate videos recursive
```

```
» Start the pre-populate job: Exinda-appliance (config) # acceleration prepopulate start
```

```
» Display the status of the pre-population job: Exinda-appliance (config) # show acceleration prepopulate videos
```

#### Schedule a Pre-population Job in the CLI

Scheduling a pre-population job can be done from the Exinda Web UI or through the CLI. The following example schedules a pre-population request daily at 12:30AM for the "Documentation" job configured in an earlier section of this guide.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).

2. Key-in the **User** and **Password**.

3. Click **Login**.

5. Click **Configuration > System > Tools** and click **Console**.
  6. Log into the console. The prompt `exinda-appliance >` is displayed.
  7. To enter configuration mode, type `enable` then `configure terminal`. The prompt `exinda-appliance (config) #` is displayed.
  8. To schedule the pre-population job type the following commands:
- ```
exinda-appliance (config) # job <job_number> command <command_number> "acceleration
prepopulate <prepopulation_name> start"
```

**EXAMPLE**

```
exinda-appliance (config) # job 100 command 1 "acceleration prepopulate
Documentation start"
exinda-appliance (config) # job 100 schedule daily time 00:30:00
```

## 0.22 Authentication

Learn the process of authenticating users and user groups on your network.

---

|                                             |     |
|---------------------------------------------|-----|
| 0.22.1 Display a List of Active Users ..... | 699 |
| 0.22.2 Local User Accounts .....            | 700 |
| 0.22.3 AAA .....                            | 701 |
| 0.22.4 LDAP Authentication .....            | 702 |
| 0.22.5 Radius Authentication .....          | 702 |
| 0.22.6 TACACS+ authentication .....         | 703 |

---

### 0.22.1 Display a List of Active Users

Active Users lists the users currently logged into either the Web UI or the CLI.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Authentication > Active Users**.

The table below shows an example of the currently logged in users along with the session type, IP address and the session idle time in seconds.

| Active Users |        |              |                |
|--------------|--------|--------------|----------------|
| Username     | Line   | Host         | Idle (seconds) |
| admin        | pts/0  | 172.16.0.239 | 1544           |
| admin        | web/73 | 172.16.0.239 | 2096           |
| monitor      | web/75 | 172.16.0.115 | 2762           |
| admin        | web/76 | 172.16.0.239 | 0              |

## 0.22.2 Local User Accounts

Local User Accounts allows you to add/remove local user accounts as well as change local user's passwords.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Authentication > Local User Accounts**.

The table at the top of the page lists the configured local users and their capabilities.

| Local Users                      |            |                                     |
|----------------------------------|------------|-------------------------------------|
| User                             | Capability | Enabled                             |
| <input type="checkbox"/> admin   | admin      | <input checked="" type="checkbox"/> |
| <input type="checkbox"/> monitor | monitor    | <input checked="" type="checkbox"/> |

6. To remove local user accounts from the Exinda appliance or to temporarily disable an account, select the checkbox for the user and click **Remove User** or **Disable User**.

7. To add a new Local User Account, specify a username and select a capability. Click **Add User**. Admin users have full read-write access to the Exinda appliance. Monitor users have read-only access.

| Add New User |                                        |
|--------------|----------------------------------------|
| User Name    | <input type="text"/>                   |
| Capability   | Admin <input type="button" value="v"/> |

8. Create a password for a new user, or change the password for an existing user by selecting the username you wish to create or change the password for and enter a new password. Click **Change Password**.

**Change Password**

User Name

New Password

Confirm Password

### 0.22.3 AAA

AAA configures how remote users should authenticate to the Exinda appliance and what privileges they should receive.

1. To configure AAA, navigate to **Configuration > System > Authentication > AAA** on the Web UI, advanced mode.
2. Specify the order in which users are authenticated. When a user logs in, the Exinda appliance will try to authenticate them using the authentication methods specified here, in the order they are configured.

**Authentication Method List**

First Method

Second Method

Third Method

Fourth Method

#### NOTE

This setting is required if you are using a remote access mechanism such as [LDAP](#), [Radius](#) or [TACACS+](#).

3. Click **Apply Changes**.
4. Control what privileges remotely authenticated users receive when they login to the Exinda appliance.

**Authorization**

Map Order

Map Default User

|                  |              |                                                                                                                                       |
|------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Map Order        | remote-first | Apply user privileges supplied by the remote authentication mechanism first. If that fails, use the 'Map Default User' setting below. |
|                  | remote-only  | Apply user privileges supplied by the remote authentication mechanism first. If that fails, the user will not be authenticated.       |
|                  | local-only   | Use the 'Map Default User' setting below.                                                                                             |
| Map Default User |              | If the 'local-only' option is selected above, the user will be given the same privileges as the local user account selected here.     |

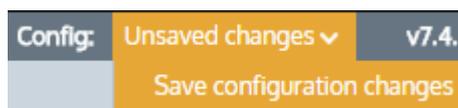
5. Click **Apply Changes**.

## 0.22.4 LDAP Authentication

LDAP authentication allows you to configure the Exinda appliance to authenticate user login attempts with a remote LDAP (including Active Directory) server.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Ensure LDAP is selected as an Authentication Method on the [AAA](#) page.
6. Click **Configuration > System > Authentication** and switch to the **LDAP** tab.
7. Define the global LDAP authentication options. Click **Apply Changes**.
8. Specify the hostname or IP address of the remote LDAP server. IPv4 or IPv6 addresses can be specified. Multiple LDAP servers may be defined.
9. Click **Add New LDAP Server**.
10. To remove an LDAP servers from the Exinda appliance, select the checkbox for the server and click **Remove Server**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.

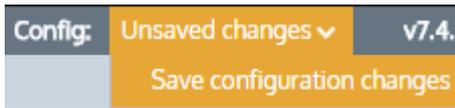


## 0.22.5 Radius Authentication

Radius authentication allows you to configure the Exinda appliance to authenticate user login attempts with a remote Radius server.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Ensure RADIUS is selected as an Authentication Method on the [AAA](#) page.
6. Click **Configuration > System > Authentication** and switch to the **Radius** tab.
7. Define the global RADIUS settings.
8. Click **Apply Changes**.
9. Specify the hostname or IP address of the remote Radius server. IPv4 addresses can be specified. Multiple Radius servers may be defined.
10. Click **Add New RADIUS Server**.
11. To remove Radius servers from the Exinda appliance, select the checkbox for the server and click **Remove Server**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.

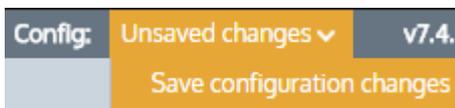


## 0.22.6 TACACS+ authentication

TACACS+ authentication allows you to configure the Exinda appliance to authenticate user login attempts with a remote TACACS+ server.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Ensure TACACS+ is selected as an Authentication Method on the [AAA](#) page.
6. Click **Configuration > System > Authentication** and switch to the **TACACS+** tab.
7. Define global TACACS+ authentication options.
8. Click **Apply Changes**.
9. Specify the hostname or IP address of the remote TACACS+ server. IPv4 addresses can be specified. Multiple TACACS+ servers may be defined.
10. Click **Add New TACACS+ Server**.
11. To remove TACACS+ servers from the Exinda appliance, select the checkbox for the server and click **Remove Server**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.



## 0.23 System Maintenance

Learn to maintain your Exinda Appliances.

The Manage System Configuration screen allows you to download, save, switch, revert and delete system configuration files. You can learn about backing up your configuration as well as importing and exporting your config system.

---

|                                          |     |
|------------------------------------------|-----|
| 0.23.1 Manage System Configuration ..... | 703 |
| 0.23.2 Factory Defaults .....            | 706 |
| 0.23.3 Reboot/Shutdown .....             | 706 |

---

### 0.23.1 Manage System Configuration

The Manage System Configuration screen allows you to download, save, switch, revert and delete system configuration files.

## NOTE

To Manage System Configuration, navigate to **Configuration > System > Maintenance > Manage Config** on the Web UI, advanced mode.

The table below lists the available system configuration files. There will be a check mark next to the active configuration. Clicking on the configuration file name will display the text-based version of the configuration file in the window at the bottom of this page. Clicking on the 'Download' icon next to the configuration file will allow you to download and save/backup the text-based version of the configuration file.

| Configuration Files                                  |                                     |                                                                                   |
|------------------------------------------------------|-------------------------------------|-----------------------------------------------------------------------------------|
| Filename                                             | Active                              | Download                                                                          |
| <input type="checkbox"/> <a href="#">initial.bak</a> |                                     |  |
| <input type="checkbox"/> <a href="#">initial</a>     | <input checked="" type="checkbox"/> |  |

- Delete the selected configuration(s).
- Make the selected configuration active and apply it to the system. (Select only one)
- Download the selected configuration as a binary file. (Select only one)

By selecting a configuration file and using the buttons above, you can delete the selected files from the system, switch to the selected configuration or download the selected configuration file in binary format.

The form below allows you to control the active and running configuration. If there are unsaved changes to the active configuration, this is known as the 'running configuration'.

| Active Configuration                   |                                                                                            |
|----------------------------------------|--------------------------------------------------------------------------------------------|
| <input type="button" value="Save"/>    | Save the running configuration to the active configuration file.                           |
| <input type="button" value="Revert"/>  | Discard the running configuration and apply the contents of the active configuration file. |
| <input type="button" value="Save As"/> | Save the running configuration to a new file and make it active.                           |
|                                        | New filename: <input type="text"/>                                                         |

You can save the running configuration and make it the active configuration, revert the running configuration back to the previously saved state of the active configuration, or save the running configuration to a new configuration file and make that the new active configuration.

## How to backup your Appliance settings

It is recommended to take a backup of your Exinda Appliance configuration during:

- » Hardware upgrade
- » RMA
- » Disk replacement
- » Re-manufacture
- » Diagnostics by TAC
- » Firmware upgrade

Generally users configure Exinda once and the configuration file doesn't have to be changed over and over again. Hence simply use the steps below to save the configuration file locally. If for any reason a backup is required periodically, you can also schedule a job for the same. Go to **System > Setup > Scheduled Jobs**.

There are two types of Exinda configuration files:

- » Binary
- » Text (recommended)

To download and save the configuration file:

1. Go to **Configuration > System > Maintenance > Manage Config**.
2. Find the configuration you wish to export. The current active configuration display a green check mark in front of it.
  - To save the configuration as a text file, click the save icon in the **Download** column beside it. It takes some time to generate. The generated text file contains all the CLI commands to replicate the configuration.
  - To save the configuration as a binary file, click the filename.

## Import System Configuration

The Import System Configuration screen allows you to import previously saved or backed-up system configuration files.

### NOTE

To Import System Configuration, navigate to **Configuration > System > Maintenance > Import Config** on the Web UI, advanced mode.

The form below can be used to upload system configurations that have been saved locally on the PC.

**Upload Configuration**

|                                                           |                                                   |                      |                                          |
|-----------------------------------------------------------|---------------------------------------------------|----------------------|------------------------------------------|
| <input checked="" type="radio"/>                          | <b>Upload local binary file:</b>                  | <input type="text"/> | <input type="button" value="Browse..."/> |
| (To be saved as separate file with its original name)     |                                                   |                      |                                          |
| <input type="radio"/>                                     | <b>Upload local text file:<br/>(CLI commands)</b> | <input type="text"/> | <input type="button" value="Browse..."/> |
| (To be executed immediately in the running configuration) |                                                   |                      |                                          |

Screenshot 298: Upload system configurations

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Upload local binary file | Use this option to upload a saved binary configuration file. This file would have been downloaded as a binary file from the <a href="#">System &gt; Maintenance &gt; Manage Config</a> page. Once this file is uploaded, it will appear in the list of available configuration files on the <a href="#">System &gt; Maintenance &gt; Manage Config</a> page.                               |
| Upload local text file   | Use this option to upload a text file containing CLI commands. The CLI commands will be executed in order and any configuration changes will be applied to the running configuration. This text file can contain one or more CLI commands or could be a complete text-based system configuration file downloaded from the <a href="#">System &gt; Maintenance &gt; Manage Config</a> page. |

Use the form below to execute a batch of CLI commands on the Web UI. The CLI commands will be executed in order and any configuration changes will be applied to the running configuration.

**Execute CLI Commands**

(To be executed immediately in the running configuration)

Execute CLI commands

### 0.23.2 Factory Defaults

The Factory Defaults screen allows you to restore the configuration of the Exinda appliance to factory default settings. This includes removing any system logs, WAN Memory cache, and monitoring statistics.

#### NOTE

To restore Factory Defaults, navigate to **Configuration > System > Maintenance > Factory Defaults** on the Web UI, advanced mode.

When restoring Factory Default settings, network connectivity settings such as the IP address, DNS servers and Default Gateway are preserved. There is also an option to preserve any monitoring data. To preserve monitoring data tick the 'Preserve monitoring' box prior to restoring the factory default settings.

Preserve monitoring data

Restore Factory Defaults

After performing a Factory Defaults, the Exinda appliance will automatically reboot.

### 0.23.3 Reboot/Shutdown

The Reboot/Shutdown screen allows you to configure Reboot options as well as gracefully shutdown the Exinda appliance in order to reboot it or power it down.

In this area of the Exinda Web UI you can:

[Reboot the Exinda Appliance](#)

[Automatically Reboot the Exinda Appliance](#)

[Shutdown the Exinda Appliance](#)

[Reboot the Exinda Appliance](#)

After a new version of the ExOS firmware is installed, you must reboot the appliance.

## CAUTION

Any unsaved configuration changes will be lost if the Exinda appliance is rebooted or shutdown without saving the changes first.

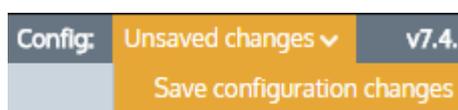
1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Maintenance** and switch to the **Reboot / Shutdown** tab.
6. (Optional) Schedule the Exinda Appliance to reboot at a specific date or time.
  - a. Check **Schedule Reboot**.
  - b. Enter the date and time when the appliance should reboot.
7. Select the reboot mode from the list.
  - **Fast Reboot**—This is a soft reboot and will reboot the operating system only. This does not reboot the hardware and does not reload the BIOS.
  - **Slow Reboot**—This is a hard reboot and will reboot the entire appliance. Use this option to access the BIOS or other start-up options.
8. Click **Reboot**. Rebooting the Exinda Appliance may take a few minutes to restart.

## Automatically Reboot the Exinda Appliance

If the Exinda Appliance becomes unresponsive, the System Watchdog can automatically reboot the appliance.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Maintenance** and switch to the **Reboot / Shutdown** tab.
6. In the System Watchdog area, select **Enable**.
7. Click **Apply Changes**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.



## Shutdown the Exinda Appliance

If the Exinda appliance needs to be powered off, shut it down from within the Exinda Web UI.

## IMPORTANT

Any unsaved configuration changes will be lost if the Exinda appliance is rebooted or shutdown without saving the changes first.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Maintenance** and switch to the **Reboot / Shutdown** tab.
6. Click **Shutdown**.

The Exinda Appliance will not restart, and must be physically powered on again.

## 0.24 System Tools

Learn about the various system tools available on your Exinda Appliance to assist you in day-to-day operations.

Exinda Appliance provides you with a set of network utilities that enables you to monitor network activity, gather network information and audit network devices.

Topics in this section:

---

|                                                   |     |
|---------------------------------------------------|-----|
| 0.24.1 Ping .....                                 | 708 |
| 0.24.2 Traceroute .....                           | 709 |
| 0.24.3 DNS Lookup .....                           | 709 |
| 0.24.4 Query a remote IPMI Exinda appliance ..... | 710 |
| 0.24.5 iPerf Client .....                         | 711 |
| 0.24.6 iPerf Server .....                         | 712 |

---

### 0.24.1 Ping

Use the Ping Tool to test network connectivity from the Exinda appliance to other hosts on the WAN or Internet.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Tools > Ping**.

IPv4 Host:

IPv6 Host:

```
PING ipv6.google.com(2404:6800:8007::63) 56 data bytes
64 bytes from 2404:6800:8007::63: icmp_seq=0 ttl=54 time=220 ms
64 bytes from 2404:6800:8007::63: icmp_seq=1 ttl=54 time=197 ms
64 bytes from 2404:6800:8007::63: icmp_seq=2 ttl=54 time=208 ms
64 bytes from 2404:6800:8007::63: icmp_seq=3 ttl=54 time=225 ms

--- ipv6.google.com ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3009ms
rtt min/avg/max/mdev = 197.239/212.949/225.904/11.118 ms, pipe 2
```

6. In the **IPv4 host** or **IPv6 host** field, specify an IP address or fully qualified domain name to attempt to ping.

7. Click **Ping**. It may take a few seconds for the ping operation to complete and display the results.

## 0.24.2 Traceroute

Use the Traceroute Tool to determine the network hops from the Exinda appliance to other hosts on the WAN or Internet.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).

2. Key-in the **User** and **Password**.

3. Click **Login**.

5. Click **Configuration > System > Tools > Traceroute**.

Host:

```
traceroute to ipv6.google.com (2404:6800:8007::68), 30 hops max, 40 byte packets
 1  2001:44b8:62:690::1  1.783 ms  1.753 ms  1.747 ms
 2  2001:44b8:61::1fc  52.539 ms  53.961 ms  54.147 ms
 3  2001:44b8:8060:8000::1  55.682 ms  56.831 ms  57.364 ms
 4  2001:44b8:8060:e::1  58.248 ms  *  *
 5  2001:44b8:8060:1::a  83.433 ms  *  *
 6  2001:4860:1:1:0:1283:0:4  86.152 ms  85.641 ms  86.588 ms
 7  2001:4860::1:0:9f7  92.365 ms  103.509 ms  2001:4860::1:0:9f8  102.835 ms
 8  2001:4860::1:0:165  210.179 ms  209.501 ms  209.033 ms
 9  2001:4860:0:1::e7  216.582 ms  215.693 ms  225.739 ms
10  2404:6800:8007::68  213.035 ms  212.868 ms  219.553 ms
```

6. In the **Host** field, specify an IPv4 or IPv6 Address, or fully qualified domain name to attempt to traceroute.

7. Click **Traceroute**. It may take a few seconds for the operation to complete and display the results.

## 0.24.3 DNS Lookup

Use the DNS Lookup Tool to have the Exinda appliance query the configured DNS servers to resolve the specified domain name.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).

2. Key-in the **User** and **Password**.

3. Click **Login**.

5. Click **Configuration > System > Tools > DNS Lookup**.

Domain:

```
www.google.com has address 173.194.77.105
www.google.com has address 173.194.77.106
www.google.com has address 173.194.77.147
www.google.com has address 173.194.77.99
www.google.com has address 173.194.77.103
www.google.com has address 173.194.77.104
www.google.com has IPv6 address 2607:f8b0:4003:c01::68
```

6. In the Domain field, specify a fully qualified domain name to look up.

7. Click **Lookup**. It may take a few seconds for the operation to complete and display the results.

### 0.24.4 Query a remote IPMI Exinda appliance

Use the IPMI Tool to query the power status, power cycle/power off, or reset a remote Exinda appliance via IPMI. The remote appliance must have [enabled IPMI access](#).

**Power Control Options**

Command

**Remote IPMI Login Details**

IPv4 Address

Username

Password

To perform an IPMI action on a remote Exinda appliance

1. Select the desired action from the **Power Control Options** drop down selection.

2. Enter the IPMI **IPv4 Address** of the remote appliance.
3. Enter the IPMI authentication details for the remote appliance.
  - The default username is admin.
  - The default password is exinda.
4. Click **Do Power Action**.

Example: Power cycle the Exinda appliance with IPMI address 192.168.110.61 -

```
ipmi power address 192.168.110.61 username admin password exinda control cycle
```

Example: Show the current power state of the Exinda appliance with IPMI address 192.168.110.61 -

```
show ipmi power address 192.168.110.61 username admin password exinda
```

### 0.24.5 iPerf Client

iPerf is a tool used for network throughput measurements. To function, it requires that two devices must be running the iPerf software to obtain bandwidth metrics between two endpoints. One device plays the role of the server and the other plays the role of the client. In Exinda, there is a Web User Interface option to configure an Exinda appliance as an iPerf client:

To configure an Exinda Appliance as the iPerf client:

1. Click **Configuration > System > Tools > Iperf Client**.

|                              |            |                               |         |           |              |              |
|------------------------------|------------|-------------------------------|---------|-----------|--------------|--------------|
| Ping                         | Traceroute | DNS Lookup                    | Console | IPMI      | Iperf Client | Iperf Server |
| Server: <input type="text"/> |            | Options: <input type="text"/> |         | Run Tests |              |              |

2. In the **Server** field, type the **IP address** or **Host Name** of an iPerf server that is already running.

|                                                 |            |                               |         |           |              |              |
|-------------------------------------------------|------------|-------------------------------|---------|-----------|--------------|--------------|
| Ping                                            | Traceroute | DNS Lookup                    | Console | IPMI      | Iperf Client | Iperf Server |
| Server: <input type="text" value="10.10.5.10"/> |            | Options: <input type="text"/> |         | Run Tests |              |              |

3. Click **Run Tests** to view the test results.

Example Results:

```

EXAMPLE
-----
Client connecting to 10.10.1.201, TCP port 5001
TCP window size: 23.2 KByte (default)
-----
[ 3] local 10.10.1.200 port 58760 connected with 10.10.1.201 port 5001
[ ID] Interval Transfer Bandwidth
[ 3] 0.0-10.0 sec 4.74 GBytes 4.07 Gbits/sec
  
```

### 0.24.6 iPerf Server

iPerf is a tool used for network throughput measurements. To function, it requires that two devices must be running the iPerf software to obtain bandwidth metrics between two endpoints. One device plays the role of the server and the other plays the role of the client. In Exinda, there is a Web User Interface option to configure an appliance as an iPerf server:

To configure an Exinda Appliance as the iPerf Server:

Use the iPerf Server tab when the Exinda appliance is the designated server.

1. Click **Configuration > System > Tools > Iperf Server**.

|              |            |             |         |                               |              |                        |
|--------------|------------|-------------|---------|-------------------------------|--------------|------------------------|
| Ping         | Traceroute | DNS Lookup  | Console | IPMI                          | Iperf Client | Iperf Server           |
| Start Server |            | Stop Server |         | Options: <input type="text"/> |              | Server status: Stopped |

2. (Optional), By default, an iPerf server listens to TCP packets on port 5001. However, you can use the following list of options to modify this condition:

Usage: `iperf [-s|-c host] [options]`

Example: `iperf [-h] --help [-v|--version]`

Options for both Clients and Servers

`-f, --format [kmKM]` format to report: Kbits, Mbits, KBytes, MBytes

`-i, --interval #` seconds between periodic bandwidth reports

- l, --len #[KM] length of buffer to read or write (default 8 KB)
- m, --print\_mss print TCP maximum segment size (MTU - TCP/IP header)
- o, --output <filename> output the report or error message to this specified file
- p, --port # server port to listen on/connect to
- u, --udp use UDP rather than TCP
- w, --window #[KM] TCP window size (socket buffer size)
- B, --bind <host> bind to <host>, an interface or multicast address
- C, --compatibility for use with older versions does not sent extra msgs
- M, --mss # set TCP maximum segment size (MTU - 40 bytes)
- N, --nodelay set TCP no delay, disabling Nagle's Algorithm
- V, --IPv6Version Set the domain to IPv6

#### Options for Servers only

- s, --server run in server mode
- U, --single\_udp run in single threaded UDP mode
- D, --daemon run the server as a daemon

#### Options for Clients only

- b, --bandwidth #[KM] for UDP, bandwidth to send at in bits/sec (default 1 Mbit/sec, implies -u)
- c, --client <host> run in client mode, connecting to <host>
- d, --dualtest Do a bidirectional test simultaneously
- n, --num #[KM] number of bytes to transmit (instead of -t)
- r, --tradeoff Do a bidirectional test individually
- t, --time # time in seconds to transmit for (default 10 secs)
- F, --fileinput <name> input the data to be transmitted from a file
- I, --stdin input the data to be transmitted from stdin
- L, --listenport # port to receive bidirectional tests back on
- P, --parallel # number of parallel client threads to run
- T, --ttl # time-to-live, for multicast (default 1)
- Z, --linux-congestion <algo> set TCP congestion control algorithm (Linux only)

#### Miscellaneous Options

- x, --reportexclude [CDMSV] exclude C(connection) D(data) M(multicast) S(settings) V(server) reports
- y, --reportstyle C report as a Comma-Separated Values
- h, --help print this message and quit
- v, --version print version information and quit

For instance, if the lperf server is to listen to UDP packets on port 319, then you must use the following options:

```
-u -p 319
```

3. In the **Options** text box, type the options you need and then click the **Start Server** button. The server must be started before triggering traffic from an lperf client.

After the server is started, you can test the connection from an lperf client by supplying the hostname as a parameter.

Example Results:

#### EXAMPLE

```
-----  
Server listening on TCP port 5001  
TCP window size: 85.3 KByte (default)  
-----
```

```
[ 4] local 10.10.1.200 port 5001 connected with 10.2.6.228 port 58665  
[ ID] Interval Transfer Bandwidth  
[ 4] 0.0-10.1 sec 112 MBytes 93.2 Mbits/sec
```

# Troubleshooting

Learn how to deal with troubleshooting issues that you might have encountered when using the Exinda Appliances. Much of the information here is also available elsewhere in the help, but is gathered here as a central location for accessing information about how to troubleshoot issues.

## 0.25 Diagnostics

Learn the various diagnostics tools available on your Exinda Appliance, You can use these tools to help troubleshoot issues you might encounter.

|                                                                 |     |
|-----------------------------------------------------------------|-----|
| 0.25.1 Diagnostics Files .....                                  | 715 |
| 0.25.2 Acceleration Diagnostics .....                           | 716 |
| 0.25.3 Monitor .....                                            | 719 |
| 0.25.4 NIC Diagnostics .....                                    | 720 |
| 0.25.5 Optimizer Diagnostics .....                              | 721 |
| 0.25.6 RAID Diagnostics .....                                   | 722 |
| 0.25.7 TCP Dump .....                                           | 723 |
| 0.25.8 View the status of an alert .....                        | 725 |
| 0.25.9 View the status of the community .....                   | 726 |
| 0.25.10 Open a case with Exinda Networks Support Services ..... | 727 |

### 0.25.1 Diagnostics Files

Diagnostics files contain system state information and can aid in troubleshooting. Diagnostics files may be requested by Exinda TAC and can be generated and downloaded using the form below.

| Diagnostics Files        |                                                                                                                                       |                              |               |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------|------------------------------|---------------|
| <input type="checkbox"/> | File Name                                                                                                                             | Timestamp                    | File Size     |
| <input type="checkbox"/> |  <a href="#">sysdump-ex240-20091120-152626.tqz</a> | Fri Nov 20 15:26:38 EST 2009 | 2498591 bytes |

System snapshots are automatically generated when a process fails. If the 'Auto Support Notifications' option is enabled, they are automatically sent to Exinda TAC for further troubleshooting.

| System Snapshot Files     |           |           |           |
|---------------------------|-----------|-----------|-----------|
| <input type="checkbox"/>  | File Name | Timestamp | File Size |
| No System Snapshot Files. |           |           |           |

| Auto Support               |                                            |
|----------------------------|--------------------------------------------|
| Auto Support Notifications | <input checked="" type="checkbox"/> Enable |

#### NOTE

Valid SMTP and DNS settings are required for diagnostics to be sent to Exinda TAC.

## 0.25.2 Acceleration Diagnostics

Acceleration diagnostics aid in troubleshooting TCP Acceleration, SMB Acceleration and WAN Memory issues by displaying the current configuration for those areas.

- » The TCP Acceleration diagnostics display the current TCP configuration settings as well as the number of new and concurrent accelerated connections and reduction statistics.
- » The SMB Acceleration diagnostics display the current SMB configuration settings. If SMB signed connections are present, the total number of signed connections is also displayed.
- » The WAN memory Acceleration diagnostics display the current configuration settings as well as reduction statistics for the individual hosts.

### Related Topics

- » [Viewing TCP Acceleration Configuration and Statistics](#)
- » [Viewing WAN Configuration and Statistics](#)
- » [Viewing SMB Acceleration Configuration and Statistics](#)

### Viewing SMB Acceleration Configuration and Statistics

Use the following instructions to view the SMB acceleration configuration and the current statistics.

1. Go to **Configuration > System > Diagnostics > Acceleration**.
2. From the **Module** drop-down, select **SMB Acceleration**. The configuration settings for SMB1 and SMB2 appear.

## EXAMPLE

The screenshot shows a web interface with a navigation bar containing tabs for System, Diagnostics, TCP Dump, Community, Acceleration, and Monitor. The 'Acceleration' tab is selected. Below the navigation bar, there is a 'Module:' dropdown menu set to 'SMB Acceleration'. A link 'View SMB Acceleration Logs' is visible. The main content area displays the following configuration and statistics:

```
SMB1 Configuration
  Enabled:          yes
  Read-ahead:      yes
  Write-behind:    yes
  Meta-data cache: yes
  Data to prefetch: 1MB

SMB1 Connections
  Concurrent:                0
  Concurrent (signed, unhandled): 0
  Concurrent (signed, handled):  0
  Total Signed (Bypassed):      0
  Total Signed (Handled):       0
  Total Signed (Unhandled):     0

SMB1 Connections compatibility (pre v6.3.0)
  Concurrent: 0
  Signed:    0

SMB2 Configuration
  Enabled:          yes

SMB2 Connections
  Concurrent: 0
  Signed:    0

SMB Signing
  Enabled:          no

No signed servers detected

Windows Authentication Credentials
  Domain          Username          Password Enabled
-----
                  <no entries>
```

The connections statistics are grouped into two categories:

- » **Concurrent** — All signed connections from the file sharing servers that are currently connected.
- » **Total Signed** — All signed connections since the SMB Acceleration service was last started, including those recorded as Concurrent.

As signed connections are processed, there are three possible results:

- » **Bypassed** — The number of connections that bypass acceleration because the first time an attempt to validate the domain credentials failed, which resulted in the connection being identified as signed, but is not accelerated. All subsequent attempts to validate credentials of a signed connection against the IP address of the server are marked as Unhandled.
- » **Handled** — The number of connections that are known to be signed and accelerated.

» Unhandled — The number of connections that, following a bypass state, had subsequent attempts to validate credentials of a signed connection against the IP address of the server.

Related topic

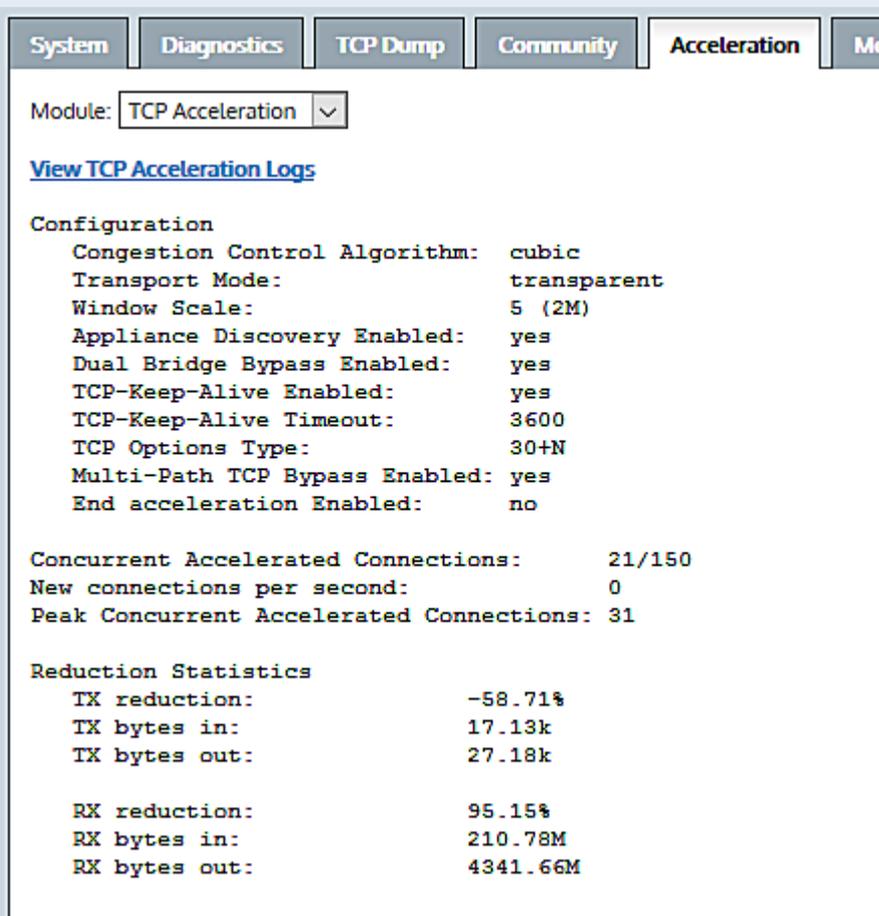
Troubleshoot issues with SMB file acceleration

## Viewing TCP Acceleration Configuration and Statistics

Use the following instructions to view the TCP acceleration configuration and current statistics.

1. Go to **Configuration > System > Diagnostics > Acceleration**.
2. From the **Module** drop-down, select **TCP Acceleration**. The configuration settings and statistics for TCP acceleration appear.

### EXAMPLE



The screenshot shows a web interface with a navigation bar containing tabs for System, Diagnostics, TCP Dump, Community, Acceleration, and Mo. The 'Acceleration' tab is selected, and a dropdown menu shows 'TCP Acceleration' as the chosen module. Below the dropdown is a link for 'View TCP Acceleration Logs'. The main content area displays configuration and statistics in a monospaced font.

```
Configuration
Congestion Control Algorithm: cubic
Transport Mode: transparent
Window Scale: 5 (2M)
Appliance Discovery Enabled: yes
Dual Bridge Bypass Enabled: yes
TCP-Keep-Alive Enabled: yes
TCP-Keep-Alive Timeout: 3600
TCP Options Type: 30+N
Multi-Path TCP Bypass Enabled: yes
End acceleration Enabled: no

Concurrent Accelerated Connections: 21/150
New connections per second: 0
Peak Concurrent Accelerated Connections: 31

Reduction Statistics
TX reduction: -58.71%
TX bytes in: 17.13k
TX bytes out: 27.18k

RX reduction: 95.15%
RX bytes in: 210.78M
RX bytes out: 4341.66M
```

## Viewing WAN Configuration and Statistics

Use the following instructions to view the WAN configuration and operational statistics.

1. Go to **Configuration > System > Diagnostics > Acceleration**.
2. From the **Module** drop-down, select **WAN Memory**. The configuration settings for WAN memory acceleration appear.

## EXAMPLE

The screenshot shows the 'Monitor' tab in the Exinda Network Orchestrator Web UI. The 'Module' dropdown is set to 'WAN Memory'. Below the navigation tabs, there is a link to 'View WAN Memory Logs'. The main content area displays configuration and statistics for WAN Memory.

```
Module: WAN Memory

View WAN Memory Logs

Configuration
Reduction LZ compression on:          yes
Reduction small matching on:         no
Reduction small matcher always list: LotusNotes
Persistence enabled:                 yes
HA cache sync enabled in cluster mode: yes

Statistics
Disk available:                       9.92G
Disk used:                            198.94M (1.958%)
Persistence active:                   yes
Persistence loading:                  no
Persistence clear pending:            no

TX reduction:                         -58.93%
TX bytes in:                          35.32k
TX bytes out:                         56.13k
TX small matcher bytes in:            0
TX small matcher bytes out:           0

RX reduction:                         97.29%
RX bytes in:                          243.28M
RX bytes out:                         8984.57M

Peer 791777123456 state
Status:                               ONLINE
Last status change:                   2016/05/05 14:43:43 (43m 17s)
IP address:                           192.168.5.20
Connections:                          0
Disk cache status:                    Active
Version info:                         Ok/28 (local:26-28 remote:26-28)

TX reduction:                         -58.93%
TX bytes in:                          35.32k
TX bytes out:                         56.13k

RX reduction:                         97.29%
RX bytes in:                          243.28M
RX bytes out:                         8984.57M
```

### 0.25.3 Monitor

The monitor diagnostics display the current monitor settings and the status of monitor and collector processes.

#### NOTE

To configure Monitor settings, navigate to **Configuration > System > Setup > Monitoring** on the Web UI, advanced mode.

```

Table size           : 50
Chart size           : 10
Realtime Window      : 10
Graphing             : flash
Detailed Monitoring  : yes
Ignore Internal-to-Internal : yes

Layer7 Monitoring    :
  Enabled            : yes
  Bittorrent Sensitivity : High
  Bittorrent Sensitivity : High
  EDonky Sensitivity   : Med
  Skype Sensitivity   : High

Host Resolution      :
  Order : DNS Rank : 2
  Order : IP Rank  : 4
  Order : Netbios Rank : 3
  Order : Network_Object Rank : 1

Monitor Status      : OK

Collector Status    : OK
Current Timestamp    : 1287546720

```

## 0.25.4 NIC Diagnostics

The NIC diagnostics page can help when troubleshooting network delay issues. NIC errors, collisions and discards indicate a negotiation problem, which can lead to dropped packets and network delay. It is recommended that negotiation issues are addressed immediately.

The first lines show a summary of installed network adapters. Detailed information is available from the CLI "show diag" command.

### NOTE

To configure NIC settings, navigate to **Configuration > System > Network > NICs** on the Web UI, advanced mode.

Slot 1: PEG2BPI-SD, 2 ports, 1G/RJ-45/1000BASE-T, 1-tx/rx queue  
Slot 2: Empty

Interface br10 state

Admin up: yes  
Link up: yes  
IP address:  
Netmask:  
Speed: N/A  
Duplex: N/A  
Interface type: ethernet  
Interface source: bridge  
MTU: 1500  
HW address: 00:E0:ED:13:73:C2  
Comment:

RX bytes: 37940508  
RX packets: 514502  
RX mcast packets: 514502  
RX discards: 0  
RX errors: 0  
RX overruns: 0  
RX frame: 0

TX bytes: 0  
TX packets: 0  
TX discards: 0  
TX errors: 0  
TX overruns: 0  
TX carrier: 0  
TX collisions: 0

## 0.25.5 Optimizer Diagnostics

The optimizer diagnostics display the current optimizer status and the optimizer configuration.

```

Optimizer Status: Started
Restart Required: no
Show VC Totals: no
Global QOS: no

```

Optimizer Configuration:

```

Chain PREROUTING (policy ACCEPT 9538K packets, 3714M bytes)
pkts bytes target prot opt in out source destination
76M 36G ACCEPT icmp -- * * 0.0.0.0/0 0.0.0.0/0 NET match ignore
20M 16G BRIDGE_PORT all -- br+ * 0.0.0.0/0 0.0.0.0/0

Chain INPUT (policy ACCEPT 13M packets, 13G bytes)
pkts bytes target prot opt in out source destination

Chain FORWARD (policy ACCEPT 83M packets, 39G bytes)
pkts bytes target prot opt in out source destination

Chain OUTPUT (policy ACCEPT 8366K packets, 5121M bytes)
pkts bytes target prot opt in out source destination
1336K 422M ACCEPT all -- * * 0.0.0.0/0 0.0.0.0/0 NET match ignore
0 0 LOCAL_OUT all -- * br+ 0.0.0.0/0 0.0.0.0/0 DIR match inbound
1715K 1183M LOCAL_OUT all -- * * 0.0.0.0/0 0.0.0.0/0 AA match accel DIR match inbound

Chain POSTROUTING (policy ACCEPT 93M packets, 45G bytes)
pkts bytes target prot opt in out source destination

Chain ACTION (1 references)
pkts bytes target prot opt in out source destination
8583K 12G AA all -- * * 0.0.0.0/0 0.0.0.0/0 socket transparent AA target
872K 779M AA all -- * * 0.0.0.0/0 0.0.0.0/0 EXPOLICY match accel AA target port 9998
0 0 COMPRESS_OLD all -- * * 0.0.0.0/0 0.0.0.0/0 EXPOLICY match compress COMPRESS algorithm 48 set-protocol 138

Chain BRIDGE_PORT (1 references)
pkts bytes target prot opt in out source destination
20M 16G DIR_MARK all -- * * 0.0.0.0/0 0.0.0.0/0
20M 16G HA all -- * * 0.0.0.0/0 0.0.0.0/0
20M 16G SETAPP all -- * * 0.0.0.0/0 0.0.0.0/0 SETAPP
6581K 3290M UNACCEL all -- * * 0.0.0.0/0 0.0.0.0/0 DIR match inbound
73440 6289K ACCEPT icmp -- * * 0.0.0.0/0 0.0.0.0/0 NET match ignore
17M 15G MON all -- * * 0.0.0.0/0 0.0.0.0/0
13M 13G ACTION all -- * * 0.0.0.0/0 0.0.0.0/0 DIR match outbound

Chain DIR_MARK (1 references)
pkts bytes target prot opt in out source destination
6581K 3290M MARK all -- * * 0.0.0.0/0 0.0.0.0/0 PHYSDEV match --physdev-in eth11 MARK xset 0x1001/0xffffffff
13M 13G MARK all -- * * 0.0.0.0/0 0.0.0.0/0 PHYSDEV match --physdev-in eth10 MARK xset 0x1000/0xffffffff

```

## 0.25.6 RAID Diagnostics

The RAID diagnostics page is available on models that support Redundant Storage. A summary of the logical volume status is shown as well as details for RAID adapters, logical volumes and physical drives.

```

Adapter: 0 Logical: 0 Size: 1429248MB State: Optimal
Adapter: 0
  Model:          PERC 6/i Integrated
  Serial:         1122334455667788
  Firmware:       6.2.0-0013
  Host Interface: PCIE
  Supported Drives: SAS, SATA
  Levels:         RAID0, RAID1, RAID5, RAID6, RAID10, RAID50, RAID60
  Memory:         Present, 256MB
  Battery:        Yes
  Alarm:          Disabled
  Current Time:   3:53:4 3/29, 2011
Logical Drive: 0
  Adapter:        0
  Size:           1429248MB
  Stripe:         64kB
  Raid Level:     Primary-1, Secondary-3, RAID Level Qualifier-0
  Drives:         2
  Span Depth:     3
  Cache Policy:   WriteBack, ReadAheadNone, Direct, No Write Cache if Bad BBU
  State:          Optimal
Drive: 0
  Adapter:        0
  Slot:           0
  Type:           SAS
  Inquiry:        SEAGATE ST3500414SS      KS679WJ01HND
  Firmware:       Online
  Raw Size:       476940MB [0x3a386030 Sectors]
  Media Errors:   0
  Other Errors:   0
  Predictive Errors: 0
  Sequence:       2
Drive: 1
  Adapter:        0
  Slot:           1
  Type:           SAS
  Inquiry:        SEAGATE ST3500414SS      KS679WJ0275D
  Firmware:       Online
  Raw Size:       476940MB [0x3a386030 Sectors]
  Media Errors:   0
  Other Errors:   0
  Predictive Errors: 0
  Sequence:       2
Drive: 2
  Adapter:        0
  Slot:           2
  Type:           SAS
  Inquiry:        SEAGATE ST3500414SS      KS679WJ033KN
  Firmware:       Online
  Raw Size:       476940MB [0x3a386030 Sectors]

```

## 0.25.7 TCP Dump

A TCP Dump captures packets being transmitted or received from the specified interfaces and can assist in troubleshooting. A TCP Dump may be requested by Exinda TAC.

[Run a TCP Dump from the Exinda appliance](#)

Click **Configuration > Diagnostics > TCP Dump**.

Make the following selections and then click Generate TCP Dumps:

|           |                                                                                                                                                                                                               |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface | Select an interface to run the TCP dump on. Select ALL to capture packets on all (link up) interfaces. Note When ALL is selected for the Interface, only those interfaces which are link up will be included. |
| Timeout   | Select the amount of time for which the TCP Dump will run.                                                                                                                                                    |
| Filter    | Set a filter if required. Refer to the Common User Case examples below for specific filters to use in common circumstances.                                                                                   |
| Status    | Shows the status of a running TCP Dump                                                                                                                                                                        |

### Common Use Cases

The following examples provide the syntax to enter in the Filter field to gather data from a particular source.

To collect traffic to/from a single host

host <IP address>

Example: host 1.2.3.4

To collect traffic from a single host who is the source of the traffic

src <IP address>

Example: src 1.2.3.4

To collect traffic from a single host who is the destination for the traffic

dst <IP address>

Example: dst 1.2.3.4

To collect traffic between two hosts

host <IP address 1> and host <IP address 2>

Example host 1.2.3.4 and host 5.6.7.8

To collect traffic to / from a subnet

net <IP subnet>

Example: net 1.2.3.0/24

To collect traffic between two subnets

src net <IP subnet> and dst net <IP subnet>

Example: src net 1.2.3.0/24 and dst net 1.2.4.0/24

## Send a TCP Dump to Exinda TAC

Saved TCP Dumps can then be downloaded and/or emailed to Exinda TAC using the form below.

| TCP Dump Files           |                                                                                                                                                |                              |                 |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|-----------------|
| <input type="checkbox"/> | File Name                                                                                                                                      | Timestamp                    | File Size       |
| <input type="checkbox"/> |  <a href="#">capture-weber-monitor-20150220-154907.tar.gz</a> | Fri Feb 20 15:49:07 EST 2015 | 308224 bytes    |
| <input type="checkbox"/> |  <a href="#">capture-weber-monitor-20150213-104642.tar.gz</a> | Fri Feb 13 10:46:43 EST 2015 | 3087354 bytes   |
| <input type="checkbox"/> |  <a href="#">capture-weber-monitor-20141217-162605.tar.gz</a> | Wed Dec 17 16:26:19 EST 2014 | 224519218 bytes |
| <input type="checkbox"/> |  <a href="#">capture-weber-monitor-20141217-133350.tar.gz</a> | Wed Dec 17 13:33:53 EST 2014 | 31631085 bytes  |
| <input type="checkbox"/> |  <a href="#">capture-weber-monitor-20141217-133348.tar.gz</a> | Wed Dec 17 13:33:50 EST 2014 | 31631085 bytes  |

For more information about TCP dump filters, refer to <https://danielmiessler.com/study/tcpdump/#common>.

### 0.25.8 View the status of an alert

System alerts notify you of any system issues that may require further attention and troubleshooting. If a system alert is raised the system health status is set to 'Warning' and an email alert is sent.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Configuration > System > Diagnostics**, and switch to the **System** tab. Anything that has generated alerts display the last time an alert was triggered, and the total number of alerts that have been sent.
5. To view the alert that has triggered the warning, click the alarm name. Use the information in this alert to help troubleshooting the issue.
6. To remove the history for an alert, click **Reset**. The system health status is returned to OK.

| Alert Name       | Description                                                                                                                                                                                                                   |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Utilization  | Alert raised when the CPU utilization threshold is reached. The trigger and clear thresholds can be altered. The defaults are 95% and 80% busy respectively.                                                                  |
| System Disk Full | Alert raised when the used disk space threshold is reached. The trigger and clear thresholds can be altered. The defaults are 7% and 10% free respectively.                                                                   |
| Memory Paging    | Alert for memory use and paging. This means that the data in RAM is swapped to disk. Excessive paging alerts could indicate a system that is running low on RAM resources. Check RAM & SWAP graphs under Monitoring > System. |
| Bridge Link      | Alert raised when one of the links of an enabled bridge is down.                                                                                                                                                              |
| Bridge Direction | Alert raised when the appliance cabling is incorrect. In most cases, it indicates the Exinda WAN interface has been incorrectly plugged into the LAN and vice versa.                                                          |

| Alert Name                           | Description                                                                                                                                                                 |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Negotiation                     | Alert raised when the speed/duplex on an interface is set to auto, but it is negotiating at half duplex and/or 10Mbps.                                                      |
| NIC Problems                         | Alert raised when errors are present on the interfaces.                                                                                                                     |
| NIC Collisions                       | Alert raised when collisions are present on the interfaces. The trigger and clear thresholds can be altered. The defaults are 20 and 1 per 30 sec respectively.             |
| NIC Dropped packets                  | Alert raised when dropped packets are present on the interfaces.                                                                                                            |
| SMB signed connections               | Alert raised when SMB signed connections are present.                                                                                                                       |
| Redundant Power                      | Alert raised when one of the power supplies fails (only available on platforms with power redundancy).                                                                      |
| Redundant Storage                    | Alert raised when one of the hard disks fails (only available on platforms with storage redundancy).                                                                        |
| Max Accelerated Connections Exceeded | Alert raised when the number of accelerated connections exceeds the licensed limit. Connections over the licensed limit pass through the appliance and are not accelerated. |
| Asymmetric Route Detection           | Alert raised when traffic from a single connection comes in to the network through one interface or node, and goes out through another interface or node.                   |
| MAPI Encrypted Connections           | Alert raised when encrypted MAPI traffic to a Microsoft Exchange server is detected on an Exinda Appliance. Encrypted MAPI traffic cannot be accelerated.                   |

### 0.25.9 View the status of the community

Display the state of the community and details of the individual hosts that have joined.

Click **Configuration > System > Diagnostics > Community**.

The status of the community is displayed, along with the details of all nodes in the community.

```
State:           Joined
Enabled:         true
Network Forwarding: true
Community Group: 10
```

#### Global Settings

#### Nodes

```
Host ID:         00e0ed13e792
IP Address:      172.16.1.240
Lost State:      found
Last Contact:    N/A
Hostname:        ex240
Version:         5.5.0.12115

Host ID:         0060e0e1c49c
IP Address:      172.16.101.3
Lost State:      found
Last Contact:    2009/12/02 14:29:37 (18s ago)
Hostname:        jib-exinda
Version:         5.5.0.12035
```

## 0.25.10 Open a case with Exinda Networks Support Services

If you are experiencing a problem or have a question about the Exinda Appliance, submit a ticket to Exinda Support.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **System > Diagnostics > Log a Case**.
6. Complete the fields in the form.
7. Ensure a brief summary of the problem or question is included in the **Subject** field.
8. Provide a detailed description of the question or the problem you are experiencing in the **Description** field.
9. Select whether to attach a diagnostics file or monitoring report to the case.
10. Click **Log Case**.

After the case is submitted, a confirmation message containing a case number is sent to the email address identified in the case.

## 0.26 Log Files

Learn about the various log files stored on an Exinda Appliance, and how you can use these logs in your efforts to troubleshoot issues that you might encounter.

---

|                                           |     |
|-------------------------------------------|-----|
| 0.26.1 Viewing System Log Files .....     | 728 |
| 0.26.2 Live Log .....                     | 728 |
| 0.26.3 Tail Log .....                     | 728 |
| 0.26.4 System Logging Configuration ..... | 729 |

---

## 0.26.1 Viewing System Log Files

The View Log Files page allows you to view the system log files and filter out various log messages. Log files provide an inside into the Exinda appliance's operation and aid in troubleshooting.

The following can be used to filter for particular messages:

- » WAN memory — `wmd`
- » TCP acceleration — `tcpad`
- » SMB acceleration — `smbad`
- » Community — `communityd`

To filter and navigate within the log file

1. Go to **Configuration > System > Logging > View**.

2. Select the log file to view. By default, the **Current Log** is displayed. The Exinda appliance periodically archives log files. These archived log files can also be viewed by selecting them from the Logfile list.

3. To filter the contents of the log file, type the criteria to filter by and click **Apply**. The following are examples of common filters that reduce the reported log lines to a single type:

- WAN memory — `wmd`
- TCP acceleration — `tcpad`
- SMB acceleration — `smbad`
- Community — `communityd`

4. If there are multiple pages of log entries, to navigate to a specific page, type the page number in the **Go to Page** field and click **Go**.

## 0.26.2 Live Log

The Live Log page allows you to view new entries to the System Log in real-time.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Logging** and switch to the **Live Log** tab.

### NOTE

A dot/period (.) character is displayed after a few seconds of inactivity to indicate the Live Log is still active.

## 0.26.3 Tail Log

The Tail Log page allows you to view the most recent entries in the system log file.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.

5. Click **Configuration > System > Logging** and switch to the **Tail Log** tab.
6. Configure how many lines to view and in which order to display the log entries.

View Last:  Lines  View Log Order:  ▼

7. To refresh this page and ensure any new log entries since the list time this page was refreshed are displayed, click **Go**.

## 0.26.4 System Logging Configuration

The System Logging Configuration page allows you customize various aspect of System Logging, including exporting to remote syslog servers.

In this area of the Exinda Web UI you can:

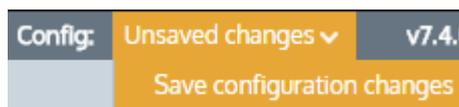
- » [Configure the appliance log files](#)
- » [Add a remote syslog server](#)
- » [Remove a remote syslog server](#)

### Configure the appliance log files

The System Logging Configuration page allows you customize various aspect of System Logging, including exporting to remote syslog servers.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Logging** and switch to the **Setup** tab.
6. Specify the format log files should be saved in. The Standard form is usually sufficient, however some external log file parsers may prefer the log file in WELF format.
7. Select the severity level of log entries that should be saved. Any log entry with this severity level or lower will be saved to the System Log file.
8. Select when the logs are rotated. To force System Log rotation immediately, click **Force Rotation Now**.
9. Specify how many log files should be kept before they are permanently removed from the Exinda appliance.
10. Click **Apply Changes**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.

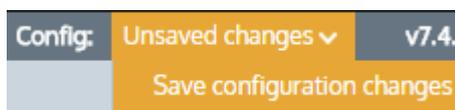


### Add a remote syslog server

Add remote syslog servers to the Exinda appliance, allowing you to forward system log entries at a defined severity level to one or more remote syslog servers.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Logging** and switch to the **Setup** tab.
6. In the Add New Remote Sink area, type the Hostname or IPv4 address of the remote syslog server. IPv6 addresses are not supported for remote sinks.
7. Select the severity level of log entries that are sent to the remote syslog server. Any log entry with this severity level or lower is sent.
8. Click **Add New Remote Sink**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.

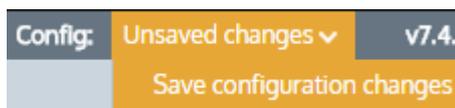


### Remove a remote syslog server

To stop forwarding system log entries to a remote syslog server, remove the server from the Exinda appliance.

1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
5. Click **Configuration > System > Logging** and switch to the **Setup** tab.
6. Select the server from the Remote Log Sinks list, and select **Remove Selected**.
7. Click **Add New Remote Sink**.

To save the changes to the configuration file, in the status bar click the Unsaved changes menu and select Save configuration changes.



### Remove Events from the Appliance System Log

The BMC processor keeps a log of systems events including power status, power redundancy, chassis intrusion. The following command can be used to periodically flush these events to the appliances system log.

```
(config) # ipmi sel enable
```

## 0.27 Troubleshoot problems with MAPI acceleration

If you are experiencing issues with MAPI acceleration, these troubleshooting topics may help resolve the issue.

---

|                                                            |     |
|------------------------------------------------------------|-----|
| 0.27.1 Outlook cannot connect to the Exchange server ..... | 731 |
|------------------------------------------------------------|-----|

---

---

|                                                                                                |     |
|------------------------------------------------------------------------------------------------|-----|
| 0.27.2 Outlook slow to send or receive emails .....                                            | 731 |
| 0.27.3 Decrease in acceleration of MAPI traffic .....                                          | 731 |
| 0.27.4 Reduction ratio for MAPI is different between Client-side and Server-side Exindas ..... | 732 |

---

## 0.27.1 Outlook cannot connect to the Exchange server

### Problem

Microsoft Outlook continuously tries to connect to the Microsoft Exchange server.

### Resolution

Verify that the Exinda is causing the connection issue by putting the appliance in to bypass mode.

1. If Microsoft Outlook continues to be unable to connect to the Exchange server, the Exinda Appliance is not causing the problem. Troubleshoot other areas of your network to find the problem.
2. If Microsoft Outlook can connect to the Exchange server while the Exinda Appliance is in bypass mode, collect a sysdump and packet captures while attempting to connect Microsoft Outlook to the Exchange server, and contact Exinda Support.

## 0.27.2 Outlook slow to send or receive emails

### Problem

When trying to send and receive emails, Microsoft Outlook takes a long time to complete the task or is unresponsive.

### Resolution

Verify that the Exinda is causing the slowness by putting the appliance into bypass mode.

1. If Microsoft Outlook continues to be slow, the Exinda Appliance is not causing the slowness. Troubleshoot other areas of your network to find the problem.
2. If Microsoft Outlook performs at an expected speed while the Exinda Appliance is in bypass mode, accelerate the MAPI traffic using only basic header marking by running the following CLI command: `acceleration mapi basic-header-marking-only`

With basic header marking, only the top level RPC header of each message is ignored when the traffic is accelerated. By accelerating with basic header marking only, the performance is improved but there is less reduction in MAPI traffic.

If performance is not improved, collect a sysdump and packet captures while using Microsoft Outlook, and contact Exinda Support.

## 0.27.3 Decrease in acceleration of MAPI traffic

### Problem

The amount of acceleration experienced on MAPI traffic is not as much as has previously been experienced. The appliance may also be sending alerts about encrypted connections.

## Resolution

Encrypted communications or encrypted email and attachments cannot currently be decrypted by the Exinda. Verify that encryption is disabled on all client computers and that the MAPI protocol is not being encrypted. See For more information, refer to [Disable encryption on the Exchange server](#) (page 677). and For more information, refer to [Turn off MAPI encryption in Microsoft Outlook](#) (page 675)..

## 0.27.4 Reduction ratio for MAPI is different between Client-side and Server-side Exindas

### Problem

The Reduction Ratio percentage reported on the client-side Exinda is not the same as the percentage reported on the server-side Exinda.

### Explanation

As the MAPI traffic passes through the client-side and server-side Exinda appliances, the traffic is decompressed and optimized at different points in the transaction. Because of the timing of the decompressing and optimizing, the appliances may report different reduction ratio percentages.

## 0.28 Troubleshoot issues with TCP acceleration

If you are experiencing issues with acceleration, the following are possible troubleshooting options to consider:

- » Ensure that the traffic is processed by the expected policy.

Go to the Real Time Conversations monitor and check the Show Policies option, which groups the traffic by the virtual circuit and policy. Look for the desired traffic in the expected policy. If the traffic is being accelerated by TCP, the background colour will be yellow.

- » If you have a mix of 7.4, 7.0, 6.4.3, and pre-6.4.3 appliances, perhaps the Acceleration TCP Option Mode is not set correctly. Exinda had used option 30 to indicate acceleration but needed to change this when option 30 was assigned to indicate multi-path TCP. A number of choices were added to ensure compatibility with earlier appliances. Ensure that you believe your choice is correct for your situation or choose another selection.
- » If you have the Multi-Path TCP Acceleration Bypass setting enabled and Acceleration TCP Option Mode is set to liberally use option 30, then when option 30 is encountered it will be interpreted as being multi-path TCP rather than Exinda acceleration and thus will not be accelerated. Ensure these settings are set correctly.
- » If you have a backhaul scenario and you have not enabled the Dual Bridge Bypass setting, then acceleration will not work properly when the SYN from the client is not processed on the same bridge as the SYN/ACK from the server.

You can also attempt to diagnose the issue by viewing the system log or the system diagnostics information.

- » The system diagnostics for acceleration can be filtered for SMB acceleration, WAN memory, or TCP acceleration. For more information, refer to [Acceleration Diagnostics](#) (page 716).
- » The system log file can be filtered for SMB acceleration (smbad), WAN memory (wmd), TCP acceleration (tcpad), or community (communityd). For more information, refer to [Viewing System Log Files](#) (page 728).

## 0.29 Troubleshoot issues with SMB file acceleration

If you are experiencing issues with SMB file acceleration, the following are possible troubleshooting options to consider:

- » Ensure that the traffic is being processed by the expected policy.



Go to **Monitor > Real Time Conversations**, and select the **Show Policies** option. This groups the traffic by the virtual circuit and policy. Look for the desired traffic in expected policy. If the traffic is being accelerated by TCP, the background colour is yellow. If the traffic is being processed by CIFS acceleration the CIFS acceleration icon is shown.

» If a client had already established a connection with the server when the SMB acceleration service was restarted, file transfers over that connection cannot take advantage of the acceleration.

There are two options for terminating the connection between the client computer and the server: Restart the client computer or on a MS Windows client computer, navigate to **Control Panel > Administrative Tools > Services**, and restart the Workstation service.

» Any of the [Troubleshoot issues with TCP acceleration](#) considerations may be applicable.

You can also attempt to diagnose the issue by viewing the system log or the system diagnostics information.

» The system diagnostics for acceleration can be filtered for SMB acceleration, WAN memory, or TCP acceleration. For more information, refer to [Acceleration Diagnostics](#) (page 716).

» The system log file can be filtered for SMB acceleration (smbad), WAN memory (wmd), TCP acceleration (tcpad), or community (communityd). For more information, refer to [Viewing System Log Files](#) (page 728).

## 0.30 Troubleshoot issues with Active Directory configuration

If you are experiencing issues with the Active Directory integration, these troubleshooting topics may help resolve the issue.

---

|                                                                                        |     |
|----------------------------------------------------------------------------------------|-----|
| 0.30.1 Exinda Appliance Reboots Every Night .....                                      | 733 |
| 0.30.2 WMI Service is not running .....                                                | 734 |
| 0.30.3 System account showing in traffic reports .....                                 | 734 |
| 0.30.4 No Communication Between the Exinda AD Connector and the Exinda Appliance ..... | 734 |
| 0.30.5 Exinda AD Connector stops running .....                                         | 734 |
| 0.30.6 Excluded Users Still Appear on the Exinda Appliance .....                       | 735 |
| 0.30.7 Changes to the Exinda Active Directory Controller have no effect .....          | 735 |
| 0.30.8 The IP addresses are not being mapped to the AD users and groups .....          | 736 |

---

### 0.30.1 Exinda Appliance Reboots Every Night

#### Problem

When multiple installations of the Exinda AD Connector have the **Send Active Directory user and group information to Exinda appliance(s) at startup** option selected, the Exinda Appliance can become overwhelmed with duplicate data from the connectors, which can cause the appliance to shut down.

#### Solution

1. On each instance of the Exinda AD Connector, check whether the **Send Active Directory user and group information to Exinda appliance(s) at startup** option is selected.
2. If the option is selected on more than one instance, deselect the option on all Exinda AD Connectors.

3. Choose one instance of the Exinda AD Connector, and select the **Send Active Directory user and group information to Exinda appliance(s) at startup** checkbox, and click **OK**.

### 0.30.2 WMI Service is not running

#### Problem

When I try to access the Exinda AD Connector, a message opens that states “The installer has detected that WMI Service is not running. Consult Windows Help files to find information on how to start WMI Service.”.

#### Solution

This message indicates that Windows Management Information (WMI) service is disabled. The Exinda AD Connector will not be able run correctly until the WMI service is started.

To start the WMI service, at a command prompt type the following command: `net start winmgmt`

### 0.30.3 System account showing in traffic reports

#### Problem

When viewing conversations, the IP address and username of an account created for signing SMB traffic is being displayed as generating traffic rather than the actual user generating the traffic.

#### Solution

When SMB signing is configured and enabled, the SMB signing account is the last user account registered as using an IP address, the Exinda AD Connector transfers the SMB signing account as the username that is generating the traffic. To ignore the SMB signing account and report the traffic as being generated by the actual user, configure the Exinda AD Connector to ignore the SMB signing account. For more information, refer to [Exclude specific usernames from reports](#) (page 561).

### 0.30.4 No Communication Between the Exinda AD Connector and the Exinda Appliance

#### Problem

You see one of the following symptoms:

- » A connection cannot be established between the Exinda AD Connector and the Exinda Appliance.
- » The Last Contact status on the **Configuration > System > Network > Active Directory** tab is blank or red.

#### Resolution

1. Ensure your firewall allows incoming and outgoing traffic on the port configured for the Exinda Appliance to communicate with the Exinda AD Connector

### 0.30.5 Exinda AD Connector stops running

#### Problem

At times after restarting the Exinda AD Connector or the Exinda AD service, the Exinda AD Connector does not continue running, and requires constant restarts.

## Solution

To fix this:

1. The Exinda AD Connector requires .NET version 4.0 for it to run successfully on a server other than the Active Directory server. Ensure .NET 4.0 or later is installed on the server running the Exinda AD Connector.
2. If the Active Directory server is running Windows 2003 R2, ensure the Exinda AD Connector is installed directly on the Active Directory server.
3. Review your event logs for .NET Run Time errors, and attempt to resolve those errors. The .NET installation may need to be reinstalled and the .NET 4.0 services and other environmental services such as WMI may need to be updated.

### 0.30.6 Excluded Users Still Appear on the Exinda Appliance

#### Problem

Even though a user name has been added to the Excluded list on the Exinda AD Connector, the username continues to appear associated with traffic on the Exinda Appliance.

#### Solution

1. Verify that the username on the Excluded tab of the Exinda AD Connector matches the username in Active Directory. The username is case sensitive. For example, if the Active Directory has the user `Domain/Test.User`, and the excluded list has the user as `Domain/test.user`, the traffic is not excluded.

#### NOTE

Regardless of the case of usernames in Active Directory, the Exinda Appliance displays the usernames with the first name capitalized and the surname in lower case; for example `Domain/test.user`. Do not use the value in the Exinda Appliance when adding a username to the Excluded list.

2. If the case matches on the usernames, restart the AD Client Service and renumerate the Exinda Appliance.

#### Related Topics

- » [Change the state of the Exinda AD Connector](#)
- » [Request updated user and group information from the Active Directory server.](#)

### 0.30.7 Changes to the Exinda Active Directory Controller have no effect

#### Problem

After making changes to the configuration of the Exinda Active Directory Controller, the information reported on the Exinda Appliance appears to be the same as before the changes.

#### Solution

Restart the AD Client Service and renumerate the Exinda Appliance to ensure the latest configuration is being used.

#### Related Topics

- » [Change the state of the Exinda AD Connector](#)
- » [Request updated user and group information from the Active Directory server](#)

## 0.30.8 The IP addresses are not being mapped to the AD users and groups

### Problem

When integrating the AD client with the Exinda appliance, the IP addresses are not being mapped to the users and groups on the Exinda appliance.

### Solution

Logon auditing must be enabled for IP address to be mapped to the users.

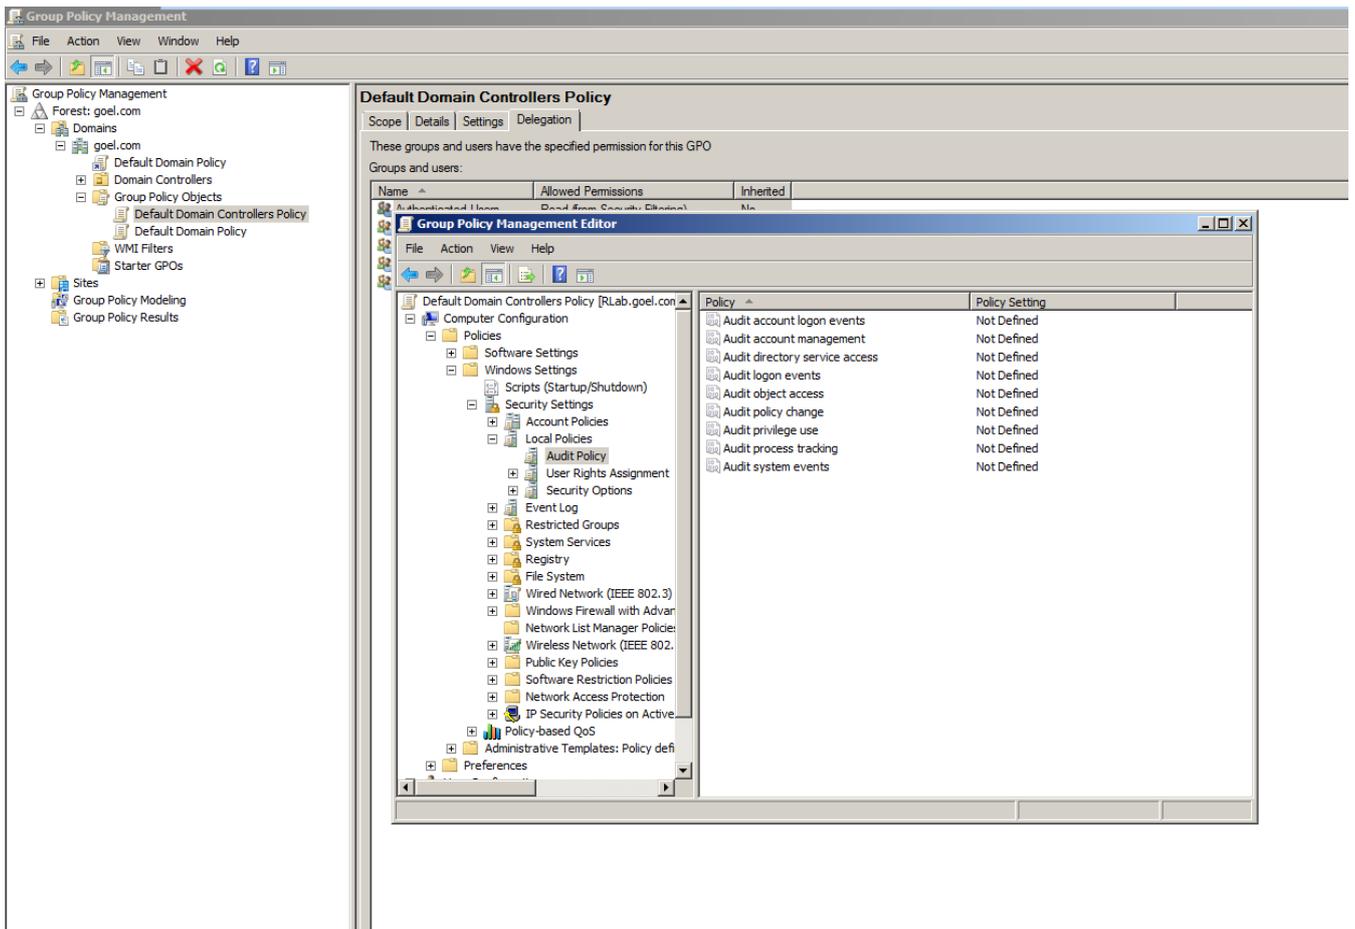
You can investigate by verifying whether the domain controller is logging particular event IDs. If these events are absent then you will need to enable logon auditing.

In the Domain Controller, go to **Event Viewer > Windows Logs > Security Logs**.

- For Windows Server 2008, 2008 R2, 2012, and 2012 R2, you should see Event ID #4624
- For Windows Server 2003, 2003 R2, you should see Event ID #528 and 540.

If the Domain Controller is not logging these events, then you need to enable **logon auditing** on the domain controller and renumerate the AD client on the Exinda appliance.

1. In the Domain Controller, go to **Start menu > Administrative Tools > Group Policy Management Snap-in**.
2. In the Group Policy Management tree, go to your domain, expand the **Group Policy Objects** node, and select **Default Domain Controllers Policy**.



3. Right click on **Default Domain Controllers Policy** and select **Edit** from the context menu.
4. In the **Group Policy Management Editor** dialog box, expand the tree and select **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > Audit Policy**.
5. In the policy list on the right, click on **Audit logon events** and ensure that **Success** is checked.
6. On the Exinda appliance, go to **Configuration > System > Network > Active Directory**.
7. Click the **Renumerate** button.
8. Apply the changes by executing the following command using a CMD console in the Domain Controller: `gpupdate /force`

## 0.31 Troubleshooting Edge Cache

I do not know if Edge Cache is caching any traffic

- » You can determine if any traffic is passing through Edge Cache by looking at the real time conversations monitor available at **Monitor > Real Time > Conversations**. Any traffic that is currently being processed by Edge Cache will have a blue background. This will tell you if Edge Cache is evaluating whether data could be retrieved from its cache or evaluating whether the data should be stored in it. However, it does not indicate whether it is successful in either retrieving or storing the data.
- » You can determine if any traffic has passed through Edge Cache over time by looking at the Edge Cache report available at **Monitor > Optimization > Edge Cache**. The Edge Cache monitor report will show the amount of data over time processed through Edge Cache served on the LAN and the amount of data retrieved from the WAN. It will also

show the requests to Edge Cache and the hits. These two charts will show you if data is going through Edge Cache and if so how many hits and how much data is served from the cache.

- » Check the logs for any errors related to Edge Cache, including DNS not being configured.
- » You should ensure that DNS is configured properly by visiting **Configuration > System > Network > DNS**. For help configuring DNS, see the main user guide.
- » If the Real Time monitor and the Edge Cache monitor indicate that data is not passing through Edge Cache, you should ensure that the Edge Cache process is running. Go to **Configuration > System > Optimization > Services** and ensure that Edge Cache is running. You may want to restart Edge Cache.

I do not know if Edge Cache is caching encrypted traffic

- » All the steps for troubleshooting caching above apply.
- » Try browsing to an encrypted site then look at the real time conversations. The encrypted traffic will be reflected in the application name, such as HTTPS. If processed through Edge Cache, the conversation will have a blue background.
- » Ensure HTTPS caching is enabled and a certificate is configured.
- » If the traffic that you are concerned about is specified in a caching whitelist or blacklist using source or destination domains, then try restarting Edge Cache at **Monitor > Optimization > Services**. If the reverse mapped domains have changed, then the domains that are specified in the whitelists and blacklists may need to be re-resolved and re-reverse mapped. This is done upon Edge Cache startup, when there is a change to the whitelist or blacklist, or when there is a change to the DNS information.

I want to increase the effectiveness of Edge Cache

- » You can restrict the size of the objects that can be cached to more closely match the type of data that you want cached by visiting **Configuration > System > Optimization > Edge Cache**.
- » You can increase the amount of storage available for Edge Cache to use by visiting **Configuration > System > Setup > Storage**.

What factors should I consider before starting Edge Cache on HTTPS for surf traffic?

- » Develop a list of sites that need to be white listed because of security concerns (for example financial sites)
- » Deploy your self-signed certificate throughout the entire network if you do not want users to accept the “false” certificate
- » Independent appliances (for example printer from various manufacturers) use HTTPS to connect to maintenance sites to order cartridges and no human interaction is needed; therefore these sites need to be white listed.
- » Payment terminals use HTTPS to accept payments and there is no option to install a self-signed certificate
- » Some browsers (typically Google Chrome) check more than only the “trusted” self-signed certificate before allowing access to a specific web page

Why are some web pages rendering incorrectly?

If Edge Cache is not rendering layouts as expected, the problem is likely due to the certificate format. Using the Certificate generator of the Exinda appliance lets you export PEM and DER Certificate formats, but some formats require a PKCS12 certificate and these cannot be exported from the Exinda appliance. To correct the problem, use openssl to generate Certificates and import them into the appliance:

1. Use openssl to create your PKCS12 + private key and CSR file. The list of common command lines for openssl to create your Certificates includes:

- Generate a new private key and Certificate Signing Request: `openssl req -out CSR.csr -new -newkey rsa:2048 -nodes -keyout privateKey.key`

- Generate a self-signed certificate (see How to Create and Install an Apache Self Signed Certificate for more info): `openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout privateKey.key -out certificate.crt`
- Generate a certificate signing request (CSR) for an existing private key: `openssl req -out CSR.csr -key privateKey.key -new`
- Generate a certificate signing request based on an existing certificate: `openssl x509 -x509toreq -in certificate.crt -out CSR.csr -signkey privateKey.key`

2. Import the certificate into the Exinda appliance:

- Login to the appliance.
- Click **Configuration > System > Certificates**.
- Type a name for the certificate, then browse to location of the Certificate file and the Private Key file.
- Click **Import**.
- Switch to the All Certificates tab to ensure your import was successful.

3. Assign the PKCS12 certificate to Edge Cache.

- Click **Configuration > System > Optimization > Edge Cache**.
- Under HTTPS Caching, click the checkbox to enable HTTPS content caching, and then select the PKCS certificate you created from the Signing Certificate list.
- Click Apply Changes.

4. Import the PFX/PKCS12 certificate to your own computer. Reload the page that was formatted incorrectly to ensure the new certificate solves the problem.

Are there any open issues I should know about?

The following open issues are known concerns:

- » Facebook does not work with Chrome
- » Facebook works with Safari but still has a pop-up
- » Google apps work with Chrome but they can only use the Google Certificate
- » Google apps work with Safari
- » Firefox still has pop-ups because it uses a different Certificate store
- » Outlook connects to Exinda without popups

## 0.32 Topology troubleshooting

| Problem                                                                                                            | Solution                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| My network traffic is blocked after deploying the Exinda appliance in line.                                        | Ensure you have used the correct cables for your environment. Some environments may require 2x straight Ethernet cables, while others may require 2x cross-over Ethernet cables. |
| My network traffic is blocked after deploying the Exinda appliance in line, after I have booted it up.             | Ensure the speed/duplex settings are correct on both the Exinda appliance and any neighboring equipment.                                                                         |
| I am experiencing significant packet loss after deploying the Exinda appliance in line, after I have booted it up. | See above regarding speed/duplex configuration. Also check Ethernet cables for defects.                                                                                          |

# Exinda Command Line Interface (CLI)

Learn how to use the Exinda Command Line Interface (CLI).

---

---

## 0.33 Using the Command Line Interface

Many of the actions available in the Exinda Web UI can also be executed through the Command Line Interface (CLI).

### TIP

- » Auto complete is available by pressing the tab key after typing the first several letters of a command. Use the tab key to view available options for any of the commands.
- » Use ? at the end of a command to view available options and descriptions.
- » Command history is available by using the up and down arrow keys. Command line editing is available, using the left and right keys to navigate.
- » Use **ctrl-w** to delete from the cursor to start of line.

See the following topics for more information:

---

|                                                   |     |
|---------------------------------------------------|-----|
| 0.33.1 Accessing the Command Line Interface ..... | 740 |
| 0.33.2 CLI Configuration Jumpstart .....          | 741 |
| 0.33.3 Configure command line options .....       | 742 |

---

### 0.33.1 Accessing the Command Line Interface

There are four ways of accessing the Exinda CLI (in order of preference):

1. Secure Shell (SSH) (recommended)
2. Exinda Web UI
3. Telnet
4. Serial Console Interface

Use this tool to connect to the Exinda appliance's Command Line Interface (CLI) from the Web UI. This tool connects to the appliance via the web interface and does not require SSH access.

### Open new fullscreen console



1. On your browser, open the Exinda Web UI ([https://Exinda\\_IP\\_address](https://Exinda_IP_address)).
2. Key-in the **User** and **Password**.
3. Click **Login**.
4. Click **Configuration > System > Tools > Console**.
5. Type the appliance username and password at the prompts. Do one of the following:
  - To enter privileged EXEC (enable) mode, at the prompt run the command: `hostname > enable`  
The `hostname #` prompt appears.
  - To enter configuration (config) mode, at the prompt run the commands: `hostname # configure terminal`  
The `hostname (config) #` prompt appears.

### 0.33.2 CLI Configuration Jumpstart

When you login to the CLI for the first time, you are presented with the option to run the CLI jump-start wizard. This is a guided wizard that helps with the initial configuration of the Exinda appliance.

#### NOTE

Changes are applied immediately after pressing **Enter** at each step. If changing network settings use the serial console or vga/keyboard to access the CLI.

1. `Enable IPv6?` - These questions allow you to enable IPv6 support for the entire system. If your network supports IPv6 then type 'Y', otherwise type 'N'.
2. `Enable IPv6 autoconfig (SLAAC) on eth1 interface?` - If you enable IPv6, you have the option of enabling IPv6 SLAAC autoconfiguration. Type 'Y' if you wish to have an address and netmask automatically configured and your network supports this option.
3. `Use eth0 for management access. Note: This disables br0 (Y/N)?` - Select whether to use eth0 for accessing management functionality.
4. `Use DHCP on eth1 (Y/N)?` - This question is asking if you want to use DHCP for automatically acquiring IP connectivity settings. If you specify 'N' here, you will be prompted to type static IP connectivity settings, such as IP address and netmask, default gateway and DNS servers.
5. `Enable br10 (Y/N)?` and `Use DHCP on br10 (Y/N)?` - These questions allow you to enable bridges and optionally configure an address manually or by using DHCP.

6.br2 IP address and netmask? [192.168.2.254/24] - Configure the IP address and netmask for the bridge.

7.Hostname? - This question is asking you to configure a hostname for the appliance.

8.SMTP server address? - In order to receive system alerts and reports, the Exinda appliance requires an SMTP server be configured so that emails can be sent.

9.An email address for reports and alerts? - If you wish to receive system alerts and reports, type an email address here.

10.Admin password (Enter to leave unchanged) : - This question is asking you if you wish to change the password of the Exinda appliance's 'admin' account. Press 'Enter' to leave the password unchanged or enter a new password and you'll be asked to re-enter the password again to confirm.

11.Do you want to configure the interface speed and duplex settings? (Y/N) ? - Enter 'Y' if you wish to configure interface settings or 'N' to leave them unchanged. If you entered 'Y', these questions will step through each interface on the Exinda appliance and ask for interface speed and duplex settings.

What is the speed of eth1 (auto, 10 or 100):

What is the duplex mode of eth1 (auto, full or half):

What is the speed of eth2 (auto, 10 or 100):

What is the duplex mode of eth2 (auto, full or half):

12.Do you want to change HTTP proxy settings (Y/N) ? - If you enter Y, these questions step through the parameters of the HTTP Proxy setup.

HTTP proxy address (0.0.0.0 to disable)?

HTTP proxy port? [3128]

HTTP proxy authentication type (N)one or (B)asic (N/B)?

Allow insecure (unverified certificate) SSL (Y/N)?

13.Do you want to check for a new license online (Y/N) ? - Enter 'Y' to have the Exinda appliance check for a newer license on the Exinda website (if the Exinda appliance has Internet connectivity). If a newer license is found, you will be asked if you wish to install it. If you enter 'N', you will be prompted for a license key.

14.Do you want to configure optimization policies (Y/N) : - Answering 'Y' here will take you through a text-based version of the Optimizer Wizard.

15.Check for new firmware (Y/N) ? - Answering 'Y' here will make the Exinda appliance check for a newer firmware version on the Exinda website (if the Exinda appliance has Internet connectivity). If a newer firmware image is found, you will be asked if you want to download and install it.

#### NOTE

You can re-run the CLI jump-start wizard at anytime by logging into the CLI (configuration mode) and typing:  
`configuration jump-start`

### 0.33.3 Configure command line options

Configure the command line interface to meet your needs.

1. Use the following command to set the terminal character width and number of lines:

```
hostname (config)# cli session terminal width <number of characters>
```

```
hostname (config)# cli session terminal length <number of lines>
```

2. Auto logout is enabled by default. To change the auto logout time use the following command:

```
hostname (config)# cli default auto-logout <minutes>
```

To disable auto-logout, set the minutes to 0.

3. To enable or disable paging use the following command:

```
hostname (config)# [no] cli default paging enable
```

4. Use the `show cli` command to see current CLI settings.

5. To save the running configuration, type `configuration write`.

# Copyright

All rights reserved. No parts of this work may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems - without the written permission of the publisher.

Products that are referred to in this document may be either trademarks and/or registered trademarks of their respective owners. The publisher and the author make no claim to these trademarks.

While every precaution has been taken in the preparation of this document, the publisher and the author assume no responsibility for errors or omissions, or for damages resulting from the use of information contained in this document or from the use of programs and source code that may accompany it. In no event shall the publisher and the author be liable for any loss of profit or any other commercial damage caused or alleged to have been caused directly or indirectly by this document.

Refer to the following:

---

---

## 0.34 Exinda End User License Agreement (EULA)

**Notice to users:** carefully read the following legal agreement. Use of the software provided with this agreement ("software") constitutes your acceptance of these terms. If you do not agree to the terms of this agreement, promptly return the complete software package (and any other devices delivered with this package) to the dealer from whom you obtained this product for a full refund. If you have any questions concerning this agreement contact [legal@exinda.com](mailto:legal@exinda.com).

**License grant:** the software is licensed, not sold. Upon the valid purchase of a license to the software and except as otherwise specified in an accompanying license summary, invoice or other documents evidencing the purchase of the software license, Exinda grants you a non-exclusive, non-transferable license to use the software during the subscription period on servers connected to a maximum number of user computers not exceeding the number of user computers specified in the packaging accompanying the software or in any supplemental agreements. This license to use the software is conditioned upon your compliance with the terms of this agreement. You agree you will only compile the software into any machine-readable or printed form as necessary to use it in accordance with this license or for backup purposes in support of your use of the software. This license is effective until terminated. You may terminate it at any point by destroying the software together with all copies of the software. Exinda has the option to terminate this agreement if you fail to comply with any term or condition of this agreement. You agree upon such termination to destroy the software together with all copies of the software.

Reverse engineering: you may not reverse engineer, decompile, modify or disassemble the software in whole or in part.

**Copyright:** all title and copyrights in and to the software, and accompanying printed materials are owned by Exinda. The software is protected by copyright laws and international treaty provisions. The software is copyright (c) 2002 Exinda networks PTY LTD, all rights reserved. The software remains the sole and exclusive property of Exinda at all times.

**Limited warranty:** Exinda warrants that for a period of thirty (30) days from the date of shipment from Exinda: (i) the software will be free of defects in workmanship under normal use, and (ii) the software substantially conforms to its published specifications. Except as expressly granted in this agreement the software is provided as is. In no event shall Exinda, or any of its affiliates, subsidiaries or suppliers (each an "Exinda party" and together the "Exinda parties") be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or any other pecuniary loss) arising out of the use of or inability to use this software, even if such Exinda party has been advised of the possibility of such damages. Because some states and jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

**No other warranties:** None of the Exinda parties warrant that the software is error free. Except for the "limited warranty" in section 4 ("limited warranty"), the Exinda parties disclaim all other warranties with respect to the software, either express or implied. Including but not limited to implied warranties of merchantability, fitness for a particular purpose and non-infringement of third party rights. Some jurisdictions do not allow the exclusion of implied warranties or limitation on how long an implied warranty may last, or the exclusion or limitation of incidental or consequential damages, so the above limitations or exclusions may not apply to you. This warranty gives you specific legal rights and you may also have other rights which vary from jurisdiction to jurisdiction.

**Severability:** in the event of invalidity of any provision of this license, the parties agree that such invalidity shall not affect the validity of the remaining portions of this license.

**Applicable law:** this license will be governed by the laws of the State of Victoria, Australia. In the event of any dispute arising out of this agreement the parties hereby agree to submit to the jurisdiction of the courts of the State of Victoria, Australia.

**Entire agreement:** this is the entire agreement between you and Exinda which supersedes any prior agreement or understanding, whether written or oral, relating to the subject matter of this license.

## 0.35 GNU General Public License (GPL)

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc.

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### 0.35.1 Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely

where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

## 0.35.2 TERMS AND CONDITIONS

### Definitions

“This License” refers to version 3 of the GNU General Public License.

“Copyright” also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

“The Program” refers to any copyrightable work licensed under this License. Each licensee is addressed as “you”.

“Licensees” and “recipients” may be individuals or organizations.

To “modify” a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a “modified version” of the earlier work or a work “based on” the earlier work.

A “covered work” means either the unmodified Program or a work based on the Program.

To “propagate” a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To “convey” a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays “Appropriate Legal Notices” to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

### 1. Source Code.

The “source code” for a work means the preferred form of the work for making modifications to it. “Object code” means any non-source form of a work.

A “Standard Interface” means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The “System Libraries” of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A “Major Component”, in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The “Corresponding Source” for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

## 2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

## 3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

## 4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

## 5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a. The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b. The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to “keep intact all notices”.
- c. You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d. If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an “aggregate” if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

## 6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c. Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and non-commercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d. Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e. Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is

expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

## 7. Additional Terms.

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a. Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b. Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c. Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d. Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e. Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f. Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits

relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

## 8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

## 9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

## 10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An “entity transaction” is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party’s predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

## 11. Patents.

A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor’s “contributor version”.

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

## 12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

## 13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special

requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

#### 14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

#### 15. Disclaimer of Warranty.

There is no warranty for the program, to the extent permitted by applicable law. Except when otherwise stated in writing the copyright holders and/or other parties provide the program “as is” without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The entire risk as to the quality and performance of the program is with you. Should the program prove defective, you assume the cost of all necessary servicing, repair or correction.

#### 16. Limitation of Liability.

in no event unless required by applicable law or agreed to in writing will any copyright holder, or any other party who modifies and/or conveys the program as permitted above, be liable to you for damages, including any general, special, incidental or consequential damages arising out of the use or inability to use the program (including but not limited to loss of data or data being rendered inaccurate or losses sustained by you or third parties or a failure of the program to operate with any other programs), even if such holder or other party has been advised of the possibility of such damages.

#### 17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

## 0.36 BSD 2.0

The BSD 2.0 License

Copyright (c) 2009 Kontron America, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- a. Re-distributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- b. Re-distributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- c. Neither the name of Kontron, nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by the copyright holders and contributors "as is" and any express or implied warranties, including, but not limited to, the implied Warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the copyright owner or contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this Software, even if advised of the possibility of such damage.

# Safety and Compliance

## NOTE

This safety and compliance information only applies to 2x61 appliances.

## 0.37 EMC Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case users will be required to correct the interference at their own expense.

See Also:

[Safety Compliance](#)

[Safety Guidelines](#)

## 0.38 Compliances

### 0.38.1 CE

This product has passed the CE test for environmental specifications. Test conditions for passing included the equipment being operated within an industrial enclosure.

In order to protect the product from being damaged by ESD (Electrostatic Discharge) and EMI leakage, we strongly recommend the use of CE-compliant industrial enclosure products.

### 0.38.2 FCC Class A

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment.

This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

## 0.39 Safety Guidelines

Follow these guidelines to ensure general safety:

- » Keep the chassis area clear and dust-free during and after installation.
- » Do not wear loose clothing or jewelry that could get caught in the chassis. Fasten your tie or scarf and roll up your sleeves.
- » Wear safety glasses if you are working under any conditions that might be hazardous to your eyes.
- » Do not perform any action that creates a potential hazard to people or makes the equipment unsafe.
- » Disconnect all power by turning off the power and unplugging the power cord before installing or removing a chassis or working near power supplies
- » Do not work alone if potentially hazardous conditions exist.
- » Never assume that power is disconnected from a circuit; always check the circuit.

### 0.39.1 Lithium Battery Caution

Risk of Explosion if Battery is replaced by an incorrect type. Dispose of used batteries according to the instructions

#### Operating Safety

Electrical equipment generates heat. Ambient air temperature may not be adequate to cool equipment to acceptable operating temperatures without adequate circulation. Be sure that the room in which you choose to operate your system has adequate air circulation.

Ensure that the chassis cover is secure. The chassis design allows cooling air to circulate effectively. An open chassis permits air leaks, which may interrupt and redirect the flow of cooling air from internal components.

Electrostatic discharge (ESD) can damage equipment and impair electrical circuitry. ESD damage occurs when electronic components are improperly handled and can result in complete or intermittent failures. Be sure to follow ESD-prevention procedures when removing and replacing components to avoid these problems.

Wear an ESD-preventive wrist strap, ensuring that it makes good skin contact. If no wrist strap is available, ground yourself by touching the metal part of the chassis.

Periodically check the resistance value of the antistatic strap, which should be between 1 and 10 megohms (Mohms).

# Predefined Applications and Application Groups

This section provides information about the predefined applications and application groups that your Exinda Appliance (s) is pre-configured to manage.

To see a list of predefined applications and application groups navigate to [http://go.gfi.com/?pageid=exos\\_help#cshid=pre\\_apps](http://go.gfi.com/?pageid=exos_help#cshid=pre_apps)

## 0.40 Predefined Applications and Supported L7 Signatures

The protocol of network traffic is identified based on a combination of deep packet inspection (DPI) and behavioral analysis, regardless of whether they are plain text, or use advanced encryption and obfuscation techniques. The following list includes all supported protocols and applications (**in bold**), sub-protocols (*in italic*) and clients (in standard characters).

0-9

1kxun

9ku

360 Mobile Security

A

Acronis

Activesync

ACTNET

AD Replication

Adobe Creative Cloud

Adobe Updates

AdobeConnect

ADrive

After School

AFSVL

Age of Warring Empire

Alibaba

AliExpress

Aimini

AirPort Utility

Amazon Cloud

Amazon Prime Music

Amazon Shop

Android OS Update

Anonymous Proxy

Antivirus Lab

ANts P2P

Apple Talk

Apple Updates

AppleJuice

Ares

*Filepipe*

*KCEasy*

Ariel

Armagetron

Audiko

Audiogalaxy

*Audiogalaxy iOS*

*Audiogalaxy Android*

AURP

AVG Antivirus

AVI

Avira Antivirus

Avira

B

Backweb

Baidu

Battlefield

BattleNet

Bebo

BGP

Biff

bini

Bitdefender

BitTorrent

*BitTorrent-plain*

*BitTorrent-encrypted*

- ABC
- Acquisition
- Anatomic P2P
- Arctic Torrent
- Ares

- Azureus/Vuze
- BTSharp
- BitComet
- BitFlu
- BitLord
- BitPump
- BitRocket
- Bits on Wheels
- BitSpirit
- BitSticks
- BitTornado
- BitTorrent/Mainline
- BitTyrant
- Blizzard Downloader
- BTG
- BtManager
- burst!
- CTorrent
- Deluge
- FlashGet
- Freeloader
- FrostWire
- G3 Torrent
- Gnome BitTorrent
- Halite
- KTorrent
- LimeWire
- MLDonkey
- MonoTorrent
- MooPolice
- Opera
- Pando
- qBittorrent
- QTorrent
- rTorrent
- Rufus

- Shareaza
- Sharktorrent
- SoMud
- SymTorrent
- Tomato Torrent
- TorrentFlux
- TorrentSpy
- Torrent Station
- Torrent Swapper
- TorrentVolve
- Transmission
- Turbo Torrent
- µTorrent
- µTorrent Mac
- WizBit
- Xtorrent
- ZipTorrent

*BitTorrent-UTP*

*BitTorrent-UTP-encrypted*

Blackberry

*Blackberry-encrypted*

*Blackberry-email*

*Blackberry-messenger*

- ORB-client
- BlackBerry Internet Service
- BlackBerry Management Center
- BlackBerry Messenger

*Unknown*

Blackberry Mail

Blackberry Messenger

Bluejeans

BOLT

- Bolt Android

Booking Com

BootPC

Box

Burner

## C

Call of Duty

Cassandra Internode Communication

Cassandra Query Language

CBT

Channel 4oD

Chargen

CIFS

CiscoDiscovery

Citrix

*Citrix-application*

*Citrix-priority*

*Citrix-user*

CloudME

Club Penguin

clussvc

CNN

CNTV

ComodoUnite

*Comodo EasyVPN*

Constant Contact

CORBA

Crashplan Pro

Crime City

Crossfire

CRS

CtiTV

CUDev

CVSpserver

CyberGhost

## D

Dailymotion

dameware

Daytime

dcc

DDM-SSL

Deezer

Demand5

*Demand5-stream*

*Demand5-web*

- Demand5 iOS

Destiny

DFS

DHCP

DHCPv6

Diameter

DICOM

DICOM TLS

DiDi IM

DiDi Voice

DirectConnect

*DirectConnect-hub*

*DirectConnect-peer*

*DirectConnect-adc-peer*

- ApexDc++
- BCDC++
- CZDC
- DC#
- DC++
- DCTC
- Dolda Connect
- Elise
- fulDC
- GtkDC
- LDCC
- LinuxDC++
- Microdc
- Microdc2
- oDC
- Revconnect
- SababaDC
- ShakesPeer

- StrongDC++
- Valknut

#### DirectDownload

*10upload.com*  
*123upload.pl*  
*139pan.com*  
*163pan.com*  
*1-clickshare.com*  
*1clickshare.net*  
*1fichier.com*  
*1-upload.com*  
*1-upload.to*  
*2shared.com*  
*4fastfile.com*  
*4share.ws*  
*4shared.com*  
*4shared.vn*  
*4us.to*  
*adrive.com*  
*alldrives.ge*  
*allshares.ge*  
*annonhost.net*  
*badongo.net*  
*banashare.com*  
*bestsharing.com*  
*bigfilez.com*  
*biggerupload.com*  
*bitshare.com*  
*boosterking.com*  
*cash-file.net*  
*chinamole.com*  
*cobrashare.sk*  
*coolshare.cz*  
*cramit.in*  
*crocko.com*  
*cum.com*  
*czshare.com*

*data.hu*  
*data-loading.com*  
*dataport.cz*  
*datei.to*  
*depositfiles.com*  
*downupload.com*  
*easybytez.com*  
*easy-share.com*  
*edisk.cz*  
*enterupload.com*  
*euroshare.eu*  
*extabit.com*  
*eyvx.com*  
*fastfileshare.com.ar*  
*fast-load.net*  
*fast-share.com*  
*fik1.com*  
*file2upload.net*  
*fileape.com*  
*filearchiv.ru*  
*filearn.com*  
*filebase.to*  
*filecloud.com*  
*filedude.com*  
*filefactory.com*  
*fileflyer.com*  
*filefrog.com*  
*filefront.com*  
*fileho.com*  
*filehook.com*  
*filehost.tv*  
*filejungle.com*  
*filemaze.ws*  
*fileover.net*  
*filepost.com*  
*filepost.ru*  
*filer.net*

*files.mail.ru*  
*files.to*  
*filesafe.to*  
*filesend.net*  
*fileserve.com*  
*filesflash.com*  
*fileshare.in.ua*  
*filesmonster.com*  
*asixfiles.com*  
*badongo.com*  
*filesonic.com*  
*filesonic.in*  
*filestube.com*  
*files-upload.com*  
*file-upload.net*  
*file-upload.to*  
*fileupyours.com*  
*filevo.com*  
*filezup.com*  
*freakshare.net*  
*freespace.by*  
*fshare.vn*  
*fsx.hu*  
*ftp2share.com*  
*gaiafile.com*  
*gamefront.com*  
*gigapeta.com*  
*gigasize.com*  
*gigeshare.com*  
*gotoupload.com*  
*hellshare.com*  
*hitfile.net*  
*hostggg.com*  
*hotfile.com*  
*hulkshare.com*  
*hyperfileshare.com*  
*ifile.it*

*i-filez.com*  
*ifolder.ru*  
*jakfile.com*  
*jumbofiles.com*  
*keepmyfile.com*  
*kewlshare.com*  
*kickload.com*  
*leteckaposta.cz*  
*letitbit.net*  
*livedepot.net*  
*live-share.com*  
*load.to*  
*loadles.in*  
*massmirror.com*  
*mediafire.com*  
*mega.1280.com*  
*megafree.kz*  
*megaporn.com*  
*megarapid.eu*  
*megarotic.com*  
*megashare.com*  
*megashare.vn*  
*megashare.vnn.vn*  
*megashares.com*  
*megaupload.com*  
*midupload.com*  
*missupload.com*  
*mofile.com*  
*mofile.net*  
*movreel.com*  
*multishare.cz*  
*mytempdir.com*  
*netload.in*  
*netuploaded.com*  
*obsiebie.najlepsze.net*  
*oron.com*  
*platinshare.com*

*przeklej.pl*  
*putshare.com*  
*qshare.com*  
*quickshare.cz*  
*quicksharing.com*  
*rapidshare.com*  
*rapidshare.de*  
*rapidshark.pl*  
*rapidupload.com*  
*sanupload.com*  
*savefile.com*  
*scribd.com*  
*sendmefile.com*  
*sendspace.com*  
*sharebase.to*  
*sharebee.com*  
*sharebig.com*  
*sharebigfile.com*  
*sharebigflie.com*  
*sharefile.net*  
*share-online.bz*  
*share-rapid.cz*  
*shrage.com*  
*simpleupload.net*  
*speedshare.org*  
*storage.to*  
*terafiles.net*  
*tunabox.net*  
*turbobit.net*  
*turbobit.ru*  
*turboupload.com*  
*u.115.com*  
*ugotfile.com*  
*uloz.to*  
*ultrashare.de*  
*unibytes.com*  
*unlimit.co.il*

*up.4share.vn*  
*up-file.com*  
*upload.com.ua*  
*uploadbox.com*  
*uploadc.com*  
*uploaded.to*  
*uploader.pl*  
*uploadfloor.com*  
*uploadpower.com*  
*uploadstation.com*  
*uploadyourfiles.de*  
*upnito.sk*  
*uptobox.com*  
*usershare.net*  
*videobb.com*  
*vip-file.com*  
*wiiupload.net*  
*wupload.com*  
*x7.to*  
*yourfilehost.com*  
*yourfiles.biz*  
*youtube.com*  
*yunfile.to*  
*ziddu.com*  
*zomgupload.to*  
*zshare.net*

Discard

DLS

DNS

Dofus

Doom3

Dook

Dota 2

Dota Legend

DriveLock

Dropbox

DRP

Duo Duo

E

EA Origin

eBay

eBuddy

- eBuddy BrowserApplet
- eBuddy iOS

echo

eDonkey

*eDonkey-plain*

*eDonkey-encrypted*

- aMule
- Hydranode
- Jubster
- Lphant
- MLDonkey
- Morpheus
- Shareaza
- WinMX MP3
- XoloX

EGP

Endomondo

enLegion

Eset

Evernote

ExindaAD

ExindaCom

ExindaSSL

ExindaWM

F

F-Prot

Facebook

Facetime

Feidian

FiCall

- FiCall iOS

Fiesta

Filemaker

Filetopia

Finger

FIX

- QuickFIX

Flash

*Flash-host*

Flickr

Florensia

FNA

Forfone

Foursquare

Freenet

Friendster

*FTP-control*

*FTP-data*

Fring

*Fring-unknown*

*Fring-video*

*Fring-voice*

- Fring Android
- Fring Symbian
- Fring iOS

FTP

FTP-SSL

Funshion

Funshion

Fuze Meeting

G

Gadu-Gadu

*Gadu-Gadu-server*

*Gadu-Gadu-file-transfer*

*Gadu-Gadu-voice*

*Gadu-Gadu-video*

- Adium
- GaduGadu Android
- GaduGadu iOS
- Instantbird
- Kopete
- Miranda
- Pidgin

Game of War

GameKit

Gmail

Gnutella

- Acqlite
- Acquisition
- Apollon
- BearFlix
- Cabos
- CocoGnut
- DM2FrostWire
- Filepipe
- giFT
- Gluz
- Gnucleus
- Gtk-gnutella
- iMesh Light
- KCEasy
- Kiwi Alpha
- LimeWire
- Morpheus
- Mutella
- Phex
- Poisoned
- Qtella
- Shareaza
- Swapper.NET
- Symella
- WinMX MP3

- XFactor
- XNap
- XoloX

Goober

Google

*Google-docs*

*Google-cloud*

*Google-drive*

*Google-plus*

*Google-encrypted*

Google Cloud

Google Docs

Google Drive

Google Earth

Google Encrypted

Google Play

Google Play Music

Google Plus

Google Talk

*Google Talk-file-transfer*

*Google Talk-unknown*

*Google Talk-video*

- Adium
- Ayttn
- GMX Multimessenger
- iChat
- Instantbird
- Miranda
- Pidgin (previously Gaim)
- Kopete
- Pidgin
- Proteus
- Psi
- Qnext
- Trillian/Trillian Pro
- Web.de Multimessenger

*Google Talk-voice*

Gopher

GoToMeeting

GoToMyPC

GRE

Grisoft

Grooveshark

Groupwise

GTP

*GTP-control*

*GTP-prime*

*GTP-user-data*

GuildWars

H

H323

- Ekiga
- NetMeeting
- OpenH.323
- XMeeting

Half-Life 2

HamachiVPN

- LogMeIn Hamachi

Hi5

HiDrive

Hike Messenger

HiNet

HL7

Hotmail

HTTP

*HTTP-advanced*

*HTTP-content type*

*HTTP-file*

*HTTP-host*

*HTTP-method*

*HTTP-referer*

*HTTP-x\_forwarded\_for*

HTTP-ALT

HTTPS

## HTTPTunnel

- HTTP-Tunnel Client

## Hulu

## HyRead

## I

## IAX

- Diax
- Firefly
- IAXComm
- IAXPhone
- Idefisk/Zoiper
- Kiax
- LoudHush
- PURtel
- YakaPhone
- ZiaxPhone

## ICAP

## Icecast

## iCloud

- Orkut iOS

## ICMP

## ICMPv6

## ICQ Webchat

## ident

## IDrive

## IEC61850

*IEC61850-goose*

*IEC61850-manufacturing message*

*IEC61850-sampled values*

## IGMP

## iHeart Radio

## iLeader

## iLive.to

## IM+

- IM+ Android
- IM+ Blackberry

- IM+ Symbian
- IM+ Windows Phone
- IM+ for Java phones
- IM+ iOS

IMAP

*IMAP-encrypted*

*IMAP-plain*

IMAP-SSL

iMesh

*BearShare*

*iMesh*

*iMesh Lite*

*Lphant*

*Shareaza*

iMessage

IMO

- IMO Android
- IMO Blackberry
- IMO Browser Applet
- IMO for iPad
- IMO iOS

IMS

Instagram

Intralinks

IOS OS Update

IP mobility

iPerf

IPIP

iPlayer

- iPlayer iOS

IPP

IPSEC

*IPSEC-control*

- Cisco VPN Client
- HotSpot Shield iOS

*IPSEC-data*

IPSEC over UDP

IPTV

- T-Home/TV2

IPX

IRC

*IRC-file-transfer*

- Adium
- Ayttn
- Fire
- Instantbird
- Mibbit Web IRC
- Miranda
- mIRC
- Pidgin

*IRC-unknown*

IRC-SSL

isakmp

iSkoot

iTunes

iTunes Radio

ITV

*ITV-stream*

*ITV-web*

- ITV Player iOS

iWow systex

J

Jabber

*Jabber-encrypted*

*Jabber-file-transfer*

*Jabber-unknown*

*Jabber-voice*

*Jabber-video*

- Adium
- Ayttn
- GMX Multimessenger
- iChat
- Instantbird

- Kopete
- Miranda
- Pidgin
- Pidgin (previously Gaim)
- Proteus
- Psi
- Qnext
- Trillian/Trillian Pro
- Web.de Multimessenger

## JAP

- JAP/JonDo

## Java-RMI

## JBK3000

## JD.com

## JDEnet

## K

## Kakao

*Kakao-generic*

*Kakao-kakaotalk*

## Kakao Talk

## Kaspersky

## Kaspersky Antivirus

## Kazaa-Fasttrack

- Apollon
- Filepipe
- iMesh Light
- Kazaa
- Kazaa Lite
- Mammoth

## Kerberos

## Kik Messenger

## Kontiki

## L

## L2TP

- OpenL2TP

## Label Distribution Protocol

last.fm

LDAP

*LDAP-plain*

LDAP-SSL

LeTV

Line

*Line-audio*

*Line-generic*

*Line-video*

LinkedIn

Live.com

LiveJasmin

LivePerson

LLMNR

lockd

LotusNotes

LoveFilmLync

*Lync-audio*

*Lync-file-transfer*

- Microsoft Lync

*Lync-unknown*

*Lync-video*

Lync

M

macfile

Magine TV

MagV

MagV Kids

Manolito

- Blubster
- Piolet

MAPI

MapleStory

MATIP

McAfee

MCK

mDNS

Meebo

*Meebo-audio*

*Meebo-unknown*

*Meebo-video*

- Meebo Android
- Meebo Blackberry
- Meebo MiniBar
- Meebo Notifier
- Meebo iOS

Meeting Maker

MEGA

Megaco H.248

Mei Lu

MGCP

*MGCP-control*

*MGCP-rtp*

Microsoft Dynamics AX

Microsoft Dynamics NAV

Microsoft Exchange

*Microsoft Exchange-generic*

*Microsoft Exchange-outlook web access*

Microsoft Media Server

Mig33

*mig33 Android*

Mitake

MMS

- Windows Media Player

Modern War

Mojo

Move Player

Moves

MPEG

*Jupster*

*Screamerradio*

*VLC*

*Winamp*

MPlus

MQTT

*Outlook Web Access*

MS Global Catalogue

MS Global Catalogue SSL

MS Local Security Architecture

MS-SQL

msmq

MSN

*MSN-file-transfer*

*MSN-unknown*

*MSN-video*

*MSN-voice*

*MSN-webchat*

- aMSN
- Adium
- Ayttn
- Fire
- GMX Multimessenger
- Instantbird
- Kopete
- Mercury
- Microsoft Messenger Mac
- Miranda
- Pidgin
- Pidgin (previously Gaim)
- Web.de Multimessenger
- Qnext
- Trillian/Trillian Pro
- Windows Live
- Windows Live iOS

MSRP

- Blink

MSRPC

mstask

Mute

MyBook

MyCard

MyFitnessPal

MyMusic TW

MyPeople

*MyPeople-audio*

*MyPeople-unknown*

MySpace

MySQL

MyVideo

*Stream*

*Web*

- Jupster
- Screamer Radio
- VLS
- Winamp

N

nat-t

National Baseball

Naver

NCP

Net2Phone

NETBIOS

Netbus

Netflix

*Netflix-generic*

*Netflix-video*

- Netflix iOS

Netflow

Netlogon

NetMotion

Next TV

NFS

Nike+

Nimbuzz

- Nimbuzz Android
- Nimbuzz Blackberry
- Nimbuzz Mac

- Nimbuzz Symbian
- Nimbuzz iOS

NNTP

NNTP-SSL

Norton Antivirus

NTLM

*NTLM-encrypted*

*NTLM-plain*

NTP

O

Octoshape

- CNN streaming

Odnoklassniki

Office 365

OFFSystem

OGG

olap

ooVoo

- ooVoo iOS

Open FT

OpenVPN

- Hotspot Shield
- MyFreedom

Opera Mini

- Opera Mini Symbian
- Opera Mini Android
- Opera Mini Blackberry
- Opera Mini iOS

Oracle

- Oracle DB

Oracle-SSL

ORB

*ORB-server*

*ORB-client*

*ORB-client-3gp*

*ORB-client-flash*

*ORB-client-flv*  
*ORB-client-rdt*  
*ORB-client-rtp*  
*ORB-client-rtsp*  
*ORB-client-upload-download*  
*ORB-client-wma*  
*ORB-client-wmv*  
*ORB-server*  
*ORB-unknown*

Orbix2000

Orbix2000 SSL

Orkut

OS Update

*Android*

*iOS*

*Windows Phone*

Oscar

*Oscar-file-transfer*

*Oscar-picture-transfer*

*Oscar-unknown*

*Oscar-video*

- AIM
- AIM Android
- AIM for Mac
- AIM for iPad
- AIM iOS
- Adium
- Ayttn
- Fire

*Oscar-voice*

- GMX Multimessenger
- iChat
- ICQ
- ICQ Android
- ICQ Mac OS X
- ICQ iOS

- ICQ mobile for Java
- Instantbird
- Kopete
- Miranda
- Pidgin
- Pidgin (previously Gaim)
- Proteus
- Qnext
- Trillian/Trillian Pro
- Web.de Multimessenger

OSPF

P

Paltalk

*Paltalk-file-transfer*

*Paltalk-server*

*Paltalk-video*

- Paltalk Android
- Paltalk iOS

*Paltalk-voice*

Panda Security

Pando

Pandora

Path

PCanywhere

PCoIP

PDProxy

PeerEnabler

Periscope

PinkFong

Pinterest

PlayTales

Pochta Rossii

Poison Ivy

POP

*POP-encrypted*

*POP-plain*

POP-SSL

POPO

PornHub

PostgreSQL

PPLive

*PPLive-generic*

*PPLive-video*

- PPTV
- PPTV Android
- PPTV iOS
- PPTV HD iOS
- PPTV Mac

PPStream

- PPStream Android
- PPStream Mac
- PPStream Ubuntu
- PPStream iOS

PPTP

Print

PS3

pubu bookbuffet

Q

QIP.RU

QQ

*QQ-file-transfer*

*QQ-unknown*

*QQ-voice*

*QQ-video*

- Instantbird
- QQ HD iOS
- QQ International
- QQ iOS
- QQ for Mac
- QQ for Pad Android

QQ Games

QQLive

- QQLive Android
- QQLive iOS
- QQLive iPad

Quake-HalfLife

- Call of Duty
- Quake HalfLife

Qualsys

QUIC

QuickTime

*Quicktime-host*

- iTunes

Quotd

QVOD

R

Radius

radmin

Rally

RC5DES

Rdio

RDP

RealDataTransport

RealMa

RealMedia

- RealMediaPlayer

Redmine

RedTube

Remote Telnet

Remote Web Workplace

Renren

Replify

rFactor

Rhapsody

RIP

rlogin

RRP

RSVP

rsync

RTCP

RTP

*codec*

RTSP

RunKeeper

Runtastic

S

SalesForce

SAP

SAP.MCAST.NET

SCTP

Scydo

- Scydo Android
- Scydo iOS

Second Life

SecureSight

SecureSight-SSL

Shell-SSL

Shoutcast

- Fishradio
- iTunes

Silverlight

*Silverfast-host*

Simfy

Sina Weibo

SinaTV

SIP

*SIP-audio*

*SIP-gadu-gadu*

*SIP-ichat-facetime*

*SIP-MagicJack*

*SIP-mplus*

*SIP-msn*

*SIP-oscar*

*SIP-oscar-video*

*SIP-unknown*

*SIP-video*

*SIP-vonage*

*SIP-yahoo*

*SIP-yahoo-video*

- ClearSea
- Ekiga
- ExpressTalk
- Facetime
- Facetime iOS
- gooberVoIP
- Gizmo
- KPhone
- MagicJack/MagicTalk
- OpenWengo
- QuteCom
- SipGate
- Twinkle
- Vonage
- XMeeting

SIP-SSL

Skinny

Skout

SkyDrive

Skyking

Skype

*Skype-chat*

*Skype-file transfer*

*Skype-out*

*Skype-unknown*

*Skype-video*

- Skype Android
- Skype for iPad
- Skype iOS
- Skype Linux
- Skype Symbian

*Skype-voice*

Skype for Business – File Transfer

Skype Video

Skype Voice

Slingbox

SMS

SMTP

*SMTP-encrypted*

*SMTP-plain*

Snapchat

SnapMirror

SNMP

SOAP

SOCKS

*Socks v4*

*Socks v5*

SoftEthernet

Sohu

SOPCAST

Soribada

Soulseek

- iSoul
- Nicotine Plus
- SolarSeek
- SoulSeeX
- pySoulSeek

SoundCloud

SPDY

Speakaboos

Speedtest

SpiderOak

SplashFighter

Spotify

- Spotify Android
- Spotify iOS
- Spotify Mac

Spreecast

Spreed

SPTG TV

sqidproxy

ssdp

SSH

SSL

- HTTPS
- IMAPS
- POP3S
- SMTPS
- Stunnel

*SSL-advanced*

*SSL-common name*

*SSL-organization name*

*SSL-SPDY*

SSTP

StealthNet

Steam

StreamWorks

StreetVoice

STUN

Sudaphone

- Sudaphone Android
- Sudaphone iOS

SugarSync

SunRPC

svrloc

Sybase

Syncplicity

SyncServer-SSL

syslog

T

T.120

Taaze

TACACS

Tango

*Tango-audio*

*Tango-file-transfer*

*Tango-im*

*Tango-unknown*

*Tango-video*

- Tango Android
- Tango iOS

Taobao

TDS

TeamSpeak

- Teamspeak Linux
- Teamspeak Mac
- Teamspeak iOS

TeamViewer

- TeamViewer Android
- TeamViewer Linux
- TeamViewer Mac
- TeamViewer for iPad
- TeamViewer HD for iPad
- TeamViewer for Meetings iPad
- TeamViewer iOS
- TeamViewer for Meetings iOS

Telegram

Telnet

*Telnet-generic*

*Telnet-tn3270*

Telnet-SSL

Teredo

textPlus

TFTP

Three Bamboo

Threema

Thrift

Tibia

TigerText

Timbuktu

Time Server

Timely TV

TMall

TN3270

TOR

- Tor Vidalia Bundle

Trendmicro

Truphone

- Truphone Android
- Truphone Mac
- Truphone iOS
- Truphone iPad

TTPoD

Tudou

Tumblr

TVAnts

TVBS

TVUPlayer

- TVUPlayer Android
- TVUPlayer Mac
- TVUPlayer iOS

Twitch

Twitter

U

Ubuntu One

UltraBac

UltraSurf

Unreal

UppTalk

Usenet

- iLoad
- UseNeXT

uucp

UUSee

V

vBulletin

Vcast

Ventrilo

VeohTV

*VeohTV-generic*

*VeohTV-flash*

- VeohTV Web Player

Vibe

Viber

*Viber-audio*

*Viber-file-transfer*

*Viber-Instant Messaging*

- Viber Android
- Viber iOS

*Viber-out*

*Viber-unknown*

Vimeo

Vine

Vippie

- Vippie Android
- Vippie iOS

VK

VNC

VoIP Tunnel

Voxer

VPN-X

VRRP

VTUN

VyprVPN

W

WAP

WAP TLS

Warcraft 3

Watchever

Wealth God

WebDAV

- Cadaver
- Goliath

Webex

- WebEx Meeting Center Applet
- Webex for iPad
- Webex iOS

WebQQ

WebSocket

Webthunder

- Thunder

WeChat

*WeChat-Chat*

*WeChat-Data Message*

*WeChat-Generic*

*WeChat-Video Call*

*WeChat-Voice Live Chat*

*WeChat-WebChat*

WhatsApp

- WhatsApp Android
- WhatsApp Symbian
- WhatsApp iOS

*WhatsApp-file transfer*

*WhatsApp-generic*

*WhatsApp-voice call*

Whois

Wickr

Wii

Windows Azure

Windows Phone OS Update

Windows Updates

WindowsMedia

*WindowsMedia-host*

WinMX

- XNap

WINNY

Wins

Workout Trainer

World of Kung Fu

World of Warcraft

Wowza

WSM Server

WSM Server-SSL

Wuala

X

x400MTA

xbox-live

XDCC

- XDCC Catcher
- XDCC Fetcher
- XDCC Klipper
- mIRC

xdmcp

Xing

XNXX

XVideos

XWindows

Y

Yahoo

*Yahoo-file-transfer*

*Yahoo-unkown*

*Yahoo-video*

*Yahoo-voice*

*Yahoo-webchat*

*Yahoo-webmail*

- Adium
- Ayttn

*GMX Multimessenger*

- Instantbird
- Kopete
- Miranda
- Pidgin
- Pidgin (previously Gaim)
- Qnext
- Trillian/Trillian Pro

*Web.de Multimessenger*

- Yahoo Messenger
- Yahoo Messenger Android
- Yahoo Messenger for Mac
- Yahoo Messenger iOS

Yahoo Games

Yahoo IM

Yahoo WebMail

Yandex

Yelp

YikYak

Youku

YouPorn

YourFreedom

YouTube

Yuanta Web

Z

Zattoo

- Zattoo HD iOS
- Zattoo iOS

Zoho Work Online

Zynga

## 0.41 Predefined Application Groups

The Exinda appliance includes predefined groups of applications that can be used when setting policy for the appliance and for monitoring. Custom application groups can also be defined, and custom applications can be added to existing application groups.

### NOTE

A given application can exist in multiple application groups. However, monitored groups must not contain applications which are already a member of another group being monitored. Any given application can only be monitored within a single application group.

### Database Services

- » Cassandra Internode Communication
- » Cassandra Query Language
- » MS-SQL
- » MySQL
- » Oracle
- » PostgreSQL
- » TDS

### Exinda

- » ExindaAD
- » ExindaCom
- » ExindaWM

- » ExindaSSL

## File Services

- » Apple Talk

- » CIFS

- » FLUTE

- » FTP

- » Intralinks

- » NCP

- » NFS

- » rsync

- » Syncplicity

- » TFTP

- » Thri

- » ft

## Games

- » Age of Warring Empire

- » Armagetron

- » Battlefield

- » Call of Duty

- » Club Penguin

- » Crime City

- » Crossfile

- » Destiny

- » Dofus

- » Doom3

- » Dota 2

- » Dota Legend

- » Fiesta

- » Florensia

- » Game of War

- » Gamekit

- » GuildWars

- » Half-Life 2

- » King of Pirates

- » MapleStory

- » Modern War

- » MyCard
- » PinkFong
- » PS3
- » QQ Games
- » Quake-HalfLife
- » rFactor
- » SecondLife
- » Skyking
- » SplashFighter
- » Steam
- » TeamSpeak
- » Tibia
- » Unreal
- » Warcraft 3
- » WealthGod
- » Wii
- » World of Kung Fu
- » World of Warcraft
- » xbox-live
- » Zynga

#### Instant Messaging

- » Blackberry Messenger
- » DiDi IM
- » eBuddy
- » enLegion
- » FETnet
- » Google Talk
- » Hike Messenger
- » ICQ Webchat
- » IM+
- » iMessage
- » IMO
- » IRC
- » Jabber
- » Kik Messenger
- » Live Person
- » Meebo

- » MPlus
- » MQTT
- » MSN
- » MSRP
- » Oscar
- » PalTalk
- » Path
- » Telegram
- » textPlus
- » TigerText
- » Vibe
- » Voxer
- » WebQQ
- » WeChat
- » WhatsApp
- » Wickr
- » Yahoo IM

#### Interactive

- » BGP
- » DNS
- » echo
- » Remote Telnet
- » rlogin
- » SSH
- » Telnet
- » Time Server
- » TN3270

#### Mail

- » Blackberry Mail
- » GMail
- » Hotmail
- » IMAP
- » IMAP-SSL
- » LotusNotes
- » MAPI
- » POP

- » POP-SSL
- » SMTP
- » Yahoo WebMail

## P2P

- » ANts P2P
- » Ares
- » BitTorrent
- » Edonkey
- » Feidian
- » Filetopia
- » Gadu-Gadu
- » Gnutella
- » Kazaa-Fastrack
- » Kontiki
- » OFFSystem
- » Open FT
- » POPO
- » PPLive
- » PPSteam
- » QQ
- » QQLive
- » Soulseek
- » StealthNet
- » UUSEE
- » Webthunder
- » ZATTOO

## Recreational

- » Anonymous Proxy
- » Audiko
- » Audiogalaxy
- » DirectDownload
- » Duokan
- » Endomondo
- » Facebook
- » Flash
- » Flickr

- » Foursquare
- » Hi5
- » Hotmail
- » Instagram
- » iTunes
- » JAP
- » Kollect Book
- » LinkedIn
- » MagV
- » MagV Kids
- » MapMyRun
- » Mig33
- » Mojo
- » Moves
- » MyFitnessPal
- » MySpace
- » National Baseball
- » Nike+
- » Orkut
- » Pinterest
- » PlayTales
- » RealMedia
- » Run Keeper
- » Runtastic
- » Sina Weibo
- » Snapchat
- » Sohu
- » Speakaboos
- » Spotify
- » SPTG TV
- » TTPod
- » Tumblr
- » UltraSurf
- » Vine
- » VoIP Tunnel
- » Watchever
- » WindowsMedia

- » Workout Trainer
- » Xing
- » YourFreedom
- » YouTube

## Secure

- » DDM-SSL
- » FTP-SSL
- » HTTPS
- » IMAP-SSL
- » IRC-SSL
- » L2TP
- » LDAP-SSL
- » MS Global Catalogue SSL
- » NNTP-SSL
- » Oracle-SSL
- » Orbix2000 SSL
- » POP-SSL
- » PPTP
- » SecureSight
- » Shell-SSL
- » SIP-SSL
- » SSH
- » SSL
- » SyncServer-SSL
- » Tacacs
- » Telnet-SSL
- » WAP TLS
- » WSM Server-SSL

## Social Networking

- » amini
- » Bebo
- » Dook
- » Facebook
- » Flickr
- » Foursquare
- » Friendster

- » Google Plus
- » Hi5
- » Instagram
- » iSkoot
- » Kakao
- » LinkedIn
- » MySpace
- » Orkut
- » Pinterest
- » Renren
- » Sina Weibo
- » Skout
- » Snapchat
- » StreetVoice
- » Tumblr
- » Twitter
- » Vine
- » Xing
- » Yelp
- » Yik Yak
- » YouTube

## Software Updates

- » Adobe Updates
- » Android OS Update
- » Antivirus Lab
- » Apple Updates
- » Avira
- » Bitdefender
- » Eset
- » F-Prot
- » Grisoft
- » iOS OS Update
- » Kaspersky
- » McAfee
- » Norton Antivirus
- » Panda Security
- » Windows Phone OS Update

» Windows Updates

## Streaming

» 1kxun

» 9ku

» Amazon Prime Music

» AVI

» BlueJeans

» Channel 4oD

» Ciaomobile HDTV

» CNN

» CNTV

» CtiTV

» Dailymotion

» Deezer

» Demand 5

» Flash

» Funshion

» Google Play Music

» Grooveshark

» Hulu

» Icecast

» iHeartRadio

» iLive.to

» iPlayer

» IPTV

» iTunes Radio

» ITV

» last.fm

» LoveFilm

» MMS

» Move Player

» MPEG

» MyMusic TW

» MyVideo

» Netflix

» Octoshape

» OGG

- » Pandora
- » Periscope
- » QUIC
- » Quicktime
- » QVOD
- » RealDataTransport
- » RealMedia
- » Rhapsody
- » RTSP
- » Shoutcast
- » Silverlight
- » Simfy
- » SinaTV
- » Soribada
- » SoundCloud
- » Spreecast
- » Timely TV
- » TTPod
- » TuDou
- » TVBS
- » VeohTV
- » Vimeo
- » Watchever
- » WindowsMedia
- » Youku

#### Thin Client

- » AdobeConnect
- » Citrix
- » GoToMeeting
- » GoToMyPc
- » PCAnywhere
- » PCoIP
- » RDP
- » Spread
- » TeamViewer
- » VNC
- » Webex

- » xdmcp
- » XWindows

## Unified Communications

- » FiCall
- » Forfone
- » Fring
- » Goober
- » Lync
- » ooVoo
- » Scydo
- » Tango
- » Truphone
- » UppTalk
- » Ventrilo

## Voice

- » Burner
- » DiDi Voice
- » Facetime
- » H323
- » IAX
- » Megaco H.248
- » MGCP
- » Net2Phone
- » Nimbuzz
- » RTCP
- » RTP
- » SIP
- » Skinny
- » Skype
- » Skype Video
- » Skype Voice
- » Sudaphone
- » Vippie

## VPN

- » CpmodoUnite
- » CyberGhost

- » GRE
- » HamachiVPN
- » HTTP Tunnel
- » IPSEC
- » IPSEC over UDP
- » isakmp
- » L2TP
- » NetMotion
- » OpenVPN
- » PDProxy
- » PPTP
- » SoftEthernet
- » SSTP
- » Teredo
- » VPN-X
- » VyprVPN

## Web

- » Baidu
- » BOLT
- » Google
- » Google Cloud
- » Google Drive
- » Google Encrypted
- » HTTP
- » HTTP-ALT
- » HTTPS
- » Live.com
- » Naver
- » OperaMini
- » SPDY
- » squidproxy
- » WAP
- » WebDAV
- » WebSocket