

Exinda Appliance User Manual



Exinda ExOS Version 6.4
© 2013 Exinda Networks, Inc.



Copyright

© 2013 Exinda Networks, Inc. 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 the 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.

Document Built on Tuesday, November 19, 2013 at 1:31 PM

Using this guide

Before using this guide, become familiar with the Exinda documentation system.

- ["Exinda documentation conventions" on page 2](#)
- ["Notes, Tips, Examples, and Cautions" on page 3](#)

Exinda documentation conventions

The Exinda documentation uses the following conventions in the documentation.

Graphical interface conventions

The following is a summary of the conventions used for graphic interfaces such as those in the Exinda Web UI and the Central Management Technical Preview UI.

Convention	Definition
bold	Interface element such as buttons or menus. For example: Select the Enable checkbox.
<i>Italics</i>	Reference to other documents. For example: Refer to the <i>Exinda Application List</i> .
>	Separates navigation elements. For example: Select File > Save .

Command line conventions

The following is a summary of the syntax used for the CLI commands.

```
(config)# command <user input> keyword {list|of|options|to|select|from} [optional  
parameter]
```

Convention	Definition
monospace text	Command line text or file names
< <i>courier italics</i> >	Arguments for which you use values appropriate to your environment.
courier bold	Commands and keywords that you enter exactly as shown.
[x]	Enclose an optional keyword or argument.
{x}	Enclose a required element, such as a keyword or argument.
	Separates choices within an optional or required element.
[x {y z}]	Braces and vertical lines (pipes) within square brackets indicate a required choice within an optional element.
command with many parameters that wrap onto two lines in the documentation	Underlined CLI commands may wrap on the page, but should be entered as a single line.

Notes, Tips, Examples, and Cautions

Throughout the manual the following text styles are used to highlight important points:

- **Notes** include useful features, important issues. They are identified by a light blue background.

Note Note text

- **Tips** include hints and shortcuts. They are identified by a light blue box.

Tip Tip text

- **Examples** are presented throughout the manual for deeper understanding of specific concepts. Examples are identified by a light gray background.

Example

Text

- **Cautions** and warnings that can cause damage to the device are included when necessary, and are highlighted in yellow.

Caution Caution text

Table of Contents

Chapter 1: Learn About	20
WAN optimization	20
Exinda Appliance product family	21
Safety and Compliance	23
Compliances	24
Safety Guidelines	24
EMC Notice	25
Chapter 2: Installation And Deployment	26
Package Contents	26
Pre-Installation Checklist	26
Deployment topologies	27
Platforms	28
IP Modes	28
In-path Topologies	29
Main Site Internet Link Topology	30
Main Site WAN Link Topology	31
Distributed Branch Topology	32
Topologies with Firewalls	33
Topologies with VPNs	35
Multiple Link Topology	37
Out of Path Topologies	38
Accelerate traffic with WCCP	38
Configure the Exinda appliance to use WCCP	39
Configure internal subnets as internal network objects	40
Display the state of the WCCP service	41

SPAN and Mirror Port Monitoring	41
Configure Mirror Port Mode	42
Enable Monitoring of Mirror/SPAN Traffic	42
Configure internal subnets as internal network objects	43
Monitor Span/Mirror Traffic	44
Directing traffic with policy-based routing	44
Basic policy-based routing	45
Policy-based routing with two subnets	49
VRRP with PBR	53
VRRP with PBR and VLANs	58
VRRP with PBR and IP SLA Tracking	66
Cluster and High Availability	73
Redundancy through multiple Exinda appliances	73
Load balancing and fail-over with multiple Exinda appliances	75
High availability mode	77
Cluster Interfaces	78
Cluster Failover	79
Cluster Terminology	79
Create a cluster of Exinda appliances	81
Add Exinda appliances to the cluster	81
Specify what data is synchronized between cluster members	82
View the status of all members of the cluster	83
Virtual Appliances	84
Virtual Appliance Deployment Options	84
Understanding virtual appliance resource requirements	85
Sizing Guidance	86
Monitor IOPS in VMware vSphere	87

Exinda model 2750	88
Exinda model 4750	88
Exinda model 6750	90
Exinda model 8750	90
Exinda Model 2850	91
Exinda model 4850	92
Exinda model 6850	93
Run the Virtual Appliance on VMware vSphere (ESX and ESXi)	93
Understanding how VMotion works	93
Install the Virtual Appliance on VMware	94
Modify the VMware Virtual Machine Configuration	95
Convert NICs into a Bridge	98
Allow Ports to Accept and Bridge Packets	99
Start the VMware Virtual Appliance	104
Install the Silicom Bypass Driver	105
Virtual appliance use cases	108
In-line deployment with externally attached LAN (VMware ESXi)	109
In-line deployment with an isolated virtual LAN and virtual applications (VMware ESXi)	111
Out-of-band (WCCP) mode (VMware ESXi)	113
Port mirroring with an external Nexus switch	114
Port mirroring with a virtual Nexus switch	115
Virtual WAN simulator in an isolated network (VMware ESXi)	117
Run the Virtual Appliance on Citrix XenServer	121
Installing the Virtual Appliance on XenServer	121
Custom Settings	123
Additional NICs	123
Add storage to the XenServer virtual appliance	126

Booting	128
Licensing	129
Generating a virtual appliance trial license	129
Purchasing a virtual appliance license	130
Hypervisor limitations	130
Topology troubleshooting	131
Maintenance	132
Manage System Configuration	132
Import System Configuration	133
Install an update to the Exinda appliance software	134
Factory Defaults	135
Reboot the Exinda appliance	135
Chapter 3: Access The Appliance	137
Access the Command Line Interface	137
Web User Interface (Web UI)	138
Switch between Exinda Web UI display modes	138
Check the Features Included in your License	139
Chapter 4: Initial Configuration	140
CLI Configuration Jumpstart	140
Web UI Basic Wizard	142
Chapter 5: Dashboards	146
System Dashboard	146
Benefits Dashboard	147
Chapter 6: System Settings, Configuration, And Diagnostics	152
Network Settings	152
NIC Settings	152
IP Address	154
Cluster and High Availability	158

Routes	158
Configure DNS and Domain Names	159
Set the host name and DNS of the Exinda appliance	160
Add a domain name	160
Remove a domain name	160
HTTP Proxy	160
Configure the appliance to send email notifications	161
Add an SMTP server for sending email notifications	161
Add a user to receive email notifications	161
Stop sending notifications to a user	162
SNMP	162
Modify the SNMP configuration	163
Remove an SNMP community	163
Download the SNMP MIB file	163
Modify the SNMP administrator user settings	164
Add an SNMP trap sink server	164
Enable or Disable an SNMP trap sink server	164
Remove an SNMP trap sink server	164
Integrate the Exinda Appliance with Active Directory	165
Install the Exinda Active Directory Connector	165
Add Exinda appliances to the Active Directory Connector	167
Change the Active Directory Connector port number	167
Select the information sent between the Exinda appliance and the Active Directory server	168
Configure the connection to the Active Directory server	169
Verify communication between the Active Directory server and the Exinda Appliance	169
Request updated user and group information from the Active Directory server	170
Change the state of the Exinda Active Directory Connector	170

Exclude specific usernames from reports	171
Identify users using applications on a Citrix XenApp server	172
Report on Network Activity by User	176
Set the Time Period Reflected in the Report	177
Understanding the Conversation Report	184
Controlling Traffic based on Users	185
Troubleshoot issues with Active Directory configuration	188
IPMI Overview	190
Configure IPMI	191
Manage Power Settings on an IPMI Enabled Appliance	192
Remove Events from the Appliance System Log	193
System Setup	193
Date and Time Configuration	193
Set the date and time of the appliance	194
Add an NTP server	194
Disable an NTP server	194
Remove an NTP server	194
UI Configuration	195
Configure when the Exinda Web UI logs out	196
Configure when the CLI console times out	196
SDP Configuration	197
Configure SQL Access	197
Download the ODBC Driver	197
Set Remote SQL Options	197
View SQL Access data in Microsoft Excel	206
SQL Schema	208
Monitoring Configuration	216

Configure monitoring settings	216
Control the order that IP addresses are resolved	217
Enable or disable Application Specific Analysis Modules	218
Identify the statistics to collect	219
Clear saved monitor statistics	220
Netflow Configuration	221
Create a Scheduled Job	224
Notify administrators of system issues	225
License	226
Control Configuration	229
Allocate Disk Storage for System Services	229
Configure Storage with CLI	231
Remove all data from a service's disk storage	234
Optimization	234
Auto Discovery	235
Manage optimization services	237
Group appliances into a community	238
Add an Exinda appliance to the community	238
Change the IP address of an appliance in the community	238
Remove an Exinda appliance from the community	239
Specify the community groups an Exinda appliance can join	239
Remove the community from a community group	240
Accelerate file transfers	240
SMB optimizations	242
Common SMB Use Cases	243
Configure file acceleration with SMB1	244
Enable file acceleration with SMB2	245

Accelerate digitally signed SMB connections	245
Allow file transfer acceleration with older versions of Exinda OS	246
Verify acceleration configuration	246
Report: Accelerated Connections	247
TCP Acceleration	248
Configure TCP Acceleration	249
WAN Memory	251
Add SSL certificates and keys on the Exinda appliances	252
View all certificates and private keys	253
Export an SSL certificate	253
Import a certificate onto an appliance	253
Generate a self-signed certificate	254
Display the contents of a certificate	255
Delete an SSL certificate	255
Add servers for SSL acceleration	255
Create policies to accelerate SSL traffic	257
Host multiple secure websites on Windows Server 2012	258
Host multiple secure websites on Apache	261
Accelerate web applications through caching	263
Create the Edge Cache Policy	264
Add the Edge Cache Policy to a Virtual Circuit	264
Configure the Edge Cache Default Settings	264
Exclude URLs from the Edge Cache	265
Add an Edge Cache Peer	265
Remove All Objects from the Edge Cache	266
View Edge Cache Statistics	266
Pre-populate the cache	267

Prerequisites for Pre-population	268
Create a pre-population job with the Exinda Web UI	268
Create a Scheduled Job	271
Accelerate Exchange and Microsoft Outlook traffic	272
Enable MAPI Acceleration on the Exinda Appliances	272
Turn off MAPI encryption in Microsoft Outlook	273
Disable encryption on the Exchange server	275
Verify MAPI traffic is being accelerated	275
Understanding the Conversation Report	278
Troubleshoot problems with MAPI acceleration	279
Add an SMTP server for sending email notifications	281
View the status of an alert	281
Specify application quality based on host	284
Set a per-host limit on bandwidth usage	284
Specify when multi-queue is activated	287
View the number of hosts on a Dynamic Virtual Circuit	287
View the data throughput on the interfaces	288
View the outbound packet rate for all traffic	289
Per Host QoS Usage Examples	290
Controlling Traffic based on Users	307
Create Network User Objects	308
Create Network Group Objects	309
Optimize Traffic Based on Users and Groups	309
Add SSL certificates and keys on the Exinda appliances	310
Virtualization	310
Virtualization overview	311
Virtualization requirements	312

Products supported on the virtualization partition	312
Installing Virtualization	313
Exinda SDP	314
Exinda SDP VA	314
Exinda Mobile Suite	315
Install the Exinda Mobile Server on an Exinda appliance	316
Install the Exinda Mobile Manager on an Exinda appliance	318
Microsoft Windows Server	321
Microsoft Windows Server 2008 R2	321
Authentication	322
Display a List of Active Users	322
Local User Accounts	323
AAA	324
LDAP Authentication	325
Radius Authentication	325
TACACS+ authentication	326
Logging	326
View System Log Files	327
Live Log	327
Tail Log	327
System Logging Configuration	327
Configure the appliance log files	328
Add a remote syslog server	328
Remove a remote syslog server	328
System Diagnostics	329
View the status of an alert	329
Diagnostics Files	331

TCP Dump	332
View the status of the community	333
Verify acceleration configuration	334
View the SMB configuration and connections	334
View the TCP acceleration configuration and connections	335
View the WAN memory configuration and reduction statistics	336
Monitor	337
Optimizer Diagnostics	338
NIC Diagnostics	339
RAID Diagnostics	339
Open a case with Exinda Networks Support Services	340
Maintenance	341
Manage System Configuration	341
Import System Configuration	342
Cluster and High Availability	343
Install an update to the Exinda appliance software	344
Return to the previously installed version of ExOS	345
Factory Defaults	345
Reboot the Exinda appliance	346
Automatically reboot the Exinda appliance	346
Shutdown the Exinda appliance	346
Tools	347
Ping	347
Traceroute	347
DNS Lookup	348
Access the Command Line Interface	348
IPMI	349

Chapter 7: Object Definitions	7
Network Objects	7
Create a static network object	7
Restrict access to management services by IP address	10
Dynamic Network Objects	10
Users and Groups	11
Create Network User Objects	11
Create Network Group Objects	12
VLAN Objects	13
Protocol Objects	14
Configuring application objects	14
Add a new application	15
Application sub-types	19
Add an application group	20
Add an application to an application group	20
Enable anonymous proxy classification	21
Schedule Objects	23
Adaptive Response Rules	25
Adaptive Response	26
Create a Source Network Object	26
Create an Adaptive Response limit rule	27
Use the Adaptive Response Rule in the Optimizer	28
Use Adaptive Response with Active Directory	29
Create Adaptive Response Rules with CLI	30
Add a Dynamic Network Object to Optimizer with CLI	31
Disable an Adaptive Response Rule	31
Exclude Hosts or Subnets from the Quota	31

Other Adaptive Response CLI Commands	32
Service Levels	32
Service Level Agreements	33
Chapter 8: Monitoring And Reporting	35
Set the Time Period Reflected in the Report	35
Interactive Reports	36
Printable Reports	36
Real Time Monitoring	37
Applications	37
Real-time Traffic by Hosts	38
View real-time inbound and outbound conversations	39
Reduction	41
Application Response	41
Host Health	41
Interface Reports	42
View the data throughput on the interfaces	42
View the outbound packet rate for all traffic	44
Network	45
Control Reports	47
Policies Report	47
Discard Report	50
Prioritization Report	51
Optimization Reports	52
Reduction Report	52
View Edge Cache Statistics	55
Service Level Reports	56
View the Application Performance Score results	56
Network Response (SLA) Reports	59

TCP Efficiency Report	60
View TCP health	61
System Reports	64
Connections Report	64
Report: Accelerated Connections	65
CPU Usage Report	66
CPU Temperature Report	66
RAM Usage Report	67
Disk IO	68
Swap Usage Report	68
Applications Report	69
View Unclassified Applications	69
Application Groups Report	70
View All Network Activity for a Specific User	71
URLs Report	73
VoIP Report	74
Report on Network Activity by User	75
Top Users Generating Traffic	75
Set the Time Period Reflected in the Report	76
View All Network Activity for a Specific User	77
Top Internal and External Users on the Network	79
Real-time Traffic by Hosts	80
View real-time inbound and outbound conversations	81
Understanding the Conversation Report	83
Hosts Report	83
Conversations Report	85
Subnets Report	87

PDF Reporting	89
CSV Reporting	93
Monitor application performance on the network	94
How are the metrics calculated?	95
Round trip time	95
Read and write transactions	96
Packet loss	99
View TCP health	100
View the Application Performance Score results	102
Calculate an Application Performance Score	105
Capture Application Performance Metrics	105
Monitor the real time application response	109
Monitor the real time TCP health	110
View a network summary of application groups	111
View TCP efficiency	112
Create an Application Performance Score object	113
Generate recommended Application Performance Score thresholds	114
Review and modify the APS threshold values	115
Generate a PDF report of APS results	116
Chapter 10: Optimizer Configuration	117
Optimizer Policy Tree	118
Create a new circuit	119
Virtual Circuits	119
Virtual Circuit Oversubscription	122
Set a per-host limit on bandwidth usage	123
Policies	125
Optimizer Policies	128

Optimizer Wizard	130
ToS and DiffServ	133
The ToS / DiffServ Field	133
How Exinda Uses the ToS/DiffServ Field	134
Match Packets to ToS/DSCP Values	134
Mark Packets with ToS/DSCP Values	135

Chapter 1: Learn about

Read through any of the topics in the **Learn about** section for conceptual information on this product.

- ["WAN optimization" on page 20](#)
- ["Exinda Appliance product family" on page 21](#)
- ["Safety and Compliance" on page 23](#)

WAN optimization

Networks are becoming more and more congested with business-critical applications, and are in direct competition with all other Internet traffic on the network. Often business applications are delayed because of unwanted and non-critical traffic to a site. Enterprises must be able to balance applications, cloud services, backups, and file transfers over the same link that simultaneously handles VoIP and collaboration. A higher level of user experience assurance is needed to boost performance and drive productivity.

WAN, or Wide Area Network, describes external network connectivity to a business. Optimizing the WAN connections allows network administrators to prioritize inbound and outbound traffic based on a variety of factors. Traffic can be prioritized and de-prioritized by application type, who is generating the traffic, and the time of day the request is being made. For example, traffic flowing between a branch office and the head office network can be prioritized over any other traffic.

An Exinda appliance provides all of the core capabilities needed to effectively manage a WAN in a single network appliance. These tightly integrated capabilities include real-time monitoring, reporting, traffic control, optimization and intelligent acceleration.



Exinda offers an assurance-centric WAN optimization solution that applies over 2,500 unique application and user profiles to directly connect WAN optimization policies with the actual priorities of the business. By intelligently applying its three-point WAN optimization technology, Exinda delivers the right mix of acceleration, caching and containment for the optimal user experience.

Exinda Appliance product family

The following table outlines the physical specifications of each Exinda appliance.

	2061	4010	4061	6062	8062	10062
Designed for	Small Office	Medium Office	Headquarters	Small Data Center	Data Center	Large Data Center
Supported Users	Up to 1,600	Up to 38,000	Up to 51,000	Up to 200,000	Up to 250,000	Up to 400,000
Software						
Wan Capacity (x800)	6 Mbps	20 Mbps	20 Mbps	50 Mbps	155 Mbps	310 Mbps
LAN Capacity (x700)	20 Mbps	500 Mbps	1 Gbps	2.5 Gbps	5 Gbps	10 Gbps
Optimized Connections	600	1,200	3,000	10,000	16,000	45,000
Max Concurrent	32,000	768,000	1,024,000	5,120,000	5,120,000	8,000,000

	2061	4010	4061	6062	8062	10062
Flows						
Hardware						
Form Factor	Desktop or 1U rack mount	1U rack mount	1U rack mount	1U rack mount	2U rack mount	2U rack mount
Data Store/Cache Size	240 GB	385 GB	385 GB	1.7 TB, redundancy built-in	2.7 TB, redundancy built-in	1.5 TB, redundancy built-in
NICs (Default)	2 Bridge Pairs, or 1 Bridge Pair plus 1 Management, 1 Cluster	3 Bridge Pairs, 1 Management / Cluster	1 Bypass Bridge, 1 Management / Cluster / IPMI	1 Bridge, 1 Management, 1 Cluster / IPMI	1 Bridge, 1 Management, 1 Cluster / IPMI	1 Bridge, 1 Management, 1 Cluster, 1 IPMI
NICs (expandable to)	-	5 Bridges	3 Bypass Bridges, 1 Management / Cluster / IPMI	4 Bridges, 1 Management, 1 Cluster / IPMI	10 Bridges, 1 Management, 1 Cluster / IPMI	14 Bridges, 1 Management, 1 Cluster, 1 IPMI
Interface NIC Slots	Built-in NIC	Built-in NIC	1 full height occupied by default	1 half height occupied by default, 1 full height available	1 half height occupied by default, 3 full height available	2 half height occupied by default, 3 full height available
Redundant Power	No	No	No	Yes	Yes	Yes



Product Licenses

Each product is available with either a x700 or x800 license.

x700	The x700 license allows access to the monitoring, reporting, optimization (bandwidth management and QoS) features on the appliance.
x800	The x800 license includes all features of the x700 license, but it also allows access to the application acceleration features on the appliance.

Product Naming Conventions

The model number of the appliance reflects the license and features of the appliance in the format *<hardware series><software license><hardware version>-<bandwidth parameters>*. For example, a sample product name might be 6762-45. The following specifications are included in the product name:

- **Hardware series**—The hardware model number, represented by the first number in the model.
- **Software license**—The license purchased for the appliance.
- **Hardware version**—The version of the hardware platform configuration.
- **Bandwidth Optimization or Bandwidth Acceleration and Bandwidth Optimization**—The amount of bandwidth used visibility and QoS, and for acceleration.

Example 1: 6762-500

- Series: 6000
- License: x700 (visibility and QoS control)
- Hardware Version: 6.2 hardware platform
- Bandwidth: 500 Mbps for visibility and QoS

Example 2: 8862-100/500

- Series: 8000
- License: x800 (visibility, QoS control and acceleration)
- Hardware Version: 6.2 hardware platform
- Bandwidth: 100 Mbps for acceleration and 500 Mbps visibility and QoS

Safety and Compliance

Note This safety and compliance information only applies to 2x61 appliances.

"Compliances" on page 24

"Safety Guidelines" on page 24

"EMC Notice" on page 25

Compliances

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.

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.

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.

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).

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.

Chapter 2: Installation and Deployment

This chapter is designed to assist first-time users in deployment and configuration of the Exinda appliance.

Package Contents

Package contents varies slightly depending on model. In general, the following items are included:

- Exinda appliance
- Quick Start Guide
- AC power cable
- Straight CAT5 ethernet cable (usually blue)
- Cross CAT5 ethernet cable (usually red)
- Serial console cable

Pre-Installation Checklist


Before deploying the Exinda appliance, it is recommended that the following basic information is collected:

Host Name	Specify a host name to assign to the Exinda appliance.
Ethernet Negotiation Settings	Any Ethernet negotiation settings - does any equipment that the Exinda appliance will be connected to require and hard-coded Ethernet speed and/or duplex settings?
IP Address and Netmask	An available IP address and netmask is required.
Default Gateway	The default gateway is required.
DNS Server (s)	At least 1 DNS server is required, so that the Exinda appliance can resolve hostnames.
SMTP Server	An SMTP server needs to be specified, if you wish to receive e-mail notifications from the Exinda appliance.
Time Zone	The Exinda appliance's time zone should be correctly set.

These settings will be configured during the initial stages of deployment.

Deployment topologies

Typically, Exinda appliances are deployed in-line, between the core switch and the WAN/Internet router:



All models come with at least 1 hardware bypass port pair, marked LAN and WAN. These ports are designed to fail-over to pass-through mode in the event of system failure or loss of power.

Exinda appliances should be deployed with the appliance powered off. This will ensure hardware bypass is working correctly. Usually, the Exinda appliance's WAN port is cabled to the WAN/Internet router using the supplied cross-over Ethernet cable. The Exinda appliance's LAN port is cabled to the core switch using the supplied straight Ethernet cable. If your appliance has a dedicated management port, this will also need to be cabled to an internal switch using an Ethernet cable. For specific information about your model, see the supplied Quick Start Guide.

Once all Ethernet cables are in place, ensure there is still network connectivity with the Exinda appliance powered off. Then, power on the Exinda appliance. Again, ensure there is network connectivity after the appliance has booted.

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 1 millisecond, this may force neighboring equipment to renegotiate their layer 2 topology, which may take several seconds.

Discover the ways that an Exinda appliance can be integrated into your network.

- ["In-path Topologies" on page 29](#)
- ["Out of Path Topologies" on page 38](#)
- ["Cluster and High Availability" on page 343](#)

Platforms

The bridge and interface configuration differs between each Exinda platform. The table below summarizes the available bridges for each platform (without expansion cards installed).

Platform	Default Bridge (interfaces) / Bypass Support
2000	br0 (eth0 eth1) / no bypass br1 (eth2 eth3) / bypass
2060	br1 (eth1 eth2) / bypass br3 (eth3 eth4) / bypass
4000	br0 (eth0 eth1) / no bypass br1 (eth2 eth3) / bypass
4060	br10 (eth10 eth11) / bypass
4061	br10 (eth10 eth11) / bypass
5000	br0 (eth2 eth3) / bypass br1 (eth4 eth5) / bypass br2 (eth6 eth7) / bypass
6000	br0 (eth1 eth2) / fiber br1 (eth3 eth4) / bypass br2 (eth5 eth6) / bypass
6010	br0 (eth2 eth3) / bypass br1 (eth4 eth5) / bypass
6060	br10 (eth10 eth11) / bypass
7000	br0 (eth2 eth3) / bypass br1 (eth4 eth5) / bypass
8060	br10 (eth10 eth11) / bypass
10060	br10 (eth10 eth11) / bypass


IP Modes

Exinda OS (ExOS) version 6 and later allows you to configure the appliances network interfaces to suit your network topology. To configure network interfaces, navigate to **System > Network > IP Address** on the Web User Interface, Advanced mode. Ensure that you understand the target network environment before

changing settings on this page.

Notes

- Interfaces that are not enslaved to a bridge may have roles assigned e.g. Cluster, Mirror or WCCP.
- The interface used to manage the appliance will depend on the network topology. In general any interface not assigned to a role may be used.



IP Settings	
eth1	Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP Autoconf: <input type="text" value="Static"/> Addresses: <input type="text" value="172.16.1.240"/> / <input type="text" value="23"/> Comment: <input type="text"/>
eth2	Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP Autoconf: <input type="text" value="Static"/> Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>
br10 <input checked="" type="checkbox"/>	Autoconf: <input type="text" value="Static"/> Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>
br12 <input checked="" type="checkbox"/>	Autoconf: <input type="text" value="Static"/> Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>
br20 <input checked="" type="checkbox"/>	Autoconf: <input type="text" value="Static"/> Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>

Default Route:

In-path Topologies

Discover the ways that an Exinda appliance can be integrated into the path of network traffic.

- ["Main Site Internet Link Topology" on page 30](#)
- ["Main Site WAN Link Topology" on page 31](#)
- ["Distributed Branch Topology" on page 32](#)
- ["Topologies with Firewalls" on page 33](#)


- "Topologies with VPNs" on page 35
- "Multiple Link Topology" on page 37

Main Site Internet Link Topology

Main site with Internet link and potentially branch offices as well. Applications are hosted in the Main Site where branch offices connect to via the Internet.

All platforms support this topology.

This topology is used when customers need to monitor and control Internet and branch traffic to and from the main site. The Exinda can guarantee performance of critical applications such as voice, VPN and extranet; monitor Internet usage and control P2P applications.



Installation

The Exinda should be plugged 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.

It is recommended that you use a bypass capable bridge, which will provide ethernet bypass in the event of hardware failure. The bridge in use also needs to be enabled on the IP Address configuration page.

Capabilities

In this topology, the Exinda appliance can:

- Monitor all traffic utilization and all applications to/from the Internet. You can distinguish between business relevant traffic and traffic used for recreational purposes.

- Monitor usage of Internet and VPN branch traffic. e.g. How much of the link is being used by each branch network?
- Control all traffic traversing the link. Allocate some bandwidth to VPN branch offices and respective priorities for Internet applications.

Limitations

- With this topology, it is not possible to monitor and control branch traffic and their respective Internet links as each branch has direct access to the Internet.
- Application Acceleration is not possible with a single appliance.

Suggestions


- Disable direct access to the Internet for branch offices. Route all Internet traffic via the main site if possible.
- Use an Exinda appliance at each branch office to monitor and control traffic and increase WAN capacity with Exinda's Application Acceleration.

Main Site WAN Link Topology

Single site with Internet link and separate WAN link to branch offices.

All platforms support this topology.

This topology is used when customers need to monitor and control Internet and WAN traffic in the main site and WAN traffic from branch offices. The Exinda can guarantee traffic for the WAN and treat applications and users from different branch offices with different priorities.



Installation

The Exinda should be plugged 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.

It's recommended that you use a bypass capable bridge, which will provide ethernet bypass in the event of hardware failure.

The bridge in use needs to be enabled on the IP Address configuration page.

Capabilities

In this topology, the Exinda appliance can:

- Monitor all traffic utilization and all applications to the Internet. You can distinguish between business relevant traffic and traffic used for recreational purposes.
- Monitor usage of Internet and WAN traffic. E.g. How much of the link is being used by the Internet and each branch office?
- Monitor and control individual applications and users from each branch office.
- Control all traffic traversing the link. Allocate bandwidth to WAN and Internet applications.

Limitations

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


Distributed Branch Topology

A distributed topology of Exinda appliances offers the most flexible control. Such topology is also required for customers using Exinda's Application Acceleration technology.

All platforms support this topology.

Note Application Acceleration is possible with this topology. An acceleration license is required for Application Acceleration to be enabled. An acceleration license is **not** available on the 1010 platform.

This topology is used to monitor and control all nodes in a distributed branch office environment. As both WAN and Internet can be accessed directly from each office, an Exinda is used to monitor and manage the performance of each branch office.



Installation

An Exinda is required at all branch offices connecting to the WAN. The Exinda will need to be installed in in-line mode at each office.

It's recommended that you use a bypass enabled bridge, which will provide ethernet bypass in the event of hardware failure. The bridge in use needs to be enabled on the IP Address configuration page.

Capabilities

In this topology, the Exinda appliance can:


- Monitor and control all traffic to/from the Internet and WAN.
- 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.

Limitations

None - this is the most flexible topology.

Topologies with Firewalls

Firewall topologies can vary significantly. Typically customers will place the Exinda between the switch and internal interface of the firewall. This ensures that the Exinda can see all hosts on the LAN.

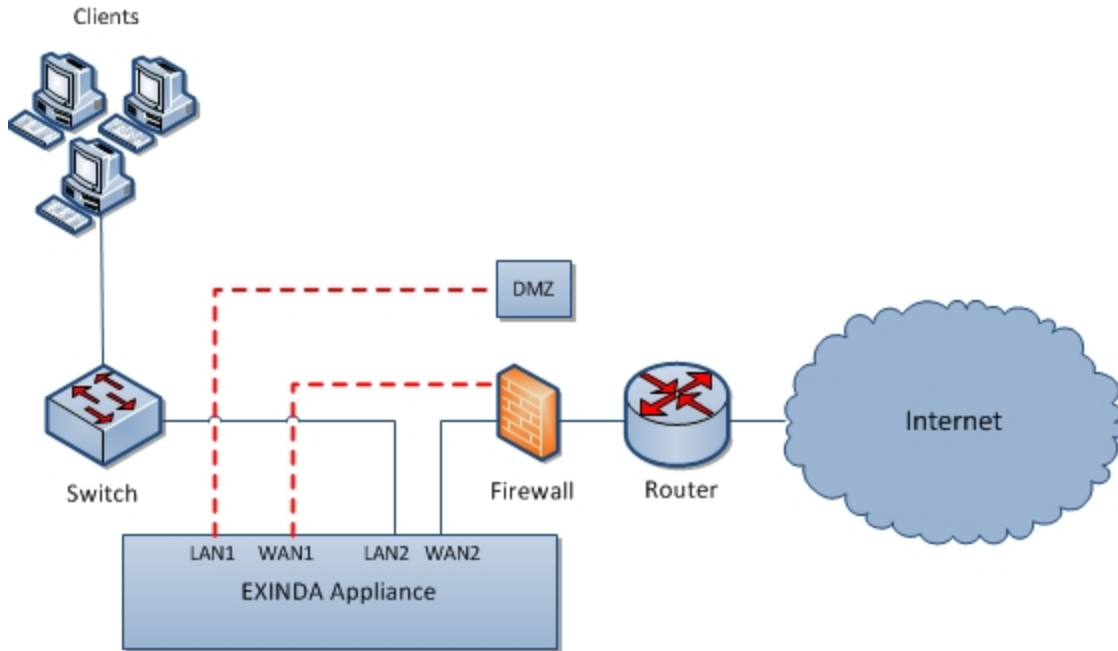


All platforms support this topology.

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), you will only see the firewall's external IP address as the source address of the monitored flows, rather than their internal addresses.

DMZ

The Exinda appliance can be deployed in-path of a DMZ, allowing for Monitoring, Optimization and Application Acceleration of traffic to/from the DMZ.




Note You will need to define a Network Object called DMZ and mark it as "Internal", so that the Exinda appliance can ignore all traffic between the local LAN and the DMZ.

Installation


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.

Topologies with VPNs

Scenario 1: Typically customers will place the Exinda between their internal LAN switch and VPN terminator. This allows for monitoring and optimization of traffic before it gets encrypted and transported across the VPN tunnel.



Scenario 2: There are scenarios where the Exinda can only plug in between the VPN terminator and the router. In this scenario only encrypted tunnel traffic will be seen by the Exinda appliance. Typically traffic of the GRE or ESP protocol will be present.



All platforms support this topology.

Installation

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.

Capabilities

In VPN scenario 2, the Exinda appliance can:

- Monitor and control any unencrypted traffic to the WAN and Internet.
- Monitor and prioritize the encrypted traffic between other VPN terminator sites. Only a single IP address will be visible per site.

Limitations

In VPN scenario 2 the Exinda appliance cannot:

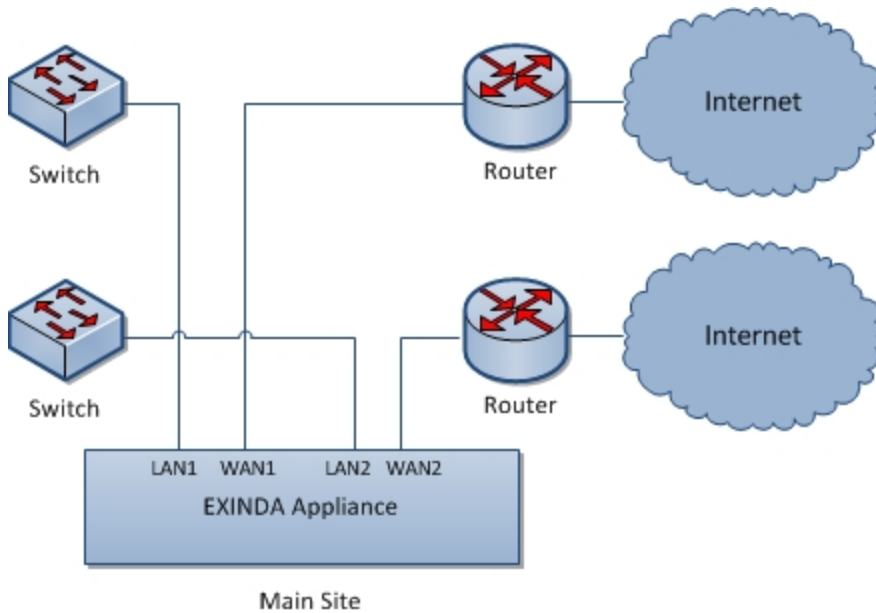
- Monitor and prioritize the encrypted traffic by application, internal hosts and servers.

Multiple Link Topology

Exinda appliances can support multiple bridges, allowing users to connect multiple links through the appliance.

All platforms support this topology, however, some platforms only have a single bypass enabled bridge, which will provide ethernet bypass in the event of hardware failure.

This topology is used when customers need to monitor and control Internet traffic to and from the main site as well as WAN traffic through a single Exinda appliance.



Installation

The Exinda should be plugged in-line 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.

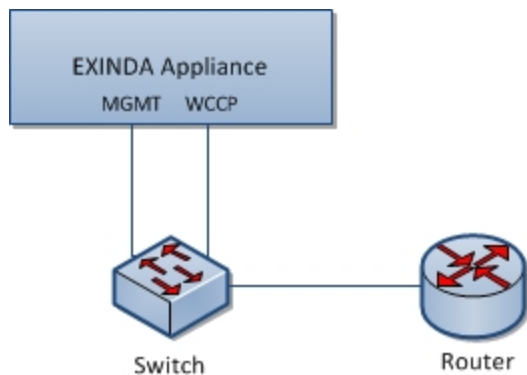
Out of Path Topologies

Discover the ways that an Exinda appliance can be integrated out-of-path.

- ["Accelerate traffic with WCCP" on page 38](#)
- ["SPAN and Mirror Port Monitoring" on page 41](#)
- ["Directing traffic with policy-based routing" on page 44](#)

Accelerate traffic with WCCP

The Exinda appliance can accelerate traffic routed using Web Cache Communication Protocol (WCCP) v2. This topology is used when customers want application acceleration, but do not wish to install the Exinda appliance in-line.



Some of the limitations of the WCCP out-of-path deployment include:

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

To configure routing using WCCP, perform the following tasks:

1. Configure router to use WCCP. Consult the documentation for 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

2. "Configure the Exinda appliance to use WCCP" on page 39.
3. "Configure internal subnets as internal network objects" on page 43.
4. "Display the state of the WCCP service" on page 41.

Configure the Exinda appliance to use WCCP

To use Web Cache Communication Protocol (WCCP) v2 to route traffic to the Exinda, configure the Exinda appliance to identify the WCCP interface, and specify any passwords required to access the router.

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

5. Assign an interface for WCCP.

- a. For a unicast configuration, set the router IP address for each WCCP v2 service.

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

For example, assign interface eth2 to accept WCCP v2 traffic with service class 10 from 192.168.0.1

```
(config) # wccp interface eth2
(config) # wccp service 10 router 192.168.0.1
```

Ensure you set the router to be 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>
```

For example, assign interface eth2 to accept WCCP v2 traffic with service class 10 from multicast address 224.1.1.1

```
(config) # wccp interface eth2
(config) # wccp service 10 group-address 224.1.1.1
```

6. 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>
```

Configure 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 AND 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 AND Packet's destination IP matches an Internal Network Object	Packet is classified as inbound traffic.
Packet's source IP matches an Internal Network Object AND	Traffic flowing from the lower IP to the higher IP is classified as outbound traffic. Traffic flowing from the higher IP to the lower IP is classified as inbound traffic.

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

1. Click **Objects > Network**.
2. To change a network object to an internal object, click **Edit** and change the location to **Internal**.
3. Click **Apply Changes**.

Display the state of the WCCP service

Display the status of the WCCP service, and verify that the connection between the Exinda and the router is active.

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

5. To display the running state of a WCCP service on the Exinda, at the `(config) #` prompt, 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.

SPAN and Mirror Port Monitoring

The Exinda appliance can operate out-of-path, or ON-LAN mode, with any hub or switch that supports port mirroring or SPAN ports.

This topology is used when customers need to monitor traffic only, without installing the Exinda in-line. The Exinda monitors and reports on all applications presented on the SPAN/mirror port. This is regularly used to perform network audits as it provides great flexibility in restricted and complex network environments.




Figure 1: 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. "Configure Mirror Port Mode" on page 42
2. "Enable Monitoring of Mirror/SPAN Traffic" on page 42
3. "Configure internal subnets as internal network objects" on page 43

After enabling Mirror/SPAN monitoring, and the appropriate Internal Network Objects have been defined, the Exinda appliance monitors traffic received on the Mirror/SPAN receiving port as if it were in-line. The only exception is the Interface Reports are blank because the Exinda appliance has no concept of packet direction at the interface level.

Configure Mirror Port Mode

Before enabling Mirror/SPAN port monitoring, you must first configure a switch port to mirror traffic to. Typically, the WAN port on the core switch is configured to mirror traffic to an unused port, which is cabled to the Exinda appliance. Alternatively, a network hub can be deployed in-path, and the Exinda appliance can be cabled directly to the hub (since 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.

Enable Monitoring of Mirror/SPAN Traffic

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

1. Click **System > Network > IP Address**.
2. To use an interface as a Mirror port, select the **Mirror** check box.

Role:	<input type="checkbox"/> Cluster	<input checked="" type="checkbox"/> Mirror	<input type="checkbox"/> WCCP
Autoconf:	IPv4: <input type="checkbox"/> DHCP	IPv6: <input type="checkbox"/> SLAAC	
eth2	Dynamic Addresses: fe80::222:19ff:fed4:8dc5/64		
Static Addresses:	<input type="text"/> / <input type="text"/>		
Comment:	<input type="text"/>		

3. 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>
```

Configure internal subnets as internal network objects

For the Exinda appliance to determine traffic direction, all internal subnets should to 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 AND 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 AND Packet's destination IP matches an Internal Network Object	Packet is classified as inbound traffic.
Packet's source IP matches an Internal Network Object AND Packet's destination IP matches an Internal Network Object	Traffic flowing from the lower IP to the higher IP is classified as outbound traffic. 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 AND 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. Traffic flowing from the higher IP to the lower IP is classified as inbound traffic.

1. Click **Objects > Network**.
2. To change a network object to an internal object, click **Edit** and change the location to **Internal**.

3. Click **Apply Changes**.

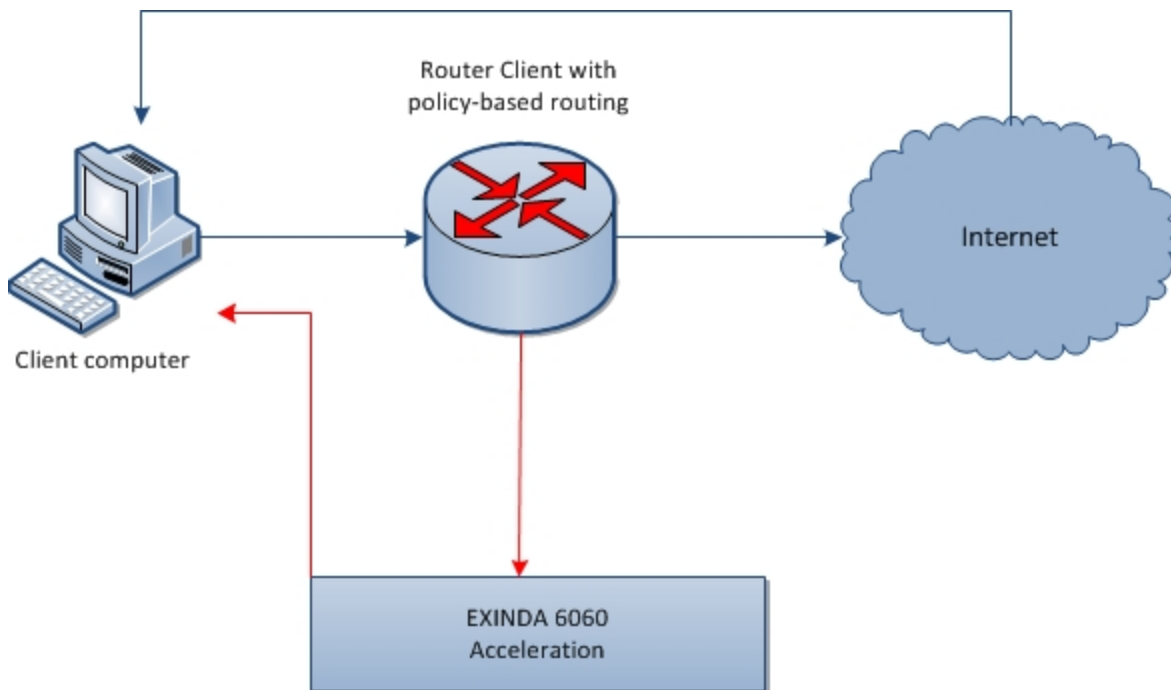
Monitor Span/Mirror Traffic

Once Mirror/SPAN 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 in-line.

The only exception is the Interface Reports will be blank, because the Exinda appliance has no concept of packet direction at the Interface level.

Directing traffic with policy-based routing

With Policy Based Routing (PBR) an Exinda appliance can be inserted in the network out-of-path, but retain in-path optimization capabilities. This is typically achieved by configuring the router with policy that determines whether traffic is sent on to the requested destination or to the Exinda appliance.



The Policy Based Routing feature offers the following benefits:

- The Exinda can be physically out-of-path, but logically in-path. This can be useful in virtual environments where it may not be desirable or possible to be in-path.
- Increased selectivity of traffic to be optimized. For example, redirect all web traffic to an Exinda.
- Increased network topology configurability. Depending on the routers capability, redirect traffic based on access-lists which may use port, source/destination address, etc.

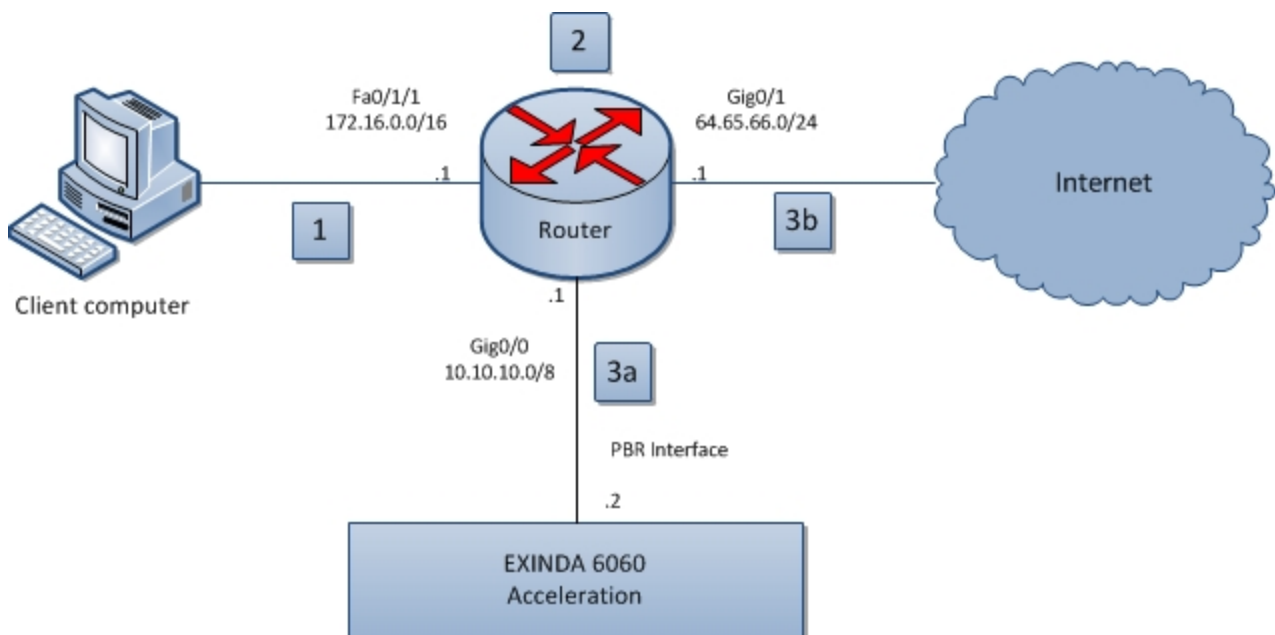
The following scenarios represent common deployment topologies in which PBR can be used with the Exinda.

- "Basic policy-based routing" on page 45
- "Policy-based routing with two subnets" on page 49
- "VRRP with PBR" on page 53
- "VRRP with PBR and VLANs" on page 58
- "VRRP with PBR and IP SLA Tracking" on page 66

Basic policy-based routing

The most basic out-of-path deployment that uses policy based routing on a router is shown below. We require three interfaces on the router.

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



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

Configure the Router for the Basic PBR topology

Specify how the router should handle traffic coming from the out-of-path Exinda appliance, the WAN, and the LAN.

Caution The IP addresses and netmasks used in this document match those used in the diagram on "Basic policy-based routing" on page 45. Use IP addresses that correspond to those in your network when configuring the router.

1. Launch the router command line interface.
2. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.
3. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.
4. 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
```
5. 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
```
6. 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
```

7. 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
```

8. 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
```

Configure the out-of-path Exinda Appliance for policy based routing

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the Exinda Web UI or [through the CLI](#).

1. Click **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, clear the BR2 checkbox. The bridge expands to display eth2 and eth3.

3. Click **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. Launch the Command Line Console.
 - a. Click **Tools > Console**.
 - b. Type the appliance username and password at the prompts.
 - c. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

- d. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

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
```

Configure the Exinda Appliance for Basic PBR topology through the CLI

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the [Exinda Web UI](#) or through the CLI.

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

5. Remove the BR10 bridge.

```
hostname (config)# no bridge BR10 enable
```

6. Set the IP address and netmask of the out-of-path Exinda appliance.

```
hostname (config)# interface eth11 ip address 10.10.10.2 /8
```

7. Identify the interface to be used for policy-based routing.

```
hostname (config)# pbr interface eth11
```

8. Set the IP address of the router.

```
hostname (config)# pbr interface ip next-hop 10.10.10.1
```

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

```
hostname (config)# no interface eth11 dhcp
```

```

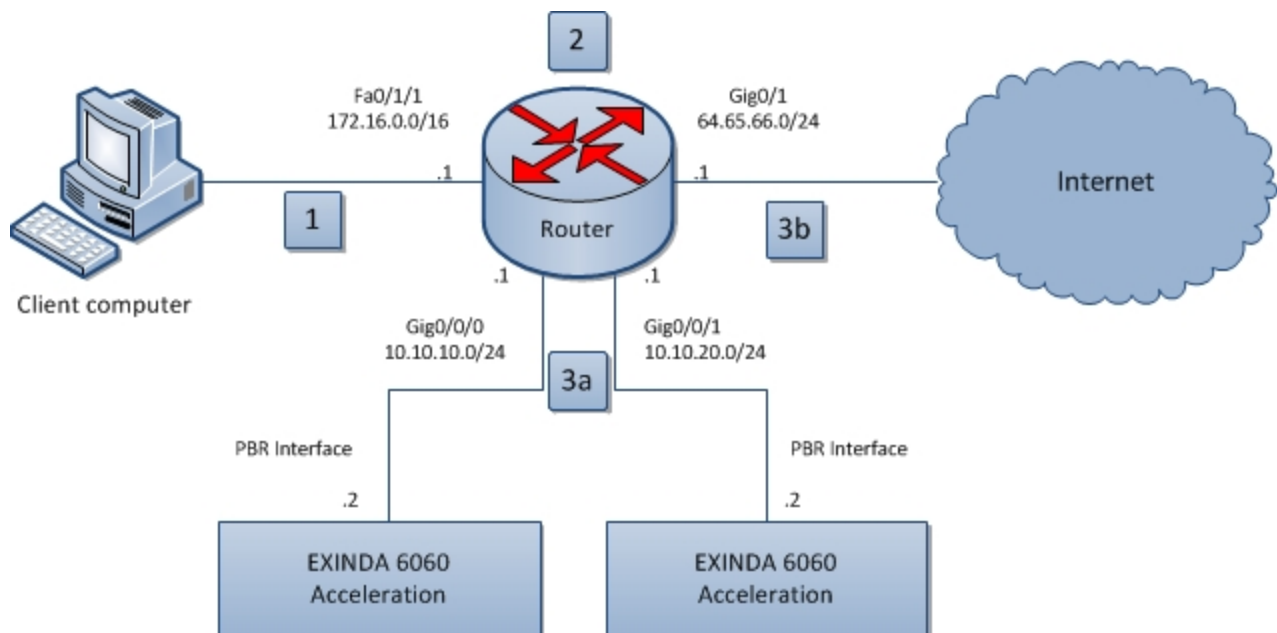
hostname (config)# interface eth11 display
hostname (config)# interface eth11 duplex auto
hostname (config)# interface eth11 mtu 1500
hostname (config)# no interface eth11 shutdown
hostname (config)# interface eth11 speed auto

```

Policy-based routing with two subnets

Policy-based routing (PBR) can be deployed on two subnets, with each subnet being serviced by its own Exinda appliance. The configuration sets rules to route traffic to another Exinda appliance should one fail.

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



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

Configure the Router for PBR with Two Subnets

Specify the behaviour of the router interfaces connecting to the the LAN, the WAN, and the out-of-path Exinda appliances.

1. Launch the router command line interface.
2. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

3. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

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 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
```

5. 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
```

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

- a. Specify the interface to configure.

```
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
```

7. 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
```

8. 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
```

9. 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
```

Configure the Exinda Appliance for PBR on Two Subnets through the Exinda Web UI

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the Exinda Web UI or [through the CLI](#).

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

- a. In a browser, enter `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 **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 appliancesoftware, clear the BR2 checkbox. The bridge expands to display eth2 and eth3.

4. Click **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. Launch the Command Line Console.

- a. Click **Tools > Console**.

- b. Type the appliance username and password at the prompts.

- c. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

- d. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

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.

Configure the Exinda Appliance for PBR on Two Subnets through the CLI

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate

settings.

Note The appliance can be configured either through the [Exinda Web UI](#) or through the CLI.

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **Tools > Console**.
3. Type the appliance username and password at the prompts.
4. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.
5. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.
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.

```
hostname (config)# interface eth11 ip address 10.10.10.2 /24
```
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 10.10.20.1
```
10. To set the parameters of eth11, type the following commands:

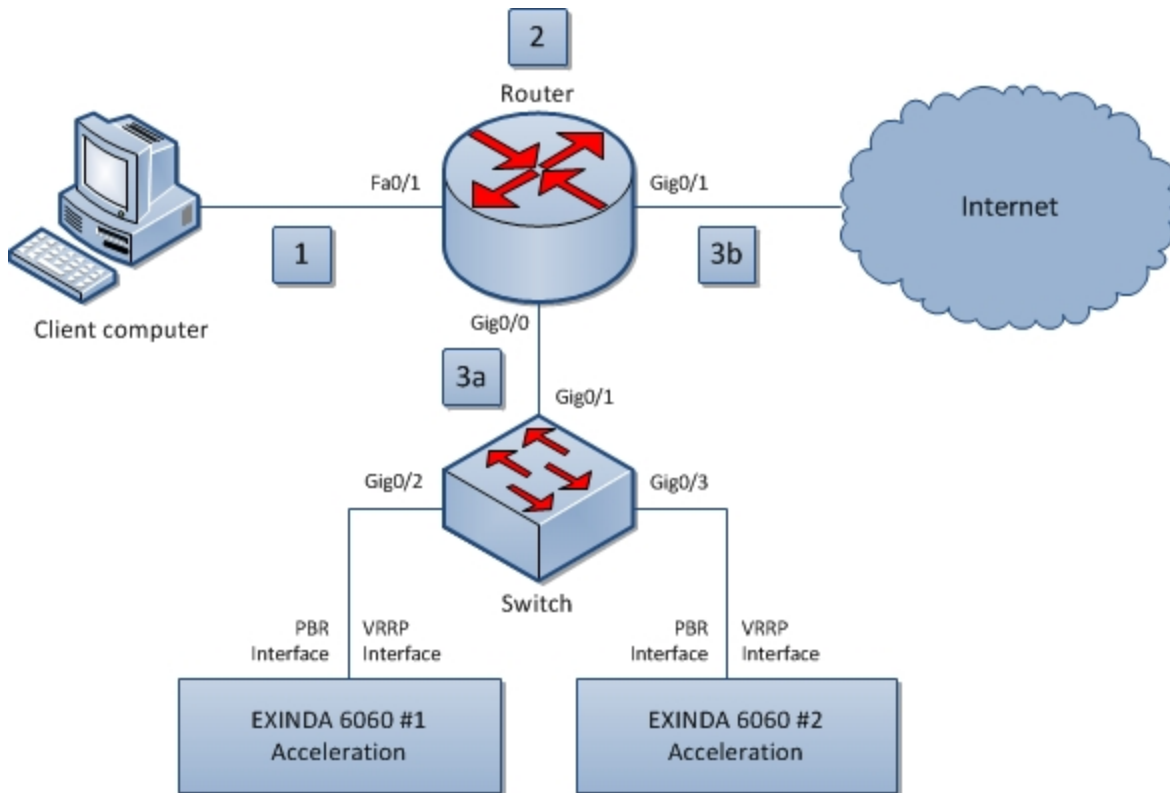
```
hostname (config)# no interface eth11 dhcp  
hostname (config)# interface eth11 display  
hostname (config)# interface eth11 duplex auto  
hostname (config)# interface eth11 mtu 1500  
hostname (config)# no interface eth11 shutdown  
hostname (config)# interface eth11 speed auto
```
11. Repeat these steps on the out-of-path Exinda appliance #2 using the appropriate IP address in Step 5.

VRRP with PBR

The Virtual Router Redundancy Protocol (VRRP) is an IP address that acts as a gateway between the router

and the appliances. The VRRP receives the traffic requests, and distributes the requests to the appliances connected to it. This allows for greater reliability and balancing traffic requests to the various appliances.

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



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

Configure the Router for VRRP with PBR

Specify the behaviour of the router interfaces connecting to the the LAN, the WAN, and the out-of-path Exinda appliances.

1. Launch the router command line interface.
2. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

3. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

4. 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
```

5. 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
```

6. 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
```

7. 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
```

8. 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
```

Configure the Exinda Appliance for VRRP with PBR

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the Exinda Web UI or [through the CLI](#).

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **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 appliancesoftware, clear the BR2 checkbox. The bridge expands to display eth2 and eth3.

4. Click **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. Launch the Command Line Console.
 - a. Click **Tools > Console**.

- b. Type the appliance username and password at the prompts.
- c. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

- d. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

13. 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
```

14. Repeat these steps on eth10 of the out-of-path Exinda appliance #2.

Configure the Exinda Appliance for VRRP with PBR through the CLI

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the [Exinda Web UI](#) or through the CLI.

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **Tools > Console**.
3. Type the appliance username and password at the prompts.
4. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.
5. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The hostname (config) # prompt is displayed.

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
```

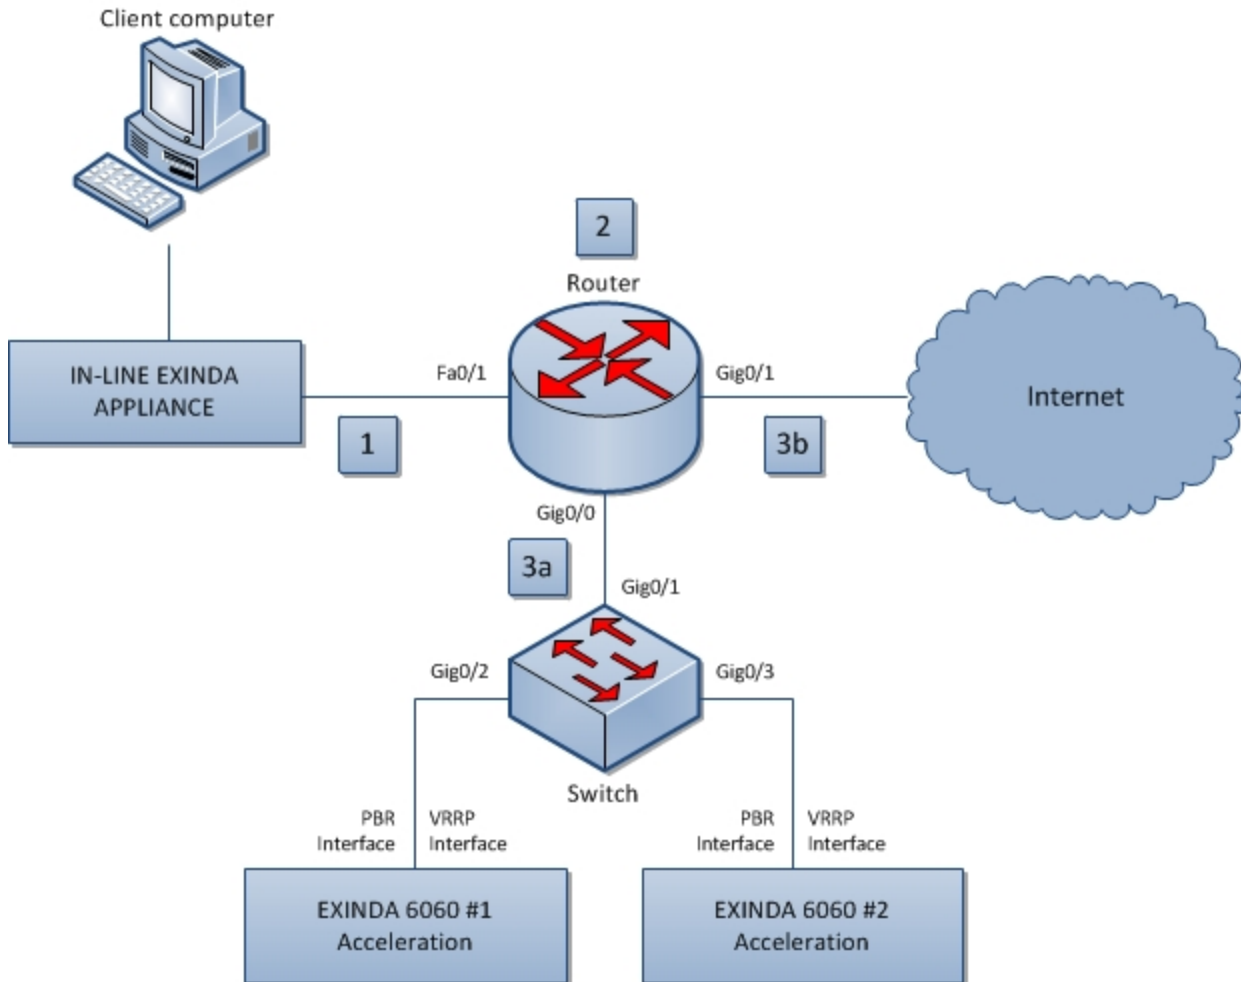
12. Repeat these steps on the out-of-path Exinda appliance #2 using the appropriate IP address in Step 5.

VRRP with PBR and VLANs

The Virtual Router Redundancy Protocol (VRRP) is an IP address that acts as a gateway between the router and the appliances. The VRRP receives the traffic requests, and distributes the requests to the appliances connected to it. This allows for greater reliability and balancing traffic requests to the various appliances.

The switch can be configured to act as a VLAN, allowing appliances that are connected to different physical switches to be grouped together. If the network has existing VLANs configured, this out-of-path Exinda appliance configuration improves latency and increases performance.

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



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

Configure the Router for 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.
2. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

3. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

4. 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
```

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. 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
```

6. 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
```

7. 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
```

8. 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
```

9. 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
```

Configure the Switch for 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.
2. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

3. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

4. Configure the interface to the router.

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

5. Configure the interface to Exinda appliance #1.

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

6. Configure the interface to Exinda appliance #2.

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

7. Set the IP addresses for the VLANs.

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

Configure the Exinda Appliance for VRRP with PBR and VLANs

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the Exinda Web UI or [through the CLI](#).

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **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 appliancesoftware, clear the BR2 checkbox. The bridge expands to display eth2 and eth3.

4. Click **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. Launch the Command Line Console.
 - a. Click **Tools > Console**.
 - b. Type the appliance username and password at the prompts.
 - c. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```


The `hostname #` prompt is displayed.
 - d. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```


The `hostname (config)#` prompt is displayed.
15. 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
```

16. Repeat these steps on eth10 of the out-of-path Exinda appliance #2.

Configure the Exinda Appliance for VRRP with PBR and VLANs through the CLI

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

Note The appliance can be configured either through the [Exinda Web UI](#) or through the CLI.

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **Tools > Console**.
3. Type the appliance username and password at the prompts.
4. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.
5. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.
6. Remove the BR10 bridge.

```
hostname (config)# no bridge BR10 enable
```
7. Identify the VLAN interface.

```
hostname (config)# vlan vlan-id 10 interface eth11
```
8. 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
```
9. Identify the interface to be used for policy-based routing.

```
hostname (config)# pbr interface eth11.10
```

10. Set the IP address of the router.

```
hostname (config)# pbr interface eth11.10 ip next-hop 10.10.10.1
```

11. 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
```

12. 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 eth11.10 vrid 50
```


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

VRRP with PBR and IP SLA Tracking

The Virtual Router Redundancy Protocol (VRRP) is an IP address that acts as a gateway between the router and the appliances. The VRRP receives the traffic requests, and distributes the requests to the appliances connected to it. This allows for greater reliability and balancing traffic requests to the various appliances.

To increase fault-tolerance, configure the router to monitor the service levels of applications (SLA). IP SLA configurations acts as a heartbeat mechanism between the router and applicable SLA'd hosts.

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



1. The client computer requests access to a location on the Internet.
2. 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.
3. Based on the results of the analysis, the request is:
 - a. sent to the switch, which evenly distributes the traffic between the connected Exinda appliances, and then back through the router to the requested destination.
 - b. sent directly to the requested destination.

Limitations

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

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

http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_pi/configuration/15-2s/iri-pbr-mult-track.html#GUID-5A5BD687-C352-4C1E-8D79-8DEAC377182D

Example configuration:

http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_pi/configuration/15-2s/iri-pbr-mult-track.html#GUID-D26DE40C-B4CA-4DCC-8A3C-4068DA11EB48

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

Configuration instructions can be found here:

http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_pi/configuration/15-2s/iri-pbr-mult-track.html#GUID-16C12D82-0E4E-4872-B9DC-D32E3EE871BA

Example configuration:

http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_pi/configuration/15-2s/iri-pbr-mult-track.html#GUID-BC718E35-0A51-49ED-8820-156B625DBA7D

Configure the Router for 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.
2. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.
3. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.
4. 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
```
5. 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
```
6. 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
```

7. 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
```

8. 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
```

9. 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
```

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

```
track 1 ip sla 5
```

```
delay down 2 up 2
```

Configure the Exinda appliances for VRRP with PBR and IP SLA Tracking

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate settings.

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **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 appliancesoftware, clear the BR2 checkbox. The bridge expands to display eth2 and eth3.

4. Click **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. Launch the Command Line Console.

- a. Click **Tools > Console**.
- b. Type the appliance username and password at the prompts.
- c. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The hostname # prompt is displayed.

- d. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The hostname (config)# prompt is displayed.

13. 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
```

14. Repeat these steps on eth10 of the out-of-path Exinda appliance #2.

Configure the Exinda appliances for VRRP with PBR and IP SLA Tracking through the CLI

To use policy-based routing, the interfaces on the Exinda appliance must be configured with the appropriate

settings.

1. On the out-of-path Exinda appliance # 1, launch the Exinda Web UI.
 - a. In a browser, enter `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 **Tools > Console**.
3. Type the appliance username and password at the prompts.
4. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

5. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

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 IP address in Step 5.

Cluster and High Availability

Clustering allows multiple Exinda appliances to operate as if they were a single appliance. This allows for seamless deployment into High Availability and Load Balanced environments. A typical deployment topology is illustrated below.



In this example, there are two physical links. An Exinda appliance is deployed between each switch and router, and a cable is connected between the two appliances for synchronization.

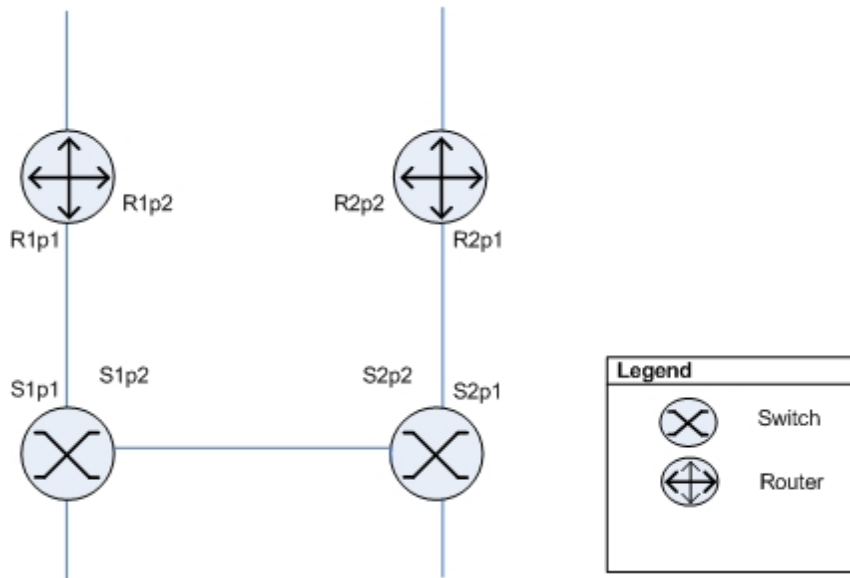
The appliances share configuration, monitoring information, and optimizer and acceleration policies, as if they were a single appliance.

Refer to the following topics for example topologies:


- ["Redundancy through multiple Exinda appliances" on page 73](#)
- ["Load balancing and fail-over with multiple Exinda appliances" on page 75](#)
- ["High availability mode" on page 77](#)

Redundancy through multiple Exinda appliances

The clustering feature allows two Exinda appliances to be connected in a redundant topology.



With the Exinda appliances installed the above topology will appear as below:



The two appliances are directly connected to each other. Both appliances will capture the same data. The appliance that receives the data directly will forward the traffic to the other appliance which will monitor it the same way. However, the copied traffic will not be forwarded onto the LAN.

Exinda's Clustering/HA framework is also responsible for automatically synchronizing configuration settings between the two appliances.

All platforms support this topology.

Installation

1. On each Exinda, 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 1 minute power up Exinda 2.
4. Connect Exinda 1 LAN into switch 1 (S1p1).
5. Connect Exinda 1 WAN into router 1 (R1p1).
6. Connect Exinda 2 LAN into switch 2 (S2p1).
7. Connect Exinda 2 WAN into router 2 (R2p1).
8. Connect Exinda 1 management interface into switch 2 (S2p2)
9. Connect Exinda 2 management interface into switch 1 (S1p2)

Capabilities

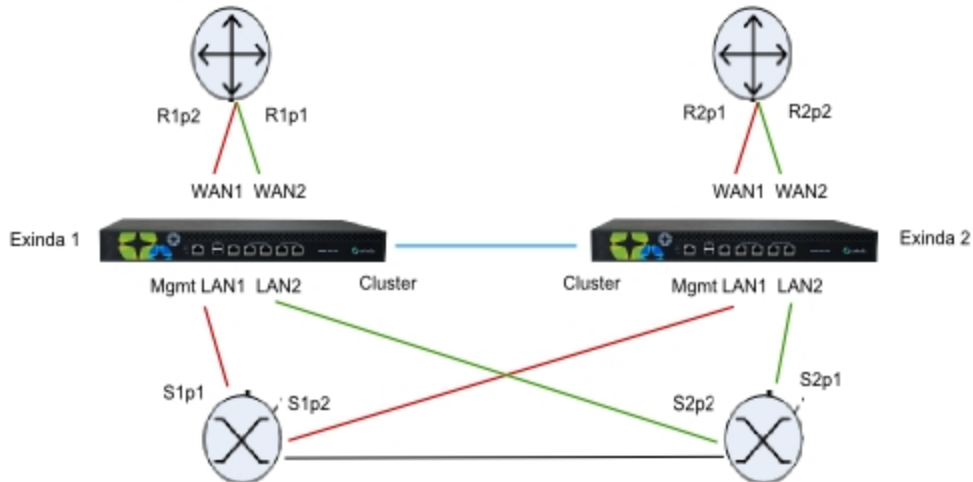
- Monitoring of both links.
- Optimization of both links.
- Redundancy of Exinda appliances.

Load balancing and fail-over with multiple Exinda appliances

Similar to the previous topology but in this case the routers are configured for load balancing. Both links in this topology act as fail-over and load balancing.



With Exinda appliances installed the above topology will appear as below:



In this topology both Exinda appliances are connected to both routers. As with the ["Redundancy through multiple Exinda appliances" on page 73](#) case, direct traffic reaching one appliance is copied to the second appliance for monitoring and optimization, but is not forwarded on.

Platforms that support this topology include the 4060¹, 4061¹, 5000, 6010, 6060¹, 7000 and 10060¹.

¹ Network expansion modules are required.

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 LAN2 into switch 1 (S1p2).
5. Connect Exinda 1 WAN2 into router 1 (R2p2).
6. Connect Exinda 1 LAN1 into switch 1 (S1p1).
7. Connect Exinda 1 WAN1 into router 1 (R1p1).
8. Connect Exinda 2 LAN2 into switch 2 (S2p1).
9. Connect Exinda 2 WAN2 into router 2 (R2p1).
10. Connect Exinda 2 LAN1 into switch 2 (S2p2).
11. Connect Exinda 2 WAN1 into router 2 (R1p2).
12. Connect Exinda 1 MGMT into switch 2.
13. Connect Exinda 2 MGMT into switch 1.

Capabilities


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

High availability mode

When Router Redundancy is not present but you would still like to configure the Exinda solution in High Availability mode, use the configuration below.



With Exinda appliances installed the above topology will appear as below:



In this topology, both Exinda appliances are connected via a WAN switch. As with the "[Redundancy through multiple Exinda appliances](#)" on page 73 case, direct traffic reaching one appliance is copied to the second appliance for monitoring, but is not forwarded.

Note Your WAN switch and LAN switch must support the Spanning Tree Protocol (STP). Configure STP with S2p1 higher priority than S2p2. If the link at S2p1 goes down (e.g. Exinda 1 loses power) then the switch will enable S2p2. Exinda1 should configure NoLink as the bridge failover option, Exinda 2 should configure Bypass.

Active Path: S2p1 to S1p1
Standby Path: S2p2 to S1p2

All platforms support this topology.

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 1 (S1p1).
5. Connect Exinda 1 WAN into switch 2 (S2p1).
6. Connect Exinda 2 LAN into switch 1 (S1p2).
7. Connect Exinda 2 WAN into switch 2 (S2p2).
8. Connect the management interface of Exinda 1 into switch 1.
9. Connect the management interface of Exinda 2 into switch 1.
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 S2p2 should have the highest STP priority.

Capabilities

- Monitoring data available on both Exinda appliances.
- Optimization available via Exinda 1 or Exinda 2.
- Redundancy of Exinda appliances.

Cluster Interfaces

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 cluster internal traffic. 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	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

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, each appliance's cluster interface must be connected to a single, dedicated switch - such that each appliance can communicate with every other appliance without requiring a route (must be on the same Layer 2 LAN segment).

Cluster Failover

In the event that a node in the cluster fails, is rebooted or powered off, it will enter bypass mode and traffic will pass through unaffected. When the appliance is brought back online, the node will be updated with the latest configuration settings from the Cluster Master and normal operation will continue. Monitoring and reporting information during the downtime will not be synchronized retrospectively.

In the event that the Cluster Master fails, is rebooted or powered off, a new Cluster Master will be automatically elected and the offline node will be treated as a regular offline node. When it is brought back online, it won't necessarily become the Cluster Master again.

Cluster Terminology

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's, cluster interface for the purposes of communicating with other nodes in the cluster.

Cluster Master—The node responsible for synchronizing configuration changes with all other cluster nodes. Configuration changes should only be made from the cluster master.

ID—The node's cluster assigned unique identifier.

Management IP—The clusters 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 node within the cluster (master or standby).

State—The node state (online or offline)

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. ["Add Exinda appliances to the cluster" on page 81](#)
2. ["Specify what data is synchronized between cluster members" on page 82](#)

After the appliances have been configured, all appliances in the cluster can be monitored. See ["View the status of all members of the cluster" on page 83](#).

Add Exinda appliances to the cluster

Configure the appliances with an internal IP address used within the cluster, as well as the IP address of the cluster master.

1. Click **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.

Cluster configuration 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
```

This command will need to be run on each node in the cluster, and each with a unique Cluster Internal IP.

2. Configure, the Cluster External IP. This command should be executed on all cluster nodes.

```
(config) # cluster master interface eth1
(config) # cluster master address vip 192.168.0.160 /24
```

The same Cluster External IP should be configured on each cluster node.

3. Enable the cluster.

```
(config) # cluster enable
```

4. As with the Web UI, the role of the node currently logged into will be displayed in the CLI prompt as shown below. Configuration changes should only be made on the Cluster Master node.

```
exinda-091cf4 [exinda-cluster: master] (config) #
```

5. It is possible to view the status of the 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
```

Specify what data is synchronized between cluster members

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

Similarly, 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 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.

The following CLI commands can be used to control how data is synchronized between cluster members:

```
(config)# [no] cluster sync {all|acceleration|monitor|optimizer}
```

`all` - Acceleration, monitor and optimizer data are synchronized. This is disabled by default.

`acceleration` - Synchronize acceleration data only

`monitor` - Synchronize monitor data only

`optimizer` - Synchronize optimizer data only

View the status of all members of the cluster

Identify the roles of each appliances in the cluster, and see basic statistics about the appliances in the Exinda Web UI.

1. Click **System > Maintenance > Clustering**.

All appliances in the cluster are displayed.

Clustering State									
Host ID	External IPv4 Address	Internal IPv4 Address	Status	Role	Uptime	Version	Memory	Operation	
0024e83dcaed	192.168.0.161	192.168.1.1	✔	Master	1h 1m 23s	6.1.0.16836	2050.5MB	Shutdown	Reboot
bc305bd453a8	192.168.0.162	192.168.1.2	✔	Standby	1h 1m 25s	6.1.0.16836	2050.5MB	Shutdown	Reboot

2. To identify the cluster master, the role is displayed in the list of all appliances.

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 as shown below.



Virtual Appliances

Exinda's range of UPM appliances are 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. If the Exinda software detects that it is running under a hypervisor, certain optimizations are automatically enabled to ensure maximum performance.

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

- [VMware vSphere](#) (ESX and ESXi) (4.1 and later)
- [Citrix XenServer](#) (5.5 and later)

This guide explains how to successfully download and deploy Exinda Virtual Appliances. The following documentation may also be useful:

- Exinda Topologies Guide
- Exinda Virtualization Deployment Guide
- Exinda Storage How to Guide
- Exinda User Manual

Note

Particular attention should be paid to the [Sizing](#) and [Licensing](#) requirements of Virtual Appliances. These are covered later in this guide.

Virtual Appliance Deployment Options

Virtual Appliances can be deployed in all the same ways hardware appliances can. There are generally 2 typical deployment topologies, inline and out-of-path.

Inline deployments involve 1 or more LAN/WAN port pairs bridged together at layer 2. Traffic must be directed through the bridge in order to be Monitored and Optimized.

Using n hardware appliance as an example, here is what an inline deployment looks like:



The challenge in a virtual environment is how to pass traffic through the bridge. There are several options:


- Dedicate 2 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 such that traffic is directed through virtual LAN/WAN ports.

How this is best achieved largely depends on the virtual environment and which hypervisor technology is used.

Out-of-path deployments are typically used in SPAN port mirroring, WCCPv2, HA and Clustering scenarios. Using a hardware appliance as an example, here is what an out-of path deployment looks like:



This is quite straight forward to setup in a virtual environment, you need one NIC for Management and another NIC for Auxiliary (which is how Virtual Appliances are configured by default).

Note

For more information, consult the [Exinda Topologies Guide](#) and the [Exinda Virtualization Deployment Guide](#).

Exinda supports a virtualization solution for most of the well-known hypervisor solutions, including:

- VMware ESX/ESXi versions 4.0, 4.1, and 5.0
- Microsoft Hyper-V Windows 2008
- Citrix XenServer versions 5.5, 5.6, and 6.0
- KVM

This document provides an overview of the supported Exinda deployments for its virtualization solution, and includes the following information:

- Diagrams of each supported deployment,
- Lists of supported software versions,
- Recommendations for sizing (CPU, Memory, and storage),
- Limitations for each hypervisors, where applicable,
- Reference links for supporting material.

Understanding virtual appliance resource requirements

The following sizing guide can be used to estimate the amount of resources required for each Virtual Appliance depending on the WAN bandwidth.

Visibility and Control	Visibility, Control and Optimization	Virtual CPUs	Minimum RAM	Storage
Up to 10 Mbps	Up to 2 Mbps	1-2	1GB	~250GB
Up to 45 Mbps	Up to 5 Mbps	2	2GB	250GB-500GB
Up to 100 Mbps	Up to 10 Mbps	2	4GB	~500GB
Up to 250 Mbps	Up to 20 Mbps	2-4	6GB	500GB-1TB *
Up to 500 Mbps	Up to 45 Mbps	4	8GB	~1TB *

* Storage at higher throughput will require higher disk I/O bandwidth, so the underlying storage should be RAID based, ideally RAID 10.

This is a guide only and there are several factors that may mean more or less 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.

Sizing Guidance

To see the recommended (and minimum where noted) hardware configuration requirements for CPU, memory, and disk storage for the available models, and determine which is right for your deployment, refer to the information below or to the Sizing Guide available in the [Exinda Virtual Appliance product sheet](#).

Exinda supports the following products and scalability:

- X700 Series – 2Mbps to 500Mbps for Traffic Shaping and Visibility
- X800 Series – 1Mbps to 45Mbps for WAN Optimization, Traffic Shaping and Visibility

Note The recommended CPU is the minimum CPU(s) 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. By increasing the virtual file sizes for each of the three caches mentioned, you can greatly improve the performance of your Exinda virtual appliance.

For Edge Cache and SMB1 Cache there is no limit on the size of the file created in the external storage. Use common sense when creating the file sizes. Exinda recommends 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 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

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 (EX-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 EX-V you will add the IOPS for each service:

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


Example: Model - 2850-6	IOPS
Edge Cache IOPS	30
Monitoring IOPS	140
Average Optimization IOPS	200
Total IOPS	370

1. On the Custom Performance Chart for the EX-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.



Exinda model 2750

Licensed Bandwidth (Full Duplex)	2	10	20
Max Concurrent Flows (K)	32K	32K	32K
Max L7 New Connection Rate	30	30	30
Reports	4	4	4
SLAs	10	10	20
APS Objects	20	20	20
Policies	64	128	128
Edge Cache Max Throughput (Mbps)	4	4	4
Edge Cache Requests per Second	100	100	100
CPU(s)*	1 x 2Ghz	2 x 2Ghz	2 x 2Ghz
Minimum Storage (GB)	160	160	160
Minimum Memory (GB)	2	2	4
EC-IOPS	30	30	30
Monitoring-IOPS	140	140	140

Exinda model 4750

Licensed	2	10	15	20	45	100	155	250
----------	---	----	----	----	----	-----	-----	-----

Bandwidth (Full Duplex)								
Max Concurrent Flows (K)	64K	128K	128K	256K	256K	384K	512K	768K
Max L7 New Connection Rate	300	300	300	300	300	300	300	300
Reports	4	6	8	10	16	20	20	20
SLAs	70	100	100	120	120	150	150	150
APS Objects	150	150	150	150	150	150	150	150
Policies	128	256	256	384	384	512	512	512
Edge Cache Max Throughput (Mbps)	20	20	20	20	20	20	20	20
Edge Cache Requests per Second	1500	1500	1500	1500	1500	1500	1500	1500
CPU(s)*	2 x 2Ghz	2 x 2Ghz	2 x 2Ghz	2 x 2Ghz	2 x 2.4Ghz	2 x 2.4Ghz	4 x 2.4Ghz	4 x 2.4Ghz
Minimum Storage (GB)	250	250	250-500	250-500	500	500	500-1000	500-1000
Minimum Memory (GB)	2	4	4	6	6	6	8	8
EC-IOPS	50	50	50	50	50	50	50	50

Monitoring-IOPS	140	140	140	140	150	150	150	150

Exinda model 6750

Licensed Bandwidth (Full Duplex)	45	100	250	500
Max Concurrent Flows (K)	512K	512K	768K	1024K
Max L7 New Connection Rate	1K	1K	1K	1K
Reports	20	20	40	60
SLAs	100	100	250	250
APS Objects	200	200	200	200
Policies	1024	1024	1536	2048
Edge Cache Max Throughput (Mbps)	125	125	125	125
Edge Cache Requests per Second	2K	2K	2K	2K
CPU(s)*	2 x 2.4Ghz	2 x 2.4Ghz	4 x 2.4Ghz	4 x 2.4Ghz
Minimum Storage (GB)	500	500	500-1000	500-1000
Minimum Memory (GB)	4	6	8	10
EC-IOPS	70	70	70	70
Monitoring-IOPS	150	150	150	150

Exinda model 8750

Licensed Bandwidth (Full Duplex)	100	250	500
Max Concurrent Flows (K)	2048K	2048K	2048K

Max L7 New Connection Rate	6K	6K	6K
Reports	80	80	80
SLAs	250	250	250
APS Objects	250	250	250
Policies	2048	2048	2048
Edge Cache Max Throughput (Mbps)	175	175	175
Edge Cache Requests per Second	2.5K	2.5K	2.5K
CPU(s)*	1 x 2Ghz	2 x 2Ghz	2 x 2Ghz
Minimum Storage (GB)	1500	1500	1500
Minimum Memory (GB)	6	8	10
EC-IOPS	80	80	80
Monitoring-IOPS	150	150	150

Exinda Model 2850

Licensed Bandwidth (Full Duplex)	1	2	6
Max Concurrent Flows (K)	32K	32K	32K
Max L7 New Connection Rate	30	30	30
Reports	4	4	4
SLAs	10	10	20
APS Objects	20	20	20
Policies	64	128	128
Edge Cache Max Throughput (Mbps)	4	4	4
Edge Cache Requests per Second	100	100	100
Optimized Connections	100	150	250

CPU(s)*	1 x 2Ghz	2 x 2Ghz	2 x 2Ghz
Minimum Storage (GB)	160	160	160
Minimum Memory (GB)	2	2	4
Average - IOPS	80	90	200

Exinda model 4850

Licensed Bandwidth (Full Duplex)	1	2	3	6	10	20
Max Concurrent Flows (K)	32K	64K	64K	128K	256K	384K
Max L7 New Connection Rate	300	300	300	300	300	300
Reports	4	6	6	8	10	12
SLAs	40	60	60	80	100	120
APS Objects	150	150	150	150	150	150
Policies	128	128	128	256	256	384
Edge Cache Max Throughput (Mbps)	20	20	20	20	20	20
Edge Cache Requests per Second	1500	1500	1500	1500	1500	1500
Optimized Connections	1000	1500	1800	2000	2500	3000
CPU(s)*	2 x 2Ghz	2 x 2Ghz	2 x 2Ghz	2 x 2Ghz	2 x 2Ghz	2 x 2Ghz
Minimum Storage (GB)	250	250	250	250	250-500	250-500
Minimum Memory (GB)	2	2	2	2	4	6
Average - IOPS	80	90	120	200	220	450

Exinda model 6850

Licensed Bandwidth (Full Duplex)	10	20	45
Max Concurrent Flows (K)	256K	384K	384K
Max L7 New Connection Rate	1K	1K	1K
Reports	20	30	40
SLAs	100	100	250
APS Objects	200	200	200
Policies	1045	1536	1536
Edge Cache Max Throughput (Mbps)	125	125	125
Edge Cache Requests per Second	2K	2K	2K
Optimized Connections	5K	6K	7K
CPU(s)*	2 x 2.4Ghz	2 x 2.4Ghz	2 x 2.4Ghz
Minimum Storage (GB)	500	500	500
Minimum Memory (GB)	4	6	6
Average - IOPS	220	400	450

Run the Virtual Appliance on VMware vSphere (ESX and ESXi)

The following sections describe how to deploy Exinda Virtual Appliance as well as customize the virtual hardware to suit your requirements.

Exinda Virtual Appliance are available for VMware ESX/ESXi hypervisors.

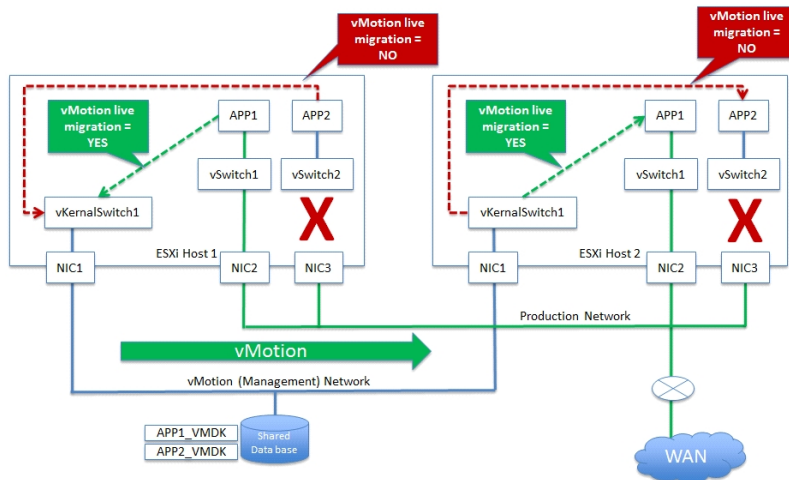
Note VMware ESX/ESXi 4.1 (or later) is required.

1. ["Install the Virtual Appliance on VMware" on page 94.](#)
2. ["Install the Silicom Bypass Driver" on page 105.](#)
3. ["Modify the VMware Virtual Machine Configuration" on page 95.](#)
4. ["Start the VMware Virtual Appliance" on page 104.](#)

Understanding how VMotion works

For isolated virtual applications on the EX-V LAN port for inline mode, the VMware vMotion feature will not

work. It is a requirement from VMware that any virtual switch must be mapped to a physical NIC, to an external network. Below is a brief illustration of the process.



- There are two types of virtual switches in the ESX/ESXi hypervisor: vKernal Switch and vSwitch. The vKernal Switch is used by the hypervisor exclusively. vKernal is the bare metal hypervisor, and provides core and memory allocation, disk and network virtualization, and 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 vKernal Switch mapped to the NIC attached to the storage.
- There are two networks:
 1. Management network where vMotion moves workloads between ESXi hosts, and
 2. Production network where the applications are accessed by the users.
- The vKernalSwitch1 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 and it 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 EX-V has at least one vSwitch mapped to it, and the vSwitch is not mapped to an external NIC interface, vMotion will not work for EX-V or workloads isolated behind it.

Install the Virtual Appliance on VMware

1. Locate the latest release of the Exinda VMware Virtual Appliance from the Firmware Downloads section of the Exinda [website](#).

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 OVF 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 appliance. If prompted, choose the location to deploy the virtual appliance. Click **Next**.
8. 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. 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.
 - 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 inline, [add additional NICs](#).
10. Click **Next**.
11. Review the deployment settings, and click **Finish**.

Review the information then click Finish to deploy the Virtual Appliance. Once it's deployed, review the following sections on [Custom Settings](#) and adding extra [NICs](#) and [Storage](#).

Modify the VMware Virtual Machine Configuration

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

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

- ["Increase the Number of CPUs on the VMware Virtual Machine" on page 95](#)
- ["Add Network Interfaces to the VMware Virtual Appliance" on page 97](#)
- ["Increase the Amount of RAM on the VMware Virtual Machine" on page 96](#)
- ["Add Storage to the VMware Virtual Machine" on page 100](#)

Increase the Number of CPUs on the VMware Virtual Machine

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

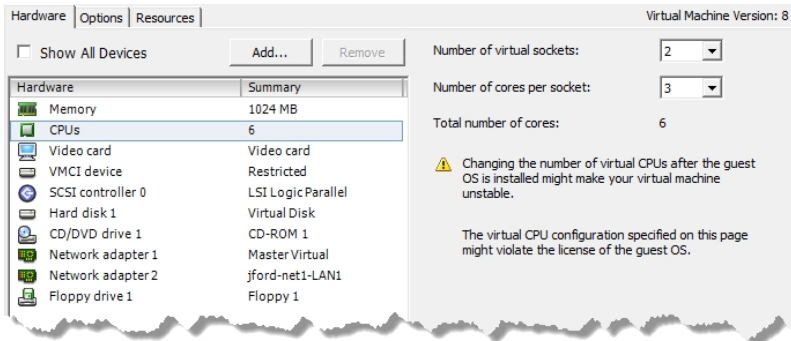
suit your requirements.

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

1. Ensure the virtual appliance is powered off.
2. Open the **VMware vSphere Client**.
3. Right-click the Exinda Virtual Appliance, and select **Edit Settings**.
4. On the **Hardware** tab, select **CPUs**.
5. Select the **Number of virtual sockets**.
6. 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.

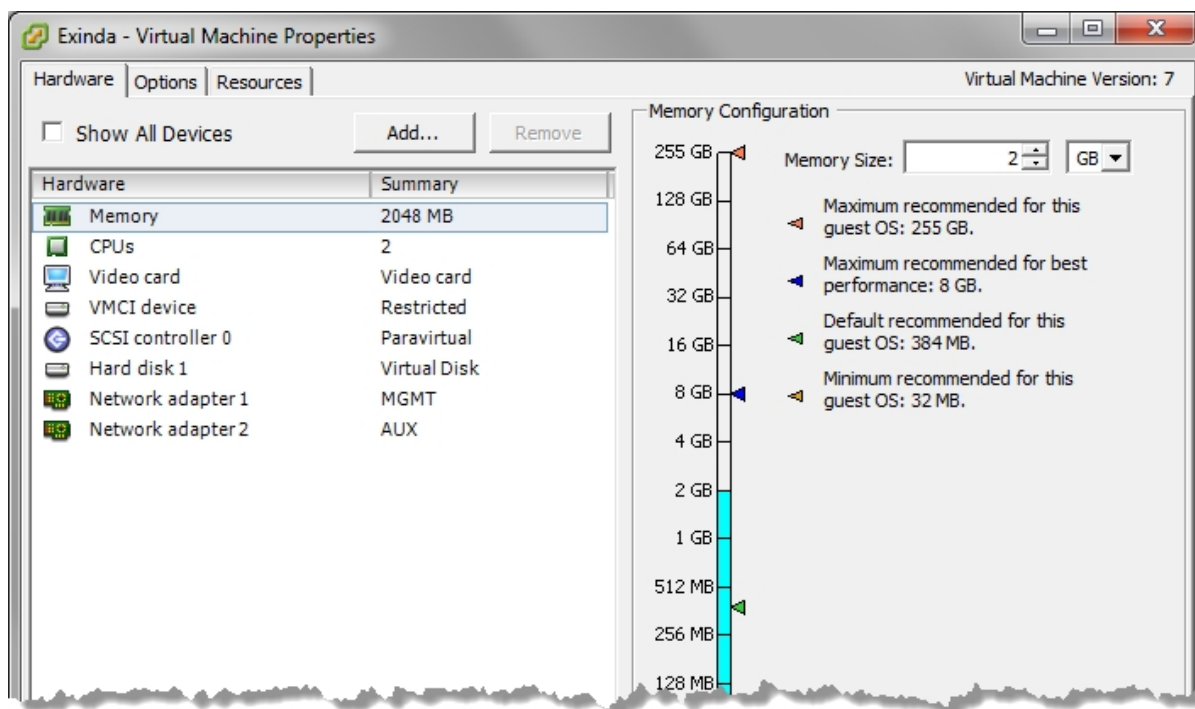


7. Click **OK**.

Increase the Amount of RAM on the VMware Virtual Machine

By default, all Virtual Appliances come configured with 2GB of RAM. Increase the amount of RAM to suit your requirements.

1. Ensure the virtual appliance is powered off.
2. Open the **VMware vSphere Client**.
3. Right-click the Exinda Virtual Appliance, and select **Edit Settings**.
4. On the **Hardware** tab, select **Memory**.
5. Select the desired **Memory Size**.



6. Click **OK**.

Add Network Interfaces to the VMware Virtual Appliance

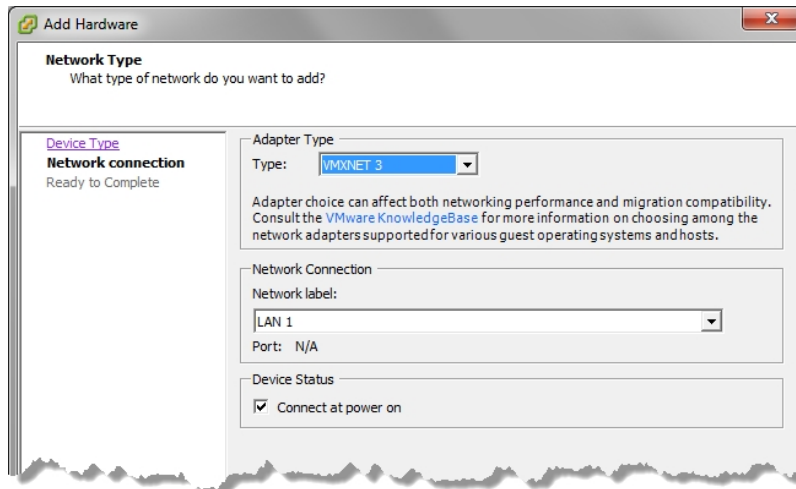
By default, all Exinda virtual appliances come with two network interface cards (NICs). The first NIC is for managing the virtual appliance through the Management Interface, and the second NIC is for managing specific configurations such as high availability, clustering, or out-of-path deployments. The second NIC is referred to as the Auxiliary Interface.

When placing the virtual appliance inline, add two extra NICs to be used as LAN and WAN ports. The two additional NICs are bridged and allow the Virtual Appliance to be placed inline. Alternatively, you can convert the two default NICs into a bridge so the Management Interface becomes a LAN Interface, and the Auxiliary Interface becomes a WAN Interface. See ["Convert NICs into a Bridge" on page 98](#).

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. Ensure the virtual appliance is powered off.
2. Open the **VMware vSphere Client**.
3. Right-click the Exinda Virtual Appliance, and select **Properties**.
4. Switch to the **Hardware** tab.
5. Click **Add**.
6. From the Device Type list, select **Ethernet Adaptor** and click **Next**.
7. In the Adapter Type list, select **VMXNET 3**.

8. Select the network to map the NIC to.




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

Convert NICs into a Bridge

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

1. Start the virtual appliance.
2. Navigate to **System > Network > IP Address**.
3. To bridge the two NICs together, select **br0**.

4. To manage the Virtual Appliance, in the **IPv4** or **IPv6** field specify an IP Address for the bridge.



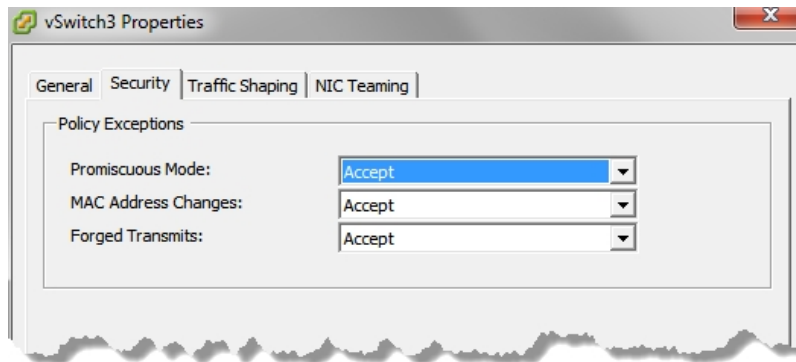
Interface Settings	
Autoconf:	IPv4: <input checked="" type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC
Dynamic Addresses:	192.168.0.221/24 fe80::20c:29ff:feaa:b541/64
br0 <input checked="" type="checkbox"/>	Static Addresses: <input type="text"/> / <input type="text"/>
Comment:	<input type="text"/>
Gateway Settings	
IPv4:	<input type="text"/>
IPv6:	<input type="text"/>

Note For inline deployments to work correctly under VMware, the virtual switches need to allow promiscuous mode. See "Allow Ports to Accept and Bridge Packets" on page 99 for more information.

Allow Ports to Accept and Bridge Packets

Any VMware virtual NIC used to deploy the virtual appliance inline 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. Ensure the virtual appliance is powered off.
2. Open the **VMware vSphere Client**.
3. Select the ESXi server, and switch to the **Configuration** tab.
4. In the list of Hardware configuration options, select **Networking**.
5. Beside the switch name, click **Properties**.
6. In the switch properties, switch to the **Security** tab.



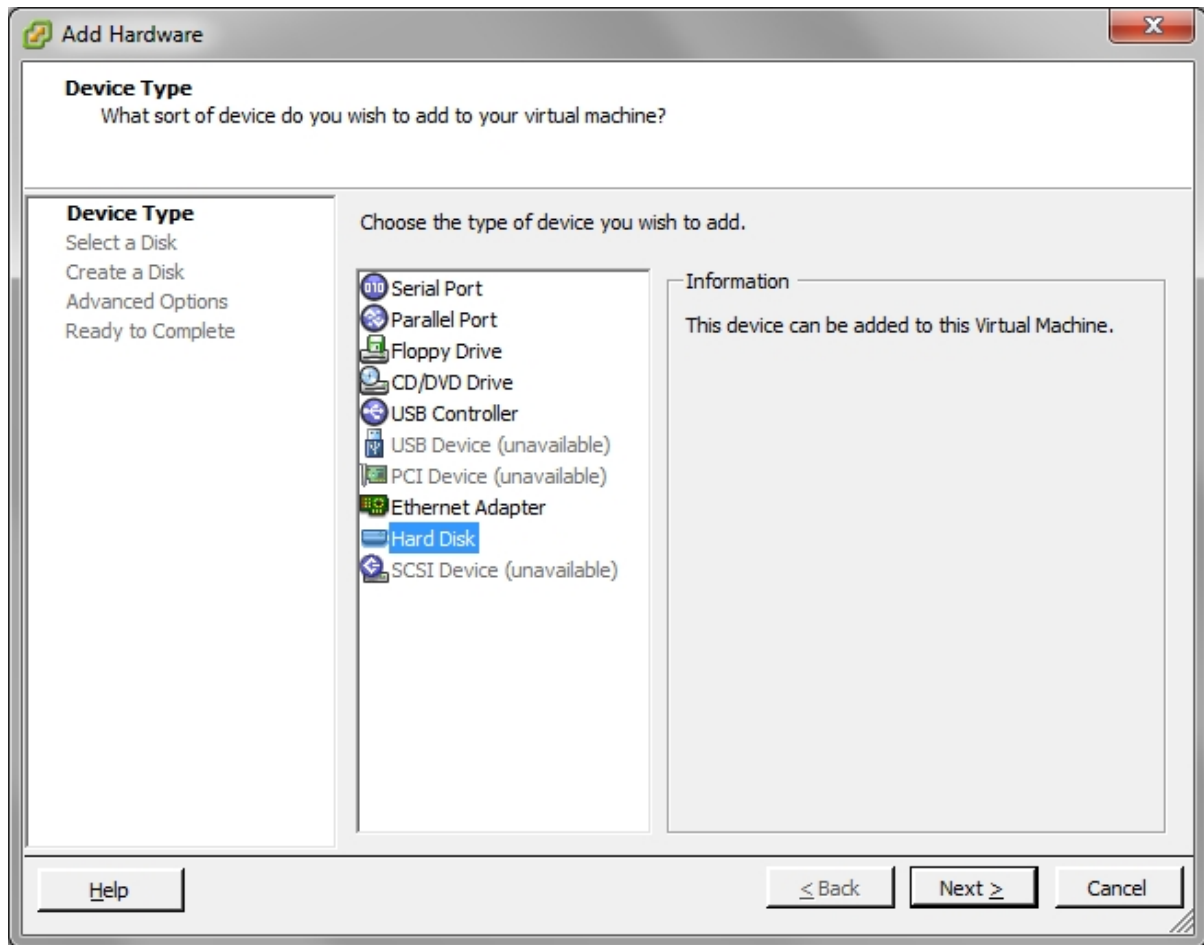
7. Set **Promiscuous Mode** to **Accept**.
8. Click **OK**.
9. Repeat these steps for each virtual switch that is attached to a NIC used in an inline deployment.

Add Storage to the VMware Virtual Machine

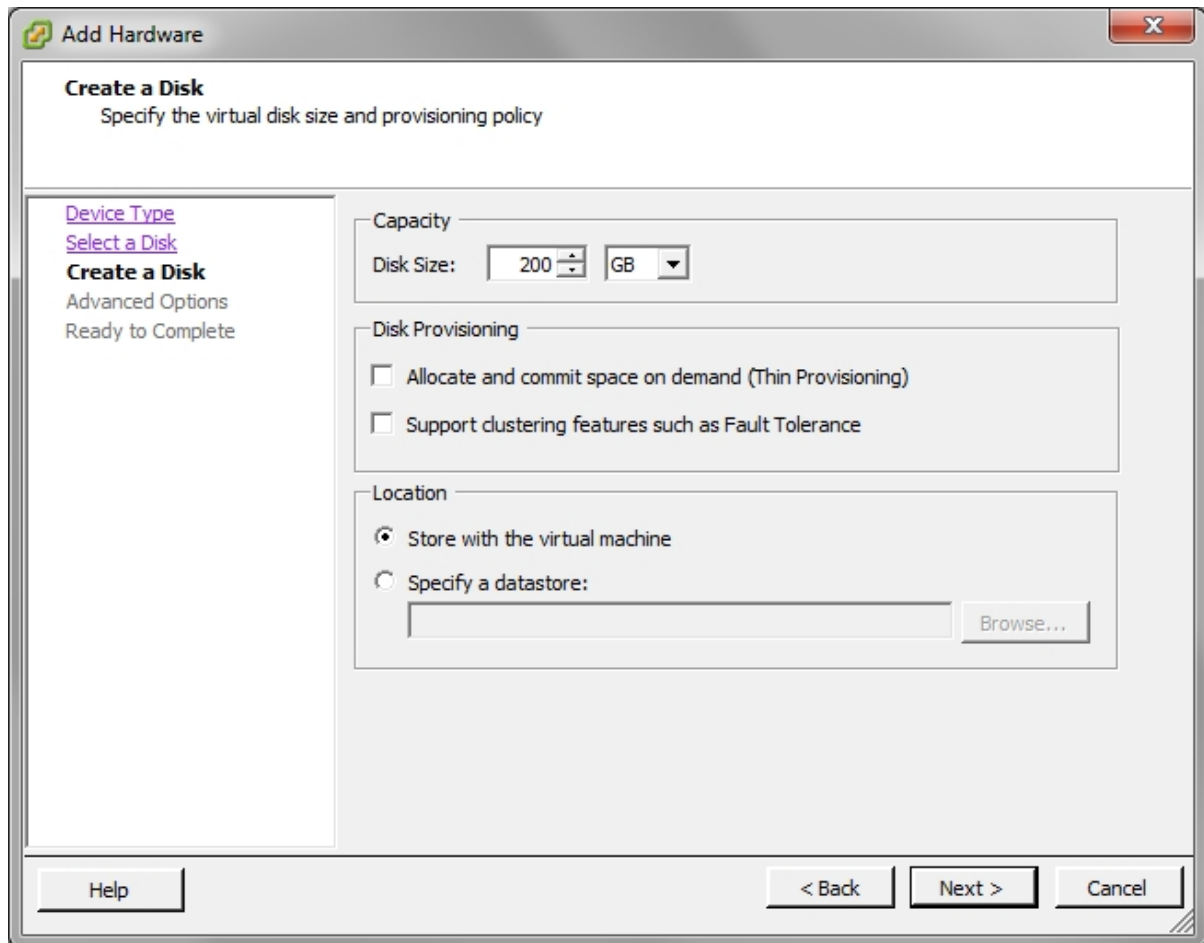
By default, all Exinda Virtual Appliances come with a single 50GB (fixed-size) disk. Usually, you will want more storage for features such as WAN Memory and Edge Cache. This is achieved by adding an additional disk to the Virtual Appliance.

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.

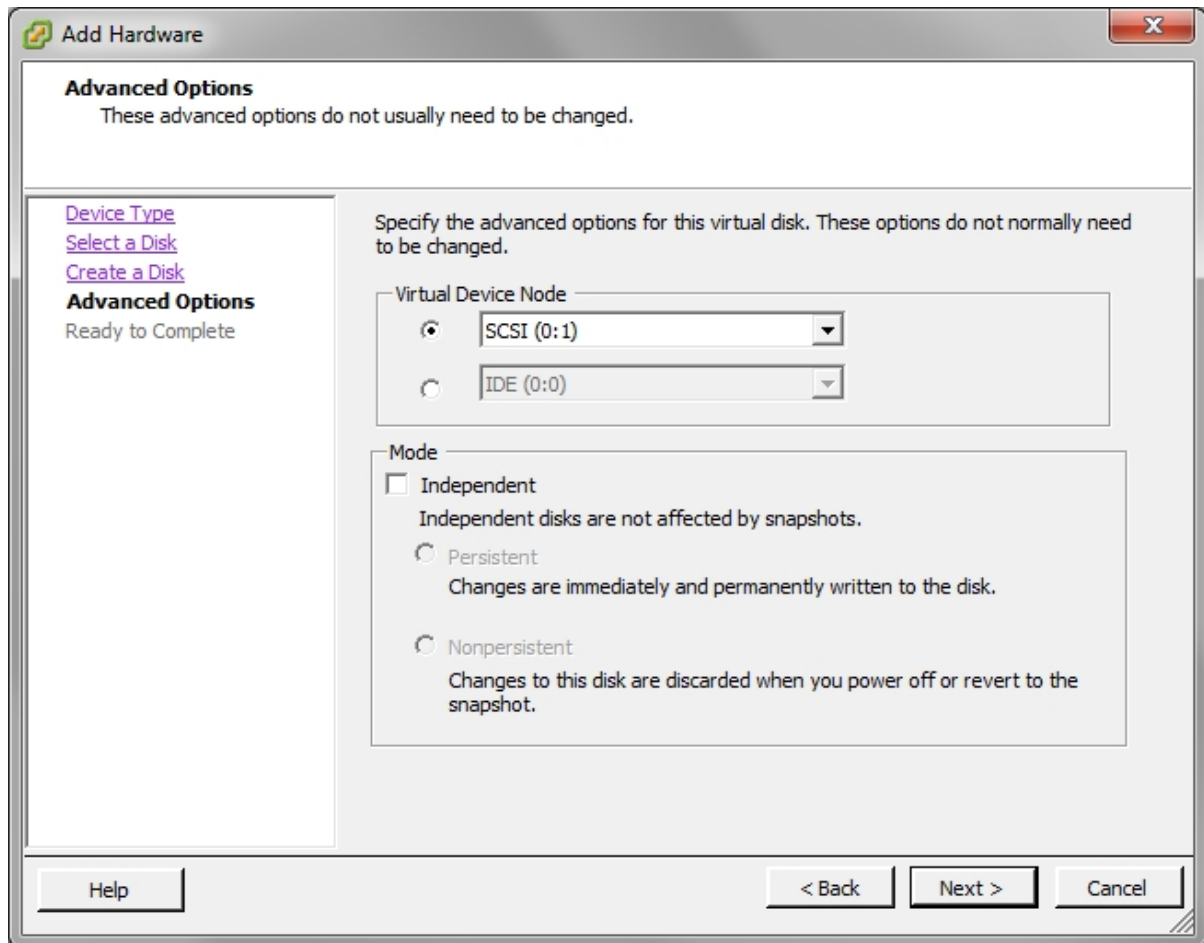
1. Ensure the virtual appliance is powered off.
2. Open the **VMware vSphere Client**.
3. From the Hardware tab in the Exinda Virtual Appliance Properties screen, click **Add**.
4. Select **Hard Disk**, then click **Next**.



5. 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.



6. Click **Next**.
7. Attach the new disk to the next available SCSI node for best performance.



8. Click **Next**.
9. Review the information and click **Finish** to add the disk.
10. 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
```

11. 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
```

12. 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:
  sdal0(internal): in use - 36.22 GB
  sdb: in use - 200.00 GB

Total:          236.21G
Unallocated: 200G
```

Start the VMware Virtual Appliance

When you're ready to start the virtual appliance for the first time, go ahead and 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>`.

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

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

Install the Silicom Bypass Driver

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

- ["Install the Silicom Bypass Driver on ESXi 4.1" on page 105](#)
- ["Install the Silicom Bypass Driver on ESXi 5.0" on page 106](#)

Install the Silicom Bypass Driver on ESXi 4.1

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.
 - Enable SSH through the CLI.
 - a. In the `/etc/ssh/sshd_config` modify the following variable:
`PermitRootLogin yes`
 - b. Restart the `sshd` service.
`# service sshd restart`
 - Enable local or remote TSM from the Direct Console User Interface (DCUI).
 - a. At the DCUI of the ESXi host, press F2 and provide credentials when prompted.
 - b. Scroll to **Troubleshooting Options**, and press **Enter**.
 - c. 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.
 - d. 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.
2. Query the existing VIBs.

If the VIB you are deploying (with `offline-bundle.zip`) exist, you must first remove the existing VIB.
3. Run the following command to determine if any existing VIB matches the VIB you are deploying.

```
# esxupdate --vib-view query | grep bpvm
```

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

4. Optional: Remove the existing VIB.

```
# esxupdate remove -b <VIB_ID> --maintenancemode
```

5. Download the VIB from the following URL:

http://updates.Exinda.com/exos/virtual/vmware/bypass/4.1/vmware-esx-drivers-net-bpvm-2.0.1.9-1OEM.x86_64.vib

6. Copy the VIB to the ESX system with SCP.

```
# scp vmware-esx-drivers-net-bpvm-2.0.1.9-1OEM.x86_64.vib root@<esx-server-
ip>:/tmp
```

7. Deploy the VIB on the ESX system.

```
# esxupdate -b /tmp/vmware-esx-drivers-net-bpvm-2.0.1.9-1OEM.x86_64.vib --nodeps
--maintenancemode --nosigcheck update
```

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

8. Reboot the appliance.

```
# reboot
```

9. After the system has restarted, check that the bpvm driver is loaded in the networks adapter list.

10. If the ESXi host is still in **Maintenance Mode**, exit the same from console or VMware client.

```
# vim-cmd /hostsvc/maintenance_mode_exit
```

11. Install Exinda as a Virtual Appliance.

Install the Silicom Bypass Driver on ESXi 5.0

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.

- Enable SSH through the CLI.

- a. In the `/etc/ssh/sshd_config` modify the following variable:

```
PermitRootLogin yes
```

- b. Restart the sshd service.

```
# service sshd restart
```

- Enable local or remote TSM from the Direct Console User Interface (DCUI).

- a. At the DCUI of the ESXi host, press F2 and provide credentials when prompted.

- b. Scroll to **Troubleshooting Options**, and press **Enter**.

- c. 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.

- d. 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.

2. Query the existing VIBs.

Make sure you are in maintenance mode:

```
# vim-cmd /hostsvc/maintenance_mode_enter
```

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

3. Run the following command to determine if any existing VIB matches the VIB you are deploying.

```
# esxcli software vib list | grep bpvm
```

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

4. Optional: Remove the existing VIB.

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

5. Download the VIB from the following URL:

http://updates.Exinda.com/exos/virtual/vmware/bypass/5.0/net-bpvm-2.0.1.13.1-1OEM.500.0.0.472560.x86_64.vib

6. Copy the VIB to the ESX system with SCP.

```
# scp net-bpvm-2.0.1.13.1-1OEM.500.0.0.472560.x86_64.vib root@<esx-server-ip>:/tmp
```

7. Deploy the VIB on the ESX system.

```
# esxcli software vib install -v /tmp/net-bpvm-2.0.1.13.1-1OEM.500.0.0.472560.x86_64.vib --no-sig-check
```

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

8. Ensure that the bypass driver gets loaded after every reboot by adding the following entry in /etc/rc.local.

```
# vmkload_mod bpvm
```

9. Reboot the appliance. -[p

```
# reboot
```

10. After reboot check that bpvm driver is loaded in the networks adapter list.

11. If the ESXi host is still in **Maintenance Mode**, exit the same from console or Vmware client.

```
# vim-cmd /hostsvc/maintenance_mode_exit
```

Note For Exinda releases 6.3 and above **Bridges with Bypass Capability** is displayed on the **System > Network Setups > NICs** page of the Exinda interface.

12. Install Exinda as a Virtual Appliance.

13. For Exinda releases 6.3 and above **Bridges with Bypass Capability** will be displayed on the **System > Network Setups > NICs** page of the Exinda interface.

Virtual appliance use cases

The following scenarios describe different ways of deploying the Exinda virtual appliance (EX-V).

- ["In-line deployment with externally attached LAN" on page 108](#)
- ["In-line deployment with an isolated virtual LAN and virtual applications" on page 110](#)
- ["Out-of-band \(WCCP\) mode" on page 112](#)
- ["Port mirroring / SPAN port Configuration" on page 114](#)
- ["Virtual WAN simulator in an isolated network" on page 117](#)
- ["VMware high availability" on page 118](#)
- ["VMware fault tolerant cluster" on page 120](#)

In-line deployment with externally attached LAN

In this use case, the EX-V is set up for In-Line mode deployment with an externally attached LAN. There are primarily two scenarios for in-line deployment of the virtual appliance:


1. The LAN-side users and applications connect to the EX-V 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 EX-V

This use case discusses the first scenario.


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

Scenario Details

- The branch users access local services (print, DHCP, Active Directory) that are virtualized.
- Two physical NIC interfaces are bridged together in the virtual EX-V.
- Users connections 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 external NIC interfaces.
- The Exinda provides Visibility, Control, and Acceleration for all traffic in this configuration.



Optionally, install the EX-V in a separate ESXi host in an Inline-mode configuration and connect through an external switch.

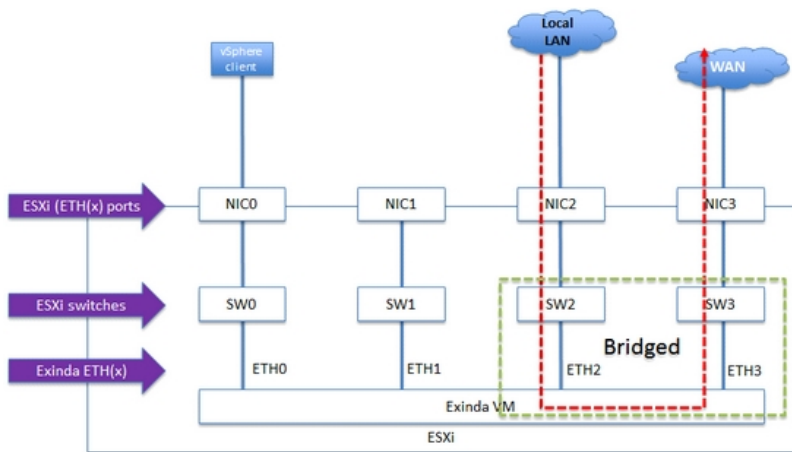


In-line deployment with externally attached LAN (VMware ESXi)

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 Visibility, Control, and Acceleration.

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

Scenario Details



- In the diagram below, a local LAN is sitting behind the host running VMware ESXi hypervisor.
- 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	NIC3
Exinda Ethernet Ports	ETH0	ETH1	ETH2	ETH3
ESXi Virtual Switches	SW0	SW1	SW2	SW3

- 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 & NIC3 are mapped to SW2 and SW3.
- SW2 & SW3 are mapped to EX-V ETH2 & ETH3.
- ETH2 & ETH3 are mapped to NIC 2 and NIC 3, and are configured and bridged together by the Exinda virtual appliance. See In-line Mode Configuration in the Exinda User Guide.
- VMware version = 4.1 & 5.0.
- The data path is from a client on the local LAN goes through the Exinda virtual appliance in In-line Mode and out to the WAN.
- Exinda EX-V software version = 6.3
- The diagram and deployment does not account for external storage. This deployment works with either local or external storage.
- vSphere Client and or vSphere Server are used to manage the system

In-line deployment with an isolated virtual LAN and virtual applications

In this use case the EX-V is set up for In-Line mode deployment with an isolated virtual LAN and virtual applications. There are primarily two scenarios for in-line deployment of the virtual appliance:

1. The LAN side users and applications connect to the EX-V via 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 EX-V.


This use case discusses the second scenario.

Note In this configuration VMware does not support vMotion and Fault Tolerant. HA is supported.

Scenario Details

- Branch users access virtualized local services (print, DHCP, Active Directory).
- One physical NIC interface is configured to the WAN side link.

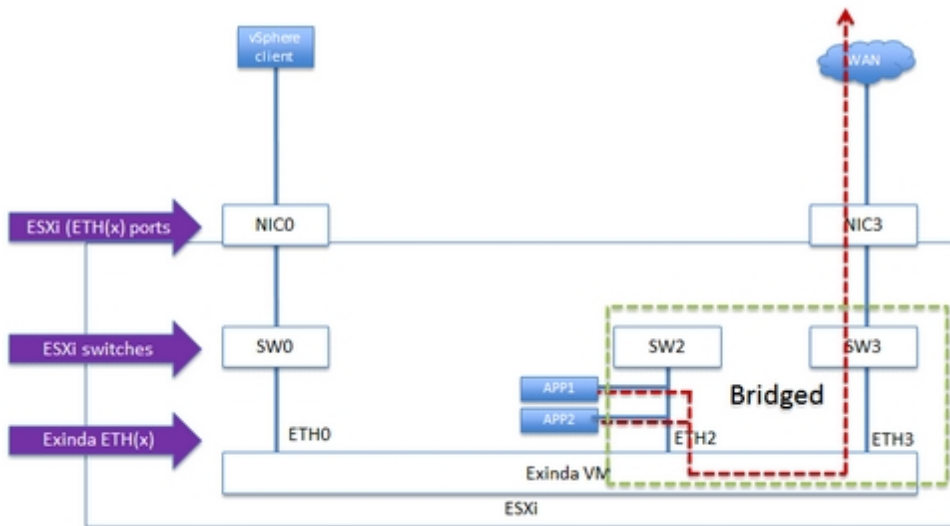
- Two virtual switches, one with virtual applications attached, and a second switch are bridged together in the virtual EX-V. The first switch is on the LAN side; the second one is for the WAN side. This results in isolating the applications behind the EX-V.
- Users connecting from the branch office 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 provides Visibility, Control, and Acceleration for all traffic in this configuration.



In-line deployment with an isolated virtual LAN and virtual applications (VMware ESXi)

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 Visibility, Control, and Acceleration.

Scenario Details



- In the diagram below, a virtual LAN is isolated and sits behind the EX-V all running on the same host running VMware ESXi hypervisor.

- 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		NIC3
Exinda Ethernet Ports	ETH0	ETH2	ETH3
ESXi Virtual Switches	SW0	SW2	SW3

- 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 ESXi to SW2.
- SW2 is configured to map to EX-V 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. See In-line Mode Configuration in the Exinda User Guide.
- The data path for any application connected to the SW2 virtual switch goes through the EX-V in in-line 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 ESXi interface and SW2/ETH2 EX-V interface and out to the WAN. This require mapping a third NIC interface.
- VMware version = 4.1 & 5.0
- Exinda EX-V software version = 6.3

Out-of-band (WCCP) mode

In this use case the EX-V is set up for out of band mode, using WCCP protocol for deployment. There are other similar deployments methods for out of band, such as PBR or policy based routing that are not supported with the 6.3 software release. 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 of a failure of the EX-V all traffic previously redirected to the EX-V will go through un-optimized and un-accelerated.

To configure for WCCP on the Exinda-V, see the Exinda How to Guide: WCCP.

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

Scenario Details

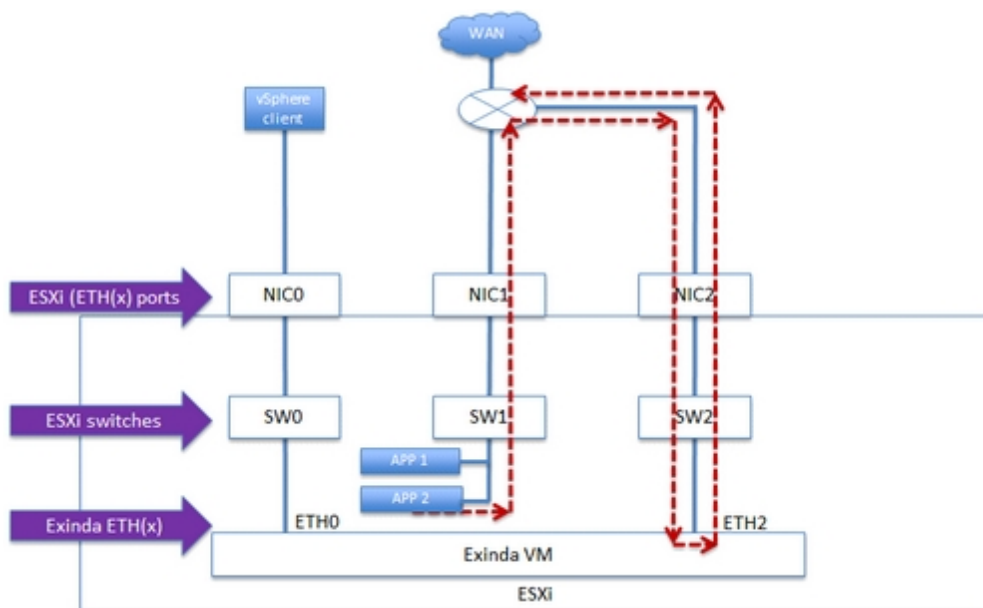
- Branch users access local services (print, DHCP, Active Directory) that are virtualized.
- One physical NIC interface is dedicated for management of the virtual machines and hypervisor.
- One physical NIC interface is configured on the hypervisor for LAN access, and has a routed connection to the WAN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

- One physical NIC interface is configured and mapped to the AUX port on the EX-V.
- 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 EX-V through WCCP on the WAN router.
- The Exinda provides Visibility, Control, and Acceleration for only redirected traffic in this configuration

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

This configuration will be in either 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 Visibility, Control, and Acceleration.

Scenario Details



- In the diagram below, a 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 EX-V 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 EX-V. See the Exinda How to Guide: WCCP.
- VMware version = 4.1 & 5.0
- Exinda EX-V software version = 6.3
- 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 EX-V through NIC2/SW2 on the AUX port of the EX-V.
- (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 EX-V virtual appliance.

Port mirroring / SPAN port Configuration

In this use case the EX-V is set up to monitor and collect traffic for reporting. This topology is used when customers need to monitor only, without installing the Exinda in in-line mode. The Exinda will monitor and report on all applications presented on the SPAN/mirror port. This is regularly used to perform network audits as it provides great flexibility in restricted and complex network environments.

To configure SPAN and Port Mirroring on the EX-V, see the Exinda How to Guide: SPAN and Port Mirroring guide.

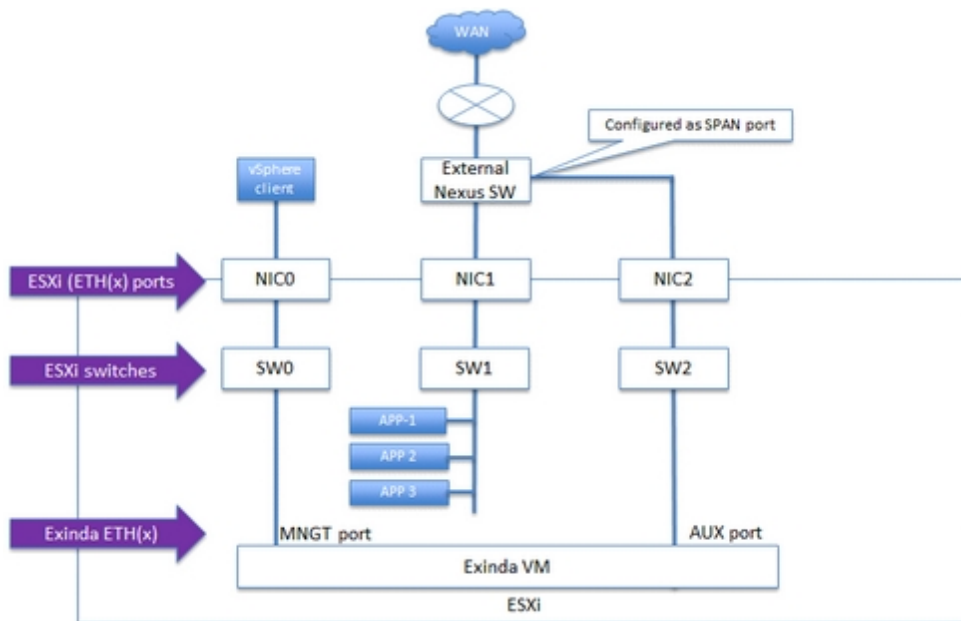
There are two deployment scenarios to consider:

- ["Port mirroring with an external Nexus switch" on page 114](#)
- ["Port mirroring with a virtual Nexus switch" on page 115](#)

Port mirroring with an external Nexus switch

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

Scenario Details

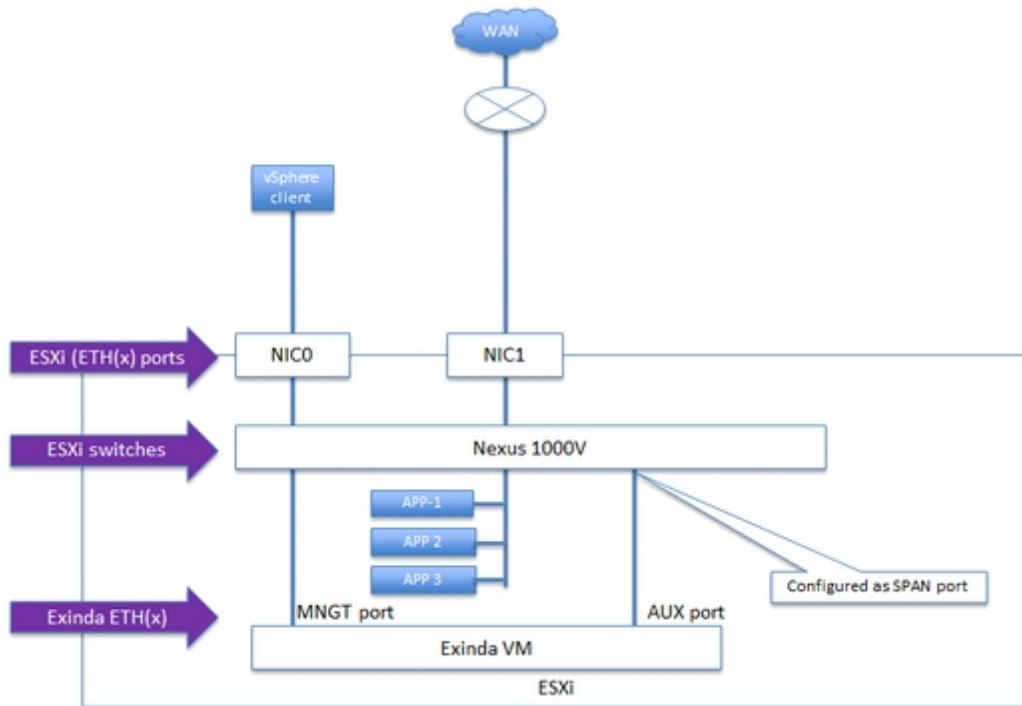


- The customer has selected Exinda for its monitoring and reporting service. The EX-V is configured as virtual machine on an ESXi hypervisor, on a dedicated NIC2 interface, and dedicated virtual switch SW2.
- The EX-V management port is mapped to SW0 and NIC0.
- The ESXi host EX-V has three four Ethernet interfaces (0-23) with the APP1-3 mapped to SW1 and NIC1 respectively, and connected to the external Nexus Switch.
- The EX-V with the AUX (ETH1) port is configured for SPAN Port Mirroring and management of the EX-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

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

Scenario Details



- 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 - 4.1 and 5.0
- EX-V versions – 6.3
- The Nexus 1000V is configured and mapped to NIC1, which has a direct connection to the WAN router.
- The EX-V has two four Ethernet interfaces (ETH0 / ETH10-3) with the AUX (ETH1) configured for Mirroring and ETH0 for management and management of the EX-V.
- The EX-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.

Virtual WAN simulator in an isolated network

The purpose of this deployment is for partners, customers, and Exinda Sales Engineers to have an isolated environment for demonstrations and functional testing of EX-V. The benefit is the portability and lack of impact to a production network. This virtual deployment can be installed on a laptop that has processor support for VT (virtual technologies) enablement and 64 Bit.

Scenario Details

- A single host that is running a hypervisor with the following components:
 - At least one Windows client (Windows XP or 7)
 - Two EX-V virtual appliances
 - One virtual WAN simulator
 - One Windows server (2003 or 2008)
 - Web Server and or FTP services configured on the Windows Server
- This deployment can be run either in VMware or XenServer.

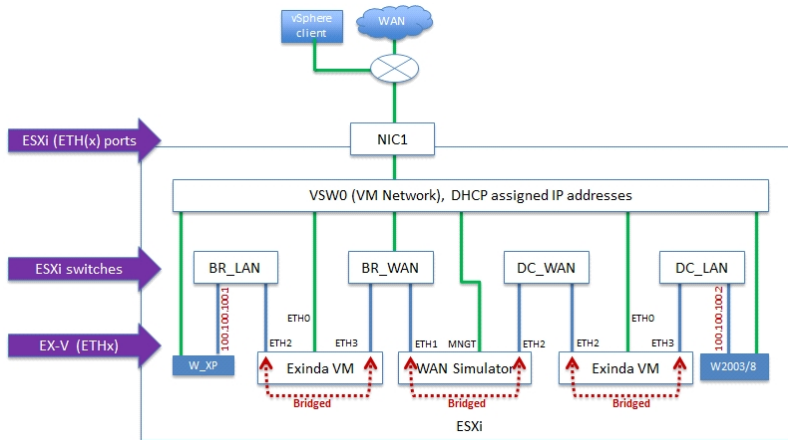
Virtual WAN simulator in an isolated network (VMware ESXi)

- 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.

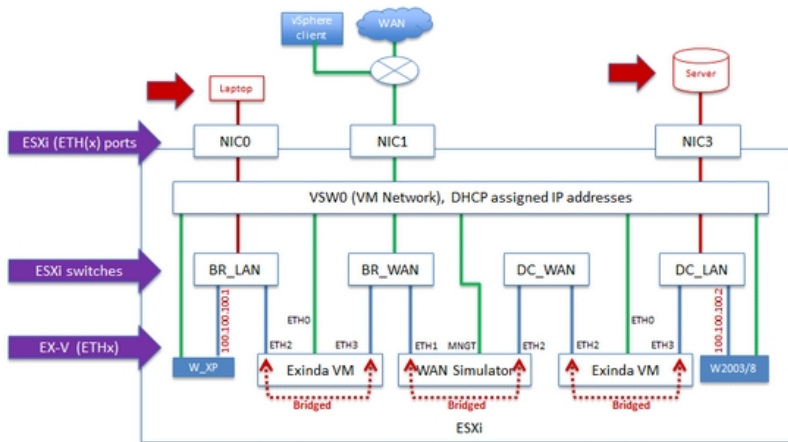
Note Any physical NIC interface can be used, NIC1 is used for illustration purposes.

- VMware software version = 4.1 & 5.0
- Exinda EX-V software version = 6.3
- 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 EX-V is configured for INLINE Mode and a single management interface on ETH0.
- DHCP is assumed on the network for management interfaces on the EX-V 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.
- EX-V 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 you can with 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 EX-V environment with no impact to a product network



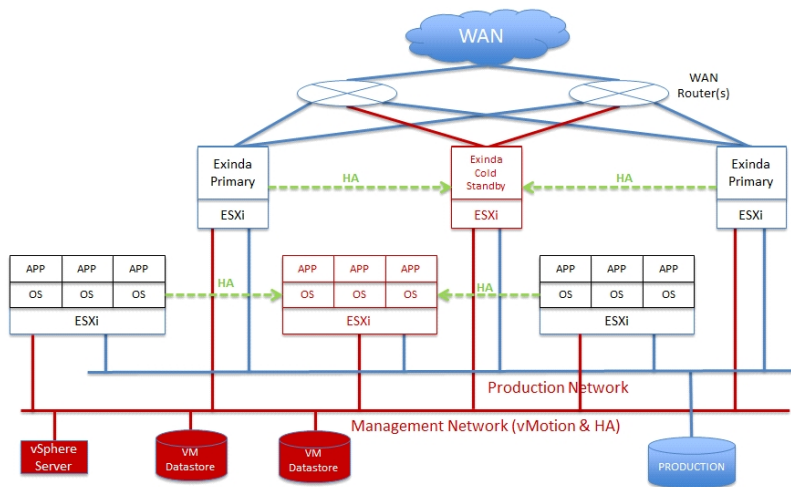
VMware high availability

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

- Exinda Active / Active (software version – 6.3),
- ESX/ESXi HA (software versions 4.0, 4.1 and 5.0),
- 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 EX-V units will include two full licenses and one cold standby license.

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

Scenario Details



- 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 EX-V appliances and a third running as a cold standby
- Optionally you can move the EX-V 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 EX-V.
- Having a separate host for the EX-V 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 EX-V to reboot.

VMware fault tolerant cluster

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

- Exinda Active / Passive (cold standby) - (software version – 6.3),
- ESX/ESXi Fault Tolerant Cluster (software versions 4.0, 4.1, and 5.0),
- vMotion support,
- vMotion = Yes for INLINE,
- VMware best practice recommends that at least 4 hosts are used for this configuration, and
- Licensing for the EX-V units will include one full license and one cold standby license.
 - Each EX-V 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 EX-V can co-reside with the Application workloads, but an external switch and an additional NIC is required for vMotion support.

Scenario Details



- There are four ESX/ESXi Hosts:
 - one running virtual workloads and a second as the Fault Tolerant system
 - one running EX-V active and a second running as a cold standby and Fault Tolerant.

- Optionally, move the EX-V 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 EX-V.
- Having a separate host for the EX-V 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.
- This configuration has 100% uptime for one virtual workload per host for VMware version 4.1, and up to four virtual workloads in 5.0.

Run the Virtual Appliance on Citrix XenServer

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

Note The Exinda Virtual Appliance requires Citrix XenServer 5.5 or later.

Installing the Virtual Appliance on XenServer

1. Locate the latest release of the Exinda XenServer Virtual Appliance from the Firmware Downloads section of the Exinda [website](#).
Download the ZIP file to your local PC and unzip the Virtual Appliance XVA file.
2. Open your Citrix XenCenter client and select **File > Import...**
Then, select the unzipped Virtual Appliance XVA file.
Click **Next**.
3. Select the target XenServer to deploy the Virtual Appliance.
Then click Next.
4. 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. See the [Additional Storage](#) section for more information.
Then click Next.
5. Choose the NIC mapping. By default, the Virtual Appliance comes with 2 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 deploy the Virtual Appliance inline, you will need to add additional NICs after deployment. See the [Additional NICs](#) section for more information.

Then click Next.

6. Review the information and uncheck the 'Start VM(s) after import' box if you want to add extra NICs or storage.

Click Finish to deploy the Virtual Appliance.

7. Select the Exinda virtual machine you are importing, and switch to the Log tab to see the progress and the completion notification. 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.
8. Right-click 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.
9. On the XenCenter Console tab of the Exinda virtual machine, enter 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.
10. Press Enter to read the EULA agreement. Press Ctrl-C to get to the EULA agreement question.
11. Press Y to accept the EULA agreement and press Enter.
12. You will 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.

- a. Select No to disable IPv6.
 - b. Select Yes to configure ETH0 for management access. This will disable the BR0 bridge.
 - c. Select Yes to use DHCP on ETH0.
 - d. Select null to default to the Exinda hostname.
 - e. Select null for SMTP server address.
 - f. Select null for email address for reports and alerts.
 - g. Select null to use the default password which is "Exinda".
 - h. Select Yes to change the interface speed.
 - i. Select AUTO to configure the interface speed on ETH0 (assumes a gigabit NIC).
- You have successfully completed the wizard setup.
13. Determine the IP address of your Exinda virtual appliance on the XenServer Network tab of the Exinda virtual machine and note IP address assigned by default to NIC 0.
 14. Browse to the Dashboard tab and find the Host-ID that the XenServer host created for this virtual machine.

Once the appliance is deployed, review the following sections on [Custom Settings](#) and adding extra [NICs](#) and [Storage](#).

Custom Settings

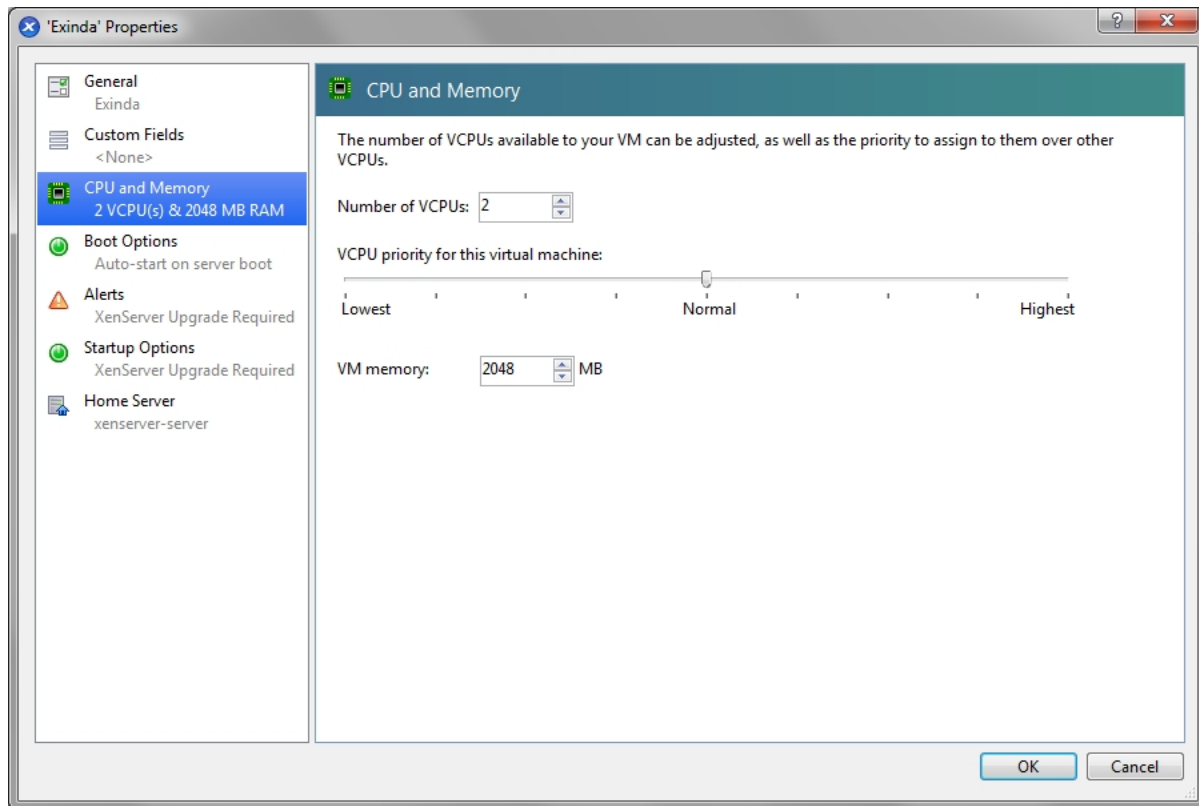
The most common customizations to Virtual Appliances are to increase the number of CPUs and the amount of RAM. By default, all Virtual Appliances come configured with 2 virtual CPUs and 2GB RAM. These should be considered minimum values, however, you can increase these to suit your requirements.

Caution

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

Note You should adjust CPU and RAM settings while the Virtual Appliance is powered off.

From the Exinda Virtual Appliance Properties screen, you can increase the number of virtual CPUs and the amount of RAM as shown below.




Additional NICs

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

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

- Convert the 2 default 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 **System > Network > IP Address** page on the Web UI, advanced mode. From this page, you can click the 'br0' checkbox to bridge the 2 default NICs together. In order to manage the Virtual Appliance, an IP Address must be specified for this bridge.



Interface Settings	
Autoconf:	IPv4: <input checked="" type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC
Dynamic Addresses:	192.168.0.221/24 fe80::20c:29ff:feaa:b541/64
br0 <input checked="" type="checkbox"/>	Static Addresses: <input type="text"/> / <input type="text"/>
Comment:	<input type="text"/>
Gateway Settings	
IPv4:	<input type="text"/>
IPv6:	<input type="text"/>
<input type="button" value="Apply Changes"/>	

Take care when using this option as this will cause the 2 default NICs to be bridged.

- Add 2 extra NICs to be used as LAN and WAN ports. The 2 additional NICs will be bridged and allow the Virtual Appliance to be placed inline.

Caution

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

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.

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

Add Virtual Interface

Select your network and MAC address for this virtual interface.
You can also optionally define a QoS limit.

Network: Network 2

MAC address:

Auto-generate a MAC address


Use this MAC address: aa:bb:cc:dd:ee:ff

QoS settings:

Enable a QoS limit of: Kbytes/s

Add **Cancel**

3. 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	
br0 <input type="checkbox"/>	Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP Autoconf: IPv4: <input checked="" type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC Dynamic Addresses: 192.168.0.225/24 fe80::10e0:9ff:fe0d:3021/64 Static Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>
eth0	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: 192.168.0.225/24 fe80::10e0:9ff:fe0d:3021/64 Static Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>
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::4430:b2ff:fe2a:c4a5/64 Static Addresses: <input type="text"/> / <input type="text"/> Comment: <input type="text"/>
br2 <input checked="" type="checkbox"/>	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC Dynamic Addresses: fe80::2c5d:7fff:fe84:9221/64 Static Addresses: <input type="text"/> / <input type="text"/> 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.

Caution

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
```

- 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

- 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 the Storage How to Guide.

Booting

When you're ready to boot the Virtual Appliance for the first time, go ahead and 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>>.

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

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

Licensing

Licensing is a little different for Virtual Appliances compared to Hardware Appliances. All Virtual Appliances are shipped unlicensed. On first-boot, they automatically generate a unique Host ID. Exinda must be notified of this Host ID before a license can be issued. All Virtual Appliances must have access to Exinda's licensing server, and must be able to access <https://license.exinda.com>. Virtual Appliances that do not have access will become unlicensed after 96 hours.

To obtain a trial license or to purchase a full license for the Exinda Virtual appliance, contact sales@exinda.com.

Generating a virtual appliance trial license

In this step you will go to https://license.exinda.com/virtual_trial/ URL 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_trial/.
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 you have installed the virtual appliance on.
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 **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.

Purchasing 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 EX-V appliance and the PO number. As with the trial license process, you must have installed the EX-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 ESXi 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 licenseYou will need this information for each EX-V virtual appliance.
4. To purchase a license, navigate to https://license.exinda.com/virtual_purchase/.
5. Complete the Virtual Appliance Purchase form as required.

Hypervisor limitations

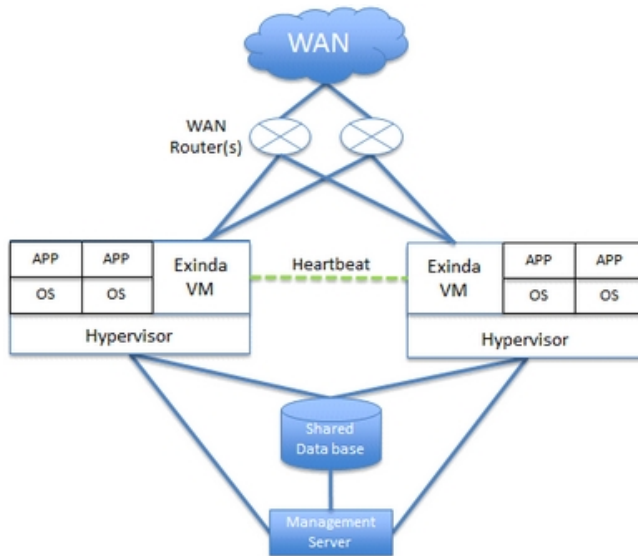
Consider these additional planning items when installing EX-V on other Hypervisors in the market including:

- XenServer 5.0, 5.5, and 6.0
- Microsoft Hyper-V on Windows 2008 R2
- KVM

What are the known limitations?

- 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 4.1 and 5.0.
- 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.
- KVM
 - Images for KVM are not published on the Exinda support website, you will need to contact support for the KVM images
- 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.



Topology troubleshooting

Q. My network traffic is blocked after deploying the Exinda appliance in-line.

A. 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.

Q. My network traffic is blocked after deploying the Exinda appliance in-line, after I have booted it up.

A. Ensure the speed/duplex settings are correct on both the Exinda appliance and any neighboring equipment.

Q. I am experiencing significant packet loss after deploying the Exinda appliance in-line, after I have booted it up.

A. See above regarding speed/duplex configuration. Also check Ethernet cables for defects.

Maintenance

The Maintenance section of the Exinda appliance System Setup allows you to perform various system maintenance tasks. These include:



- ["Manage System Configuration" on page 341](#): Allows you to save, activate, switch, revert and delete system configuration files.
- ["Import System Configuration" on page 342](#): Allows you to import previously saved or backed-up system configuration files.
- ["Cluster and High Availability" on page 343](#): View the status of Exinda clustering.
- ["Install an update to the Exinda appliance software" on page 344](#): Upgrade the ExOS software on the Exinda appliance.
- ["Factory Defaults" on page 345](#): Restore the Exinda appliance to factory default settings.
- ["Reboot the Exinda appliance" on page 346](#): Reboot or Shutdown the Exinda appliance.

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 **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"/> initial.bak		
<input type="checkbox"/> initial	<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.

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 **System > Maintenance > Import Config** on the Web UI, advanced mode.

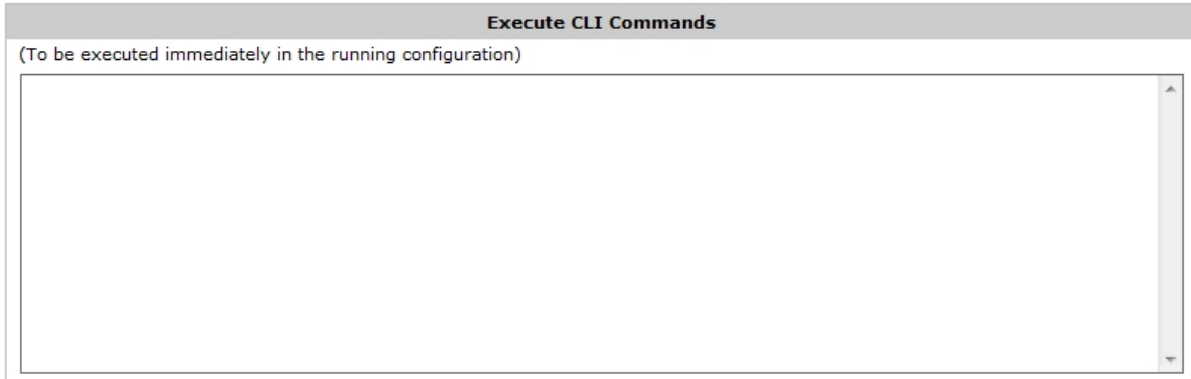
The form below can be used to upload system configurations that have been saved locally on the PC.

Upload Configuration	
<input checked="" type="radio"/> Upload local binary file:	<input type="text"/> <input type="button" value="Browse..."/> (To be saved as separate file with its original name)
<input type="radio"/> Upload local text file: (CLI commands)	<input type="text"/> <input type="button" value="Browse..."/> (To be executed immediately in the running configuration)

Upload Configuration

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 System Maintenance Manage Config page. Once this file is uploaded, it will appear in the list of available configuration files on the System Maintenance Manage Config 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 System Maintenance Manage Config 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

Install an update to the Exinda appliance software

Exinda software is called ExOS and is updated regularly with new product features as well as system and performance improvements. The Firmware Update screen allows you to update the software on your Exinda appliance.

Exinda appliances have two partitions for installing ExOS updates. The current, running ExOS version will be installed on one partition, which means you can install a newer ExOS update on the other, unused partition.

Note Valid Software Subscription (SS) is required to install new ExOS updates. You can view your SS expiry date at the top of this page.

1. Click **System > Maintenance** and switch to the **Firmware Update** tab.
2. Locate the new image to install.
 - If you know the URL for the image file, select **Install from URL** and type the URL.
This URL is usually published in the Release Notes as part of the ExOS release. If you click **Check for Latest Update**, and a newer ExOS software update is available, this field is populated automatically.
 - If the image file has previously been downloaded 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 file.
3. (Optional) Schedule the download and installation of the ExOS update for a later date or time.
 - a. Select the **Schedule Installation** checkbox and specify the **Date** and **Time**.
 - b. By default, the download of the image will happen straight away and only the installation will be scheduled. To schedule the download of the ExOS image to happen at the scheduled time, select **Schedule Image Download**.

- c. By default, the Exinda appliance will not reboot following a scheduled installation. To restart the appliance after the scheduled installation, select **Reboot After Installation**.
4. Before installing or scheduling a new ExOS update, you must accept the End User License Agreement (EULA).
5. Click **Install**.

The image is installed on the appliance. This process may take a few minutes to complete.

Note If the network connection fails while retrieving the latest ExOS 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.

6. To finalize the ExOS install, you must reboot the appliance. See "[Reboot the Exinda appliance](#)" on page 346.

Factory Defaults

The Factory Defaults screen allows you to restore the Exinda appliance's configuration back to factory default settings. This includes removing any system logs, WAN Memory cache and monitoring statistics.

Note To restore Factory Defaults, navigate to **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

After performing a Factory Defaults, the Exinda appliance will automatically reboot.

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 Reboot or Shutdown without saving the changes first.

1. Click **System > Maintenance** and switch to the **Reboot / Shutdown** tab.
2. (Optional) Schedule the Exinda appliance to reboot at a specific date or time.
 - a. Select the **Schedule Reboot** checkbox.
 - b. Type the date and time that the appliance should reboot.
3. 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.

4. Click **Reboot**.

Rebooting the Exinda appliance may take a few minutes to restart.

Chapter 3: Access the Appliance

There are two ways to access the appliance:

- ["Access the Command Line Interface" on page 348](#)
- ["Web User Interface \(Web UI\)" on page 138](#)

By default, the Exinda appliance will attempt to automatically obtain an IP address on its management interface using DHCP. Unless a valid DHCP address is obtained, the appliance will be assigned a default IP address of 172.14.1.57. If DHCP is unavailable or fails, you will need to change the IP address on your PC to 172.14.1.X in order to connect to the Exinda appliance on its default IP.

If DHCP is available, you will need to know the IP address that has been assigned to the appliance (unless physically accessing the appliance with a monitor/keyboard or serial console). A convenient way to determine the appliance's IP is to visit the www.findmyexinda.com website. This site loads a small Java client that interrogates your local LAN, looking for Exinda appliances.

By default, there are 2 predefined user accounts:

Username	Default Password	Privileges
admin	exinda	read-write
monitor	exinda	read-only

When logging into the appliance for the first time using either method, you will be required to accept the End User License Agreement (EULA) before continuing.


Access 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

A screenshot of a terminal window with a white background and a black border. The text "login:" is displayed in a monospaced font, followed by a small red square cursor. The terminal has a dark, jagged shadow effect on its right and bottom edges.

```
login: 
```

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

Web User Interface (Web UI)

Access to the Exinda appliance's Web UI is available via a compatible Web browser. Currently, the following browsers are supported:

- Windows Internet Explorer 7 or 8
- Mozilla Firefox 3.x

By default, only https access to the appliance is enabled. You'll need to enter the following address into your browser's address bar in order to connect to the appliance:

```
https://<IP address of appliance>
```

After accepting the default security certificate, you will be presented with the login screen.

The first time you access the Web UI, you will enter 'basic' mode. This mode allows you to view the appliance dashboards as well as complete the initial configuration wizard. You can also switch to 'advanced' mode, which will present you with the full user interface. The next time you login, you will enter the mode chosen last. You can always switch modes after you have logged in.

Note To enable regular http access, navigate to the [System | Setup | Access](#) page on the Web UI, advanced mode.

Switch between Exinda Web UI display modes

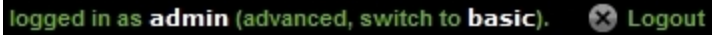
The Exinda Web UI has two display modes: Basic, which displays only the Dashboard and the configuration

wizard; and Advanced, which gives access to the full user interface. The first time you access the Exinda Web UI, the interface is in Basic mode.

1. Launch the Exinda Web UI. [Show steps...](#)
 - a. In a browser, enter `https://Exinda_IP_address`.
 - b. Enter the appliance **User Name** and **Password**.

The Exinda Web UI is displayed.

2. In the Status Bar at the top of the Exinda Web UI, select the appropriate mode.



The next time you login the last selected mode is displayed.

Note The procedures in this manual assume the Exinda Web UI is in **Advanced** mode.

Check the Features Included in your License

Some features, such as SSL Acceleration, Virtualization, and Edge Cache, require a separate license. Check your license to see what features are supported.

1. Click **System > Setup > License**.
2. 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.

Licensed	Host ID	Model	SS Expiry
<input checked="" type="checkbox"/>	f129ade8c50f	Exinda x800 (1000/2500Mbps - HP)	Sep 30, 2012
		Max Bandwidth: 1000 Mbps	
		Optimizer: <input checked="" type="checkbox"/>	
		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: <input checked="" type="checkbox"/>	
		Virtualization: <input checked="" type="checkbox"/>	
		Max Edge Cache Connections: 10000	

3. Please contact your local Exinda representative if you wish to enable a feature.

Chapter 4: Initial Configuration

When you login to the CLI or Web UI for the first time, you will be presented with an option to run a configuration wizard in order to assist with initial setup.

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 enter 'Y', otherwise enter 'N'.

2. Enable IPv6 autoconfig (SLAAC) on eth1 interface?

If you enable IPv6, you have the option of enabling IPv6 SLAAC autoconfiguration. Enter '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 enter static IP connectivity settings, such as IP address and netmask, default gateway and DNS servers.

5. Enable br10 (Y/N)?

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, enter 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. For more information about the Optimizer policy wizard, see the [Optimizer | Optimizer Wizard](#) page.

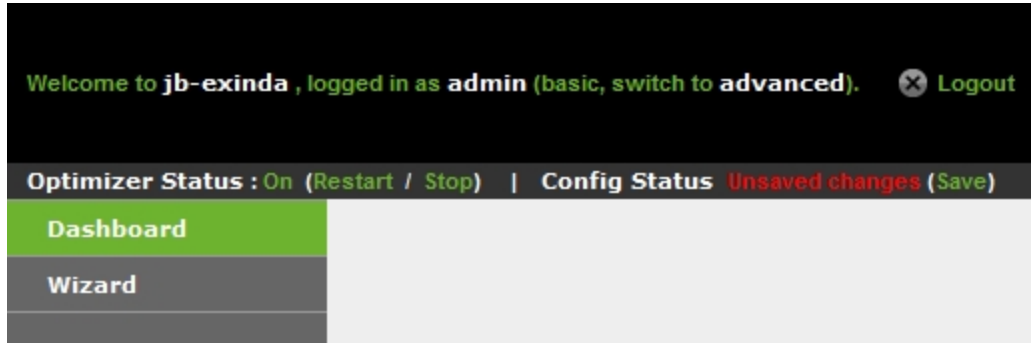
15. Check for new firmware (Y/N)?

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`

Web UI Basic Wizard

When you login to the Web UI for the first time, you enter the 'basic' mode. In this mode, you have access to the appliance dashboards as well as the initial configuration wizard.



To access the wizard, click on the 'Wizard' link on the left-hand-side menu. The wizard assists you to get the appliance up and running in 6 steps:

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

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	<input type="text" value="exinda-3dcaed"/>
Primary DNS	172.16.1.254
Secondary DNS	

* Required field

Back
Next

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	<input type="text" value="lab-services.wat.exinda.com"/>	HTTP(S) Proxy Authentication	<input type="text" value="None"/>
HTTP(S) Proxy Port	<input type="text" value="3128"/>	HTTP(S) Proxy Username	<input type="text"/>
		HTTP(S) Proxy Password	<input type="password" value="....."/>
		Do not verify SSL certificates	<input checked="" type="checkbox"/>

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:

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

Software Subscription Expiry: **Dec 31, 2009 (45d)**

License Expiry: No license expiry date

Host ID: 0010f305cd54

Monitor:

Optimize:

Accelerate:

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.

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

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](#)

8. Optimization

The final screen allows 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? Yes No
Selecting YES will overwrite any existing policies you have configured.

Step 3: Do you want to enable Optimization? Yes No
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.

Step 4: Do you want to enable Control? Yes No
Selecting YES will apply traffic shaping.

Step 5: Select the topology type WAN or WAN + Internet? WAN WAN + Internet



The diagram illustrates a network topology. On the left, there is a green WAN appliance. A green line connects it to a central green router. Another green line connects the router to a cloud labeled 'Internet' on the right.

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)

[Back](#) [Finish](#)

For more information about the Optimizer policy wizard, see the [Optimizer > Optimizer Wizard](#) page.

Note Settings on each step are automatically applied when clicking the 'Back' or 'Next' buttons.

Chapter 5: Dashboards





There are two types of dashboards available to view on the Exinda appliance:

- "System Dashboard" on page 146: Shows system information.
- "Benefits Dashboard" on page 147: Show performance information.

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.

Hostname: appliance-beta-exinda-02		Alarm	Status	Last Triggered	Count
Hardware Series:	4061	CPU Utilization	OK		
Licensed Model:	Exinda 4861 (1000/2500Mbps - HP)	System Disk Full	OK		
SS Expiry Date:	Dec 31, 2016	Memory Paging	OK		
Host ID:	b8ac6f874f7c	Bridge Link	OK		
Serial Number:	70MN1P1	Bridge Direction	OK		
Timezone:	Canada/Eastern	Link Negotiation	OK		
System Uptime:	2d 21h 45m 9.088s	NIC Problems	OK		
Scheduled Jobs:	No scheduled jobs.	NIC Collisions	OK		
Memory Usage:	51.54% of 3881MB	NIC Dropped Packets	OK		
CPU Usage:	6%	SMB Signed Connections	OK		
Database Status:	Running	Redundant Power	Not Available		
		Redundant Storage	Not Available		
		Accelerated Connections	DISABLED		
		Asymmetric Route Detection	OK		
		MAPI Encrypted Connections	OK		

WAN Reduction per Peer						
Hostname	IP Address(es)	Version	Status	LAN Data	WAN Data	Reduction Ratio (%)
				16MB	12MB	 21.48
jl-home	192.168.0.208		ONLINE	3MB	1MB	 59.44
beavers42	10.0.26.228	6.1.3.1706	ONLINE	0MB	0MB	 38.03
con-home3	10.40.88.212	6.1.8.1788	ONLINE	0MB	0MB	 46.94
forta-home	58.26.32.137	6.1.0.17088	ONLINE	12MB	11MB	 9.28

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 you exactly what the Exinda appliance is doing. Each widget can be hidden so you can customize the display to only include the widget that are relevant to you. To add a hidden widget click on 'Add More'. The widget settings are retained between logins. The dashboard can also be converted to PDF by clicking on the PDF icon in the top, right-hand corner of the interface.

Visibility

Visibility is an essential ingredient to maintaining clean network pipes. These graphs show the applications that are utilizing the network. This information is critical to the IT Manager to better manage the network and to make informed decisions.



Example

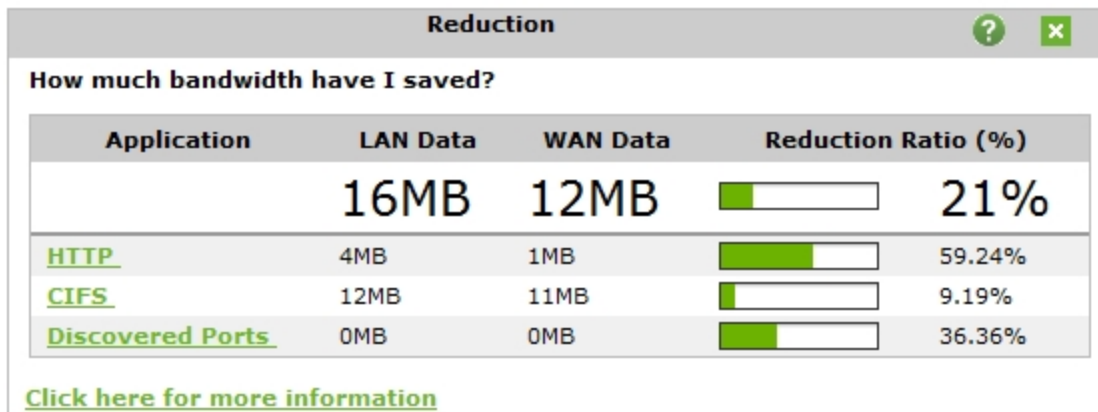
This graph can tell you if the network is being miss-used. Users downloading music and videos are choking the network; miss-configured user profiles are being downloaded every day from the wrong location causing congestion and delays; data backups are running overtime and into normal

business hours.

Reduction

Reduction refers to the amount of redundant data that has been removed from the network and has therefore increased free capacity. It is a ratio that compares After Exinda (AE) to Before Exinda (BE). 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.

$$\text{Reduction Ratio} = (\text{Data Transfer Size BE} - \text{Data Transfer Size AE}) / \text{Data Transfer Size BE}.$$



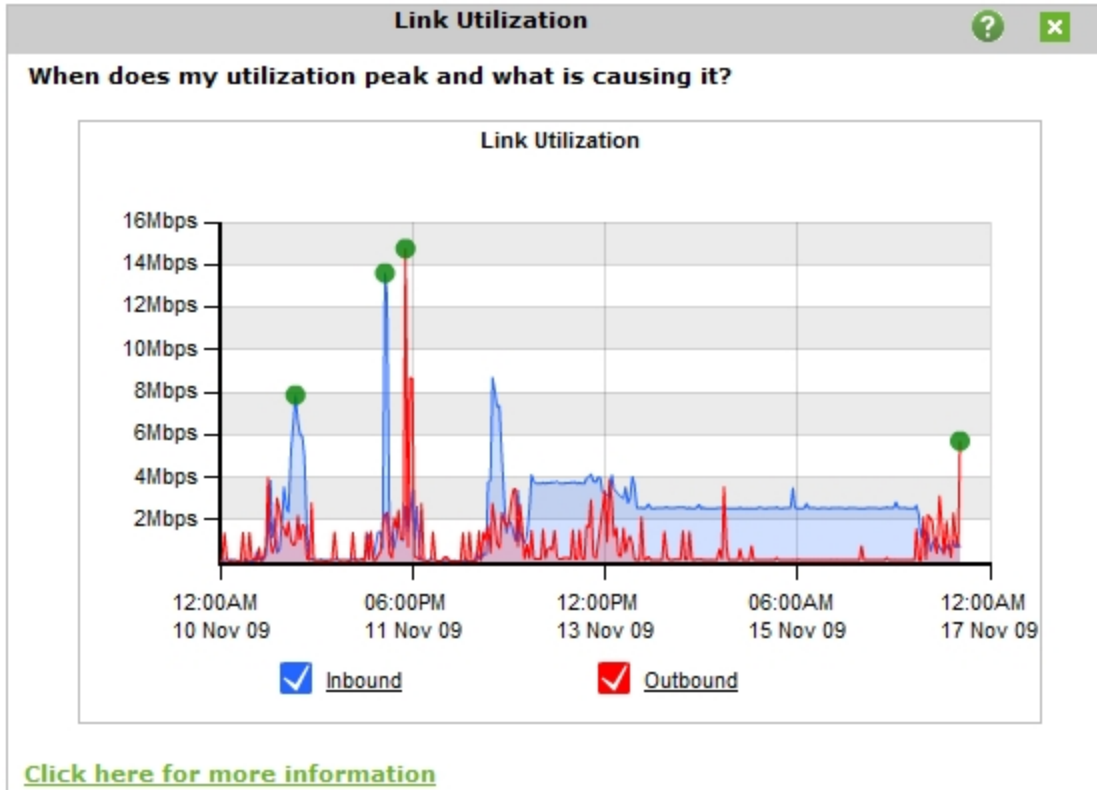
Example

A ratio of 40% means a transfer that used to put 100MB of load onto the WAN, now puts 60MB of load on the WAN. I.e. 40% less.

Link Utilization

Without the right visibility and drill down capability, it is very difficult to provide effective capacity planning for your network. It is important to know:

- you are receiving the bandwidth you are paying for,
- when the network usage peaks, and
- the causes of bottlenecks on your network.



Utilization graphs with conversation drill down allow you to see in one glance the symptom and causes of network bottlenecks. Without this level of visibility IT professionals may mistakenly treat the symptoms (for example buy more bandwidth) instead of addressing the cause which is a short term solution at best.

Recreational

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.

Recreational			
How much recreational usage is there?			
Application	Hosts	Time	Data
	29	6h 57m 50s	1798MB
Games	1	1m 10s	0MB
Instant Messaging	25	5h 10m 10s	89MB
P2P	6	29m 30s	263MB
Social Networking	18	46m 20s	221MB
Streaming	17	30m 40s	1223MB

[Click here for more information](#)

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.

$$\text{Prioritization Ratio} = \frac{\text{Number of Packets Re-ordered}}{\text{Number of Total Packets}}$$

Prioritization	
How often were my critical applications prioritized?	
Prioritization Rate: 7.5%	
Click here for more information	

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.

$$\text{Before} = \frac{\text{LAN Data}}{\text{WAN Throughput}}$$

Time Savings  **How much time is my optimization saving?**

Application	Before	After	Saving	%
Quicktime	18m	8m	9m	52%
YouTube	24m	18m	6m	27%
Total	1h01m	26m	35m	56%

Example

A Savings value of 50% means that the time taken to transfer an application's data was reduced by half.

Chapter 6: System Settings, Configuration, and Diagnostics

More advanced system settings, configuration and diagnostics are available by using the Web UI in Advanced Mode.

Network Settings

The Network Settings section of the Exinda appliance System Setup allows you to configure basic and advanced network settings. The various configuration pages include:

- [NIC Settings](#): Configure network interface cards settings.
- [IP Address](#): Configure an IP mode, IP address and gateway.
- [Routes](#): Configure a static route.
- [DNS](#): Configure a hostname, DNS and domain name.
- [HTTP Proxy](#): Configure HTTP proxy settings.
- [Email](#): Configure SMTP sever settings and email address.
- [SNMP](#): Configure SNMP settings.
- [Active Directory](#): Enable and fine-tune Active Directory.
- [IPMI](#): Enable and configure IPMI on selected hardware.

NIC Settings

Interface Settings

Use the form below to set the speed/duplex and MTU of the System NICs. 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 there are collisions and/or errors, then it is an indication that the Exinda is not auto negotiating with neighboring equipment. As a result you might notice packet loss and network delays. To resolve this check if the router or switch is hard-coded to a speed or duplex setting. If hard-coded then set the Exinda device to a desirable speed/duplex.

Note Collisions, errors and dropped packets on the Exinda NICs will set the System health status to "Warning" and the offending interface(s) will be highlighted. For further troubleshooting click on the system warning or view the NIC Diagnostics.

[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)

Interface	This is the interface number. Each interface corresponds to a physical port.
Media	Specifies the interface media. This could be Twisted pair or Fibre.
HW Address	This is the MAC address of the interface.
Speed	This is the speed at which the Exinda will negotiate with neighboring equipment.
Duplex	This is the duplex at which the Exinda will negotiate with neighboring equipment.
MTU	This is the maximum transmission unit size in bytes.
Link Status	The status of the interface 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.

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

Fail to Wire (bypass)

The Fail to Wire (bypass) settings control the behaviour of the Exinda appliance's 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. The image below shows independently controllable bypass bridges.

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

The image below show globally controllable bypass bridges.

Bridge	Status	Running Mode	Enable Failover	On Failover
br0,br1	Active	Active ▾	<input checked="" type="checkbox"/>	Bypass ▾

Apply Changes

Bridge	The bridge that the bypass settings apply to. Where available, bridges can be controlled independently, otherwise they will be controlled globally.
Status	The current status of the bridge - see below for definitions.
Running Mode	Specify the current status of the bridge. This allows you to change the current status of the bridge on the fly (e.g. manually put the bridge in and out of bypass) - see below for definitions.
Enable Failover	Enable failover in the event of failure, power outage or reboot. If not enabled, no action will be taken on failover.
On Failover	If failover is enabled, specify what action to take when failing over - see below for definitions.

The table below lists the various statuses and failover modes that are available. Depending on your hardware, the following options may or may not be available.

Active	The bridge is active (not in bypass) and traffic is been intercepted by the Exinda appliance.
Bypass	The bridge is in bypass and traffic is NOT been intercepted by the Exinda appliance.
Nolink	The bridge interfaces are both forced to link state down (as if the cables are not plugged into the interfaces).

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. Use the form below to enable or disable link state mirroring.


Link State Mirroring
Link State Mirroring <input type="checkbox"/> Enable

Apply Changes

IP Address

The Exinda appliance allows you to configure bridges and network interfaces as required. The form displays an image showing the available physical interfaces, physical interface to I/O slot and physical interface to

bridge assignments. Bridges can be enabled, roles assigned to an interface (Cluster, Mirror or WCCP) and IP settings applied.



Interface Settings	
eth1	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::a6ba:dbff:fe1e:9eec/64 Static Addresses: 172.24.1.80 / 24 Comment:
eth2	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::a6ba:dbff:fe1e:9eee/64 Static Addresses: / Comment:
eth3	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: / Comment:
eth4	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: / Comment:
br10 <input checked="" type="checkbox"/>	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC Dynamic Addresses: fe80::2e0:edff:fe17:3870/64 Static Addresses: / Comment:
br20 <input checked="" type="checkbox"/>	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC Dynamic Addresses: fe80::2e0:edff:fe16:c00e/64 Static Addresses: / Comment:
br30 <input checked="" type="checkbox"/>	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC Dynamic Addresses: fe80::2e0:edff:fe0e:9eea/64 Static Addresses: / Comment:
br40 <input checked="" type="checkbox"/>	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC Dynamic Addresses: fe80::2e0:edff:fe1a:e8a4/64 Static Addresses: / Comment:
Gateway Settings	
IPv4:	172.24.1.1
IPv6:	

A bridge consists of a LAN and WAN interface. To enable a bridge, select the brXX checkbox above the interface pair.

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.

To configure an interfaces address and netmask automatically, select either the DHCP checkbox for IPv4 networks or SLAAC for IPv6 networks. When SLAAC is selected for IPv6 networks, the following options are available:

Privacy Address	Enable SLAAC privacy extensions. Selecting this option will periodically change the automatically assigned IPv6 address.
Gateway	Assign an IPv6 gateway dynamically.

To configure a static address, enter an IPv4 or IPv6 address and netmask.

Enter the address of your networks default IPv4 and IPv6 gateways.

You can optionally add a comment describing how the interface is to be used in the Comment field.

The DHCP option is enabled by default on the Exinda appliance. If a DHCP server is available, an IP address will be automatically assigned. From a web browser go to www.findmyexinda.com. This will download a Java applet and automatically find the Exinda appliance. Click on the Exinda appliance that has been found to access it. If a DHCP address is not picked up, the Exinda will default to the IP address of 172.14.1.57.

Model	Factory default DHCP enabled interface
2000, 2061, 4000	br1

Model	Factory default DHCP enabled interface
5000, 6000, 6010	eth0
2060, 4010, 4060, 4061, 6060, 6062, 8060, 8062, 10060, 10062	eth1

The VLAN configuration allows an 802.1Q VLAN ID to be set on an interface. The VLAN ID can be between 1 and 4094.

VLAN Settings

Interface: ▼

ID:

Add VLAN

The Cluster Master address is the external address used to access an appliance in HA environments.

Cluster Master Settings

Interface: ▼


Master Address: /

Apply Changes

Further information on Clustering/HA, Mirroring and WCCP is available in the associated How To guides.

Cluster and High Availability

Clustering allows multiple Exinda appliances to operate as if they were a single appliance. This allows for seamless deployment into High Availability and Load Balanced environments. A typical deployment topology is illustrated below.



In this example, there are two physical links. 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 and acceleration policies, as if they were a single appliance.

Refer to the following topics for example topologies:

- ["Redundancy through multiple Exinda appliances" on page 73](#)
- ["Load balancing and fail-over with multiple Exinda appliances" on page 75](#)
- ["High availability mode" on page 77](#)

Routes

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

Destination	The IPv4 or IPv6 address and netmask of the destination
Gateway (Next Hop)	The IPv4 or IPv6 address of the gateway (next hop).

Configure DNS and Domain Names

Set a Host Name for your Exinda appliance and configure the DNS servers and domain names. The hostname should be unique to this appliance on the network.

In this area of the Exinda Web UI you can:

- "Set the host name and DNS of the Exinda appliance" on page 160
- "Add a domain name" on page 160
- "Remove a domain name" on page 160

Set the host name and DNS of the Exinda appliance

Configure the appliance hostname and DNS servers. The hostname should be unique to this appliance on the network.

Note A valid DNS server is required for Edge Cache, system alerts, scheduled reports, firmware updates, license updates and Anonymous Proxy updates.

1. Click **System > Network**, and switch to the **DNS** tab.
2. Type a **Host Name** for the Exinda appliance.
3. Type the IP address of up to three DNS servers.
4. Click **Apply Changes**.

The IP address is added to the list of Static and Dynamic Name Servers.

Add a domain name

Add any domain names that the Exinda appliance should try to resolve automatically.

1. Click **System > Network**, and switch to the **DNS** tab.
2. In the Add New Domain Name area, type the new domain name.
3. 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.

Remove a domain name

Maintain the list of domain names by removing any obsolete entries.

Note Only manually added domain names can be removed.

1. Click **System > Network**, and switch to the **DNS** tab.
2. Select the domain to remove from the Static and Dynamic Domain Names list.
3. Click **Remove Selected**.

The domain is removed.

HTTP Proxy

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.

1. Click **System > Network** and switch to the **HTTP** tab.
2. Specify the hostname or IP address of the HTTP proxy.
IPv4 or IPv6 addresses can be specified.
3. Select the type of authentication for the HTTP proxy.
4. Type the **Username** and **Password** for the HTTP proxy.
5. To verify SSL certificates, clear the **Do not verify SSL certificates** checkbox.
6. Click **Apply Changes**.
7. To write the changes to the configuration file, in the status bar click **Save**.

Config Status *Unsaved changes* (Save)

Configure the appliance to send email notifications

An SMTP server is required for receiving scheduled reports, system alerts and auto-support notifications. Configure the SMTP server, and manage the users who receive system notifications and reports.

In this area of the Exinda Web UI you can:

- ["Add an SMTP server for sending email notifications" on page 281](#)
- ["Add a user to receive email notifications" on page 161](#)
- ["Stop sending notifications to a user" on page 162](#)

Add an SMTP server for sending email notifications

An SMTP server is required for receiving scheduled reports, system alerts and auto-support notifications.

1. Click **System > Network**, and switch to the **Email** tab.
2. In the SMTP Server area, type the SMTP server name.
IPv4 and IPv6 addresses can be used.
3. Type the SMTP server port.
The default port number is 25.
4. Type the email address that system alerts and report notifications will appear to have been sent from.
5. To require authentication against the SMTP server before emails can be sent, select the **SMTP Authentication** checkbox.
After selecting SMTP Authentication, you must provide the username and password for the SMTP server, and select the authentication method.
6. Click **Apply Changes**.
7. To ensure the users can successfully receive notification emails, click **Send Test Email to Add**.

Add a user to receive email notifications

Add each user that is monitoring the Exinda appliances, and should be receiving any system alerts and report emails.

Note The emails being received by a user cannot be modified. To change which emails a user receives, delete the user, and then add the email address with the appropriate emails selected.

1. Click **System > Network**, and switch to the **Email** tab.
2. In the Add New Notify Recipients area, type the email address of the user who should receive notifications.
3. Select the emails 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 this recipient.
4. Click **Add New Recipient**.
The new recipients are added to the Notify Recipients list.
5. To ensure the users can successfully receive notification emails, click **Send Test Email to Add**.

Stop sending notifications to a user

Maintain the list of users receiving notifications to ensure that only the necessary users are receiving emails.

Note The emails being received by a user cannot be modified. To change which emails a user receives, delete the user, and then add the email address with the appropriate emails selected.

1. Click **System > Network**, and switch to the **Email** tab.
2. In the Notify Recipients 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.

SNMP

The Exinda appliance allows data export to SNMP systems. Configure the SNMP settings or download the Exinda SNMP MIB.

In this area of the Exinda Web UI you can:

- ["Modify the SNMP configuration" on page 163](#)
- ["Remove an SNMP community" on page 163](#)
- ["Download the SNMP MIB file" on page 163](#)
- ["Modify the SNMP administrator user settings" on page 164](#)
- ["Add an SNMP trap sink server" on page 164](#)

- "Remove an SNMP trap sink server" on page 164
- "Enable or Disable an SNMP trap sink server" on page 164

Note To disable or enable SNMP traps for system alerts, see "Notify administrators of system issues" on page 279.

Modify the SNMP configuration

Update the SNMP settings to identify where SNMP traps are sent, and whether there can be multiple SNMP communities.

1. Click **System > Network** and switch to the **SNMP** tab.
2. Select whether **SNMP Traps** are sent from the appliance.
3. Select whether **Multiple Communities** can be configured.
When Multiple Communities is disabled, the Community list area is not displayed.
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.
When the Read-only community is changed to have a value that doesn't match an existing community, a new SNMP community is added to the list.
7. Click **Apply Changes**.
8. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Remove an SNMP community


Remove any unwanted SNMP communities.

1. Click **System > Network** and switch to the **SNMP** tab.
2. In the SNMP Communities area, select the community from the list and click **Remove Selected**.
3. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Download the SNMP MIB file

Download the SNMP MIB file to gain additional monitoring information.

1. Click **System > Network** and switch to the **SNMP** tab.
2. In the SNMP Configuration area, click **Download SNMP MIB** .
The EXINDA-MIB.txt is downloaded to the location specified within the browser.

Modify the SNMP administrator user settings

Change the authentication type, the privacy type, or the password for the SNMP admin user.

1. Click **System > Network** and switch to the **SNMP** tab.
2. In the SNMP v3 Admin User area, modify the settings as needed.
3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Add an SNMP trap sink server

Define the server where SNMP traps should be sent.

1. Click **System > Network** and switch to the **SNMP** tab.
2. In the Add New Trap Sink area, specify the hostname or IP address of the SNMP trap sink server. IPv4 or IPv6 addresses can be specified.
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**.
6. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Enable or Disable an SNMP trap sink server

To temporarily stop sending SNMP traps to a sink server, disable the server. When the server can be used again, re-enable the server.

1. Click **System > Network** and switch to the **SNMP** tab.
2. In the Trap Sinks area, select the server from the list and click **Enable Trap Sink** or **Disable Trap Sink**.
3. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Remove an SNMP trap sink server

Remove any unnecessary SNMP trap sink servers.

1. Click **System > Network** and switch to the **SNMP** tab.
2. In the Trap Sinks area, select the server from the list and click **Remove Server**.
3. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Integrate the Exinda Appliance with Active Directory

Using the Exinda Active Directory Connector, customers can:

- 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.

The Exinda Active Directory Connector needs to be installed onto a server in the network that has access to the Active Directory server. Each install of the Exinda Active Directory Connector can talk to up to 20 Exinda appliances.

Complete the following tasks to connect the Exinda Active Directory Connector to the Active Directory server, and view user names in monitoring reports.

1. ["Install the Exinda Active Directory Connector" on page 165](#)
 - a. ["Add Exinda appliances to the Active Directory Connector" on page 167](#)
 - b. (Optional) ["Configure the connection to the Active Directory server" on page 169](#)
 - c. ["Select the information sent between the Exinda appliance and the Active Directory server" on page 168](#)
 - d. ["Change the Active Directory Connector port number" on page 167](#)
2. ["Identify users using applications on a Citrix XenApp server" on page 172](#)
3. ["Verify communication between the Active Directory server and the Exinda Appliance" on page 169](#)
4. ["Report on Network Activity by User" on page 75](#)
5. ["Controlling Traffic based on Users" on page 307](#)

Note If you encounter any issues, see ["Troubleshoot issues with Active Directory configuration" on page 188](#)

Install the Exinda Active Directory Connector

The Exinda Active Directory Connector needs to be installed onto a Windows server that can connect to the Active Directory server. Each Exinda Active Directory Connector can talk to up to 20 Exinda appliances. When you first install the Exinda Active Directory Connector, it may take 24 hours or longer to get all user to IP address mappings as users progressively login.

Notes

- The Exinda Active Directory Connector is supported on the following platforms:
 - Windows Server 2003 SP2

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.

- Windows Server 2008 SP2
- Windows Server 2008 R2
- Windows Server 2012
- The Exinda Active Directory Connector requires .NET Framework 4.0.
- The Logon Auditing must be enabled on the Active Directory server to install the Exinda Active Directory Connector.
- The WMI service must be started on the Active Directory server and on the server where the Exinda Active Directory Connector is installed.
- The Active Directory server and the server where the Exinda Active Directory 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>.

1. Download the installer the Exinda appliance.
 - a. Click **System > Network**, and switch to the **Active Directory** tab.
 - b. Download the **Microsoft Installer Executable**.
2. Save the Exinda Active Directory Connector install to a location that can be accessed by all Windows servers in the network.
3. On the server where the Exinda Active Directory Connector should be installed, locate and double-click installation file.
4. On the Welcome dialog, click **Next**.
5. Read the End-User License Agreement. Select the **I accept...** checkbox and click **Next**.
6. Specify the directory where the Exinda Active Directory Connector should be installed.
7. 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.

Caution 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.

8. Click **Next**.
9. (Optional) Type the Exinda appliance IP address or hostname, port number, and administrator password.

Adding an Exinda appliance can be completed after the Exinda Active Directory Connector is installed.
10. 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.

11. Click **Next**.

12. On the Check for Required Services dialog, click **Next**.

If any warnings are displayed on the page, resolve the issues as specified in the dialog.

13. Click **Install**.

14. 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.

Add Exinda appliances to the Active Directory Connector

Identify the Exinda appliance using this Active Directory 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.

1. In the **Start** menu click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.

2. Switch to the **Exinda Appliances** tab.

3. To add the Exinda appliances that communicate with this Active Directory Connector, type the IP Address or Hostname into an empty row, and type the Admin password for the appliance.

4. Modify the port number as needed. The default port number of the Active Directory Client is 8015.

When changing the port number, all Exinda appliances using the Exinda Active Directory Connector must use the same port number. See "[Change the Active Directory Connector port number](#)" on page 167.

5. In the **Sync interval** field, identify how frequently the Exinda Active Directory Connector contacts the Exinda appliances to synchronize Active Directory user and group information. The default is 5 minutes.

6. Click **OK**.

Change the Active Directory Connector port number

Identify the port on which the Exinda Active Directory Connector is communicating to the connected Exinda appliances. Changing the port number is optional, and the default port 8015 automatically communicates with the Exinda Active Directory Connector and Exinda appliances.

Note The port number must be the same on all installed Exinda Active Directory Connector instances, and all Exinda appliances using the Exinda Active Directory Connector.

Ensure firewall on the server running the Exinda Active Directory Connector is configured to allow inbound and outbound traffic on configured port.

1. Change the port number on the Exinda Active Directory Connector.
 - a. In the **Start** menu click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
 - b. Switch to the **Exinda Appliances** tab.
 - c. Type the port number in the field.
The default port number is 8015.
 - d. Click **OK**.
2. Change the port number on the Exinda appliances.

Note The port number must be changed on each Exinda appliance using the Exinda Active Directory Connector.

- a. Launch Exinda Web UI.
 - a. In a browser, enter `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.
- b. Click **System > Network**, and switch to the **Active Directory** tab.
- c. In the **Active Directory** area, type the port number of the Exinda Active Directory Connector.
The default port number is 8015.
- d. Click **Apply Changes**.

When the Exinda appliance successfully communicates with the Exinda Active Directory Connector, the following information is displayed in the table:

- **Agent Name**—The Exinda Active Directory Connector name.
- **IP Address**—The IP address of the server running the Exinda Active Directory Connector.
- **Version**—The Exinda Active Directory Connector version.
- **Windows Version**—The version of Windows on the Active Directory server.
- **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 Active Directory Connector, it may take up to 24 hours (or longer) to get all user to IP address mappings as users progressively login.

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 Active Directory 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, Send users/groups to Exinda appliances on startup should only be selected on one of the domain controllers.

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 appliances**.
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 Active Directory Connector service starts.
4. Click **OK**.

Configure the connection to the Active Directory server

The Exinda Active Directory 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, specify the location and authentication credentials of the Active Directory server.

1. In the **Start** menu click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
2. Switch to the **AD Server** tab.
3. In the Active Directory Server area, select whether the Active Directory server is **this server** or **another server**.
4. Type the IP address or hostname of the Active Directory server.
5. Type the username and password of the Administrator account on the Active Directory server.
6. Click **OK**.

Verify communication between the Active Directory server and the Exinda Appliance

Ensure the communication between the Active Directory server and the Exinda appliance is active.

Note User accounts that have been disabled on the Active Directory server are not included in the data sent to the Exinda appliances.

1. Click **System > Network**, and switch to the **Active Directory** tab.
2. 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.
3. If the service is not visible on the list, run the Event Viewer program on your Active Directory server, and examine Windows logs.
 - a. In the Start menu select **Control Panel > Administrative Tools**.
 - b. Double-click **Services**, and verify the status of the **Exinda AD** service. If the service is stopped, restart the service.
 - c. In the **Windows Logs > Application** area, the “Service started successfully” message should be displayed from Exinda Networks Active Directory Connector.

If the communication between the Active Director and the Exinda appliance is 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. Click **System > Network**, and switch to the **Active Directory** tab.
2. 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 Active Directory Connector

Temporarily stop or disable the Active Directory integration to help with troubleshooting and to avoid errors when modifying the Exinda Active Directory Connector settings.

1. Click **System > Network**, and switch to the **Active Directory** tab.
2. Modify the state of the Active Directory service.
 - To temporarily stop the Exinda Active Directory Connector, click **Stop**.
 - If you are experiencing issues with the Exinda Active Directory Connector, **Restart** the service.
 - If you no longer need the Exinda Active Directory 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. SMB signing was introduced with the 6.4.1 release. Configure the Exinda Active Directory Connector to prevent the IP address to username mapping being sent to the Exinda appliance.

1. In the **Start** menu click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
2. Switch to the **Excluded Users** tab.
3. Click in the Ignored Users area, type the full username of each user to ignore.

The username is 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.

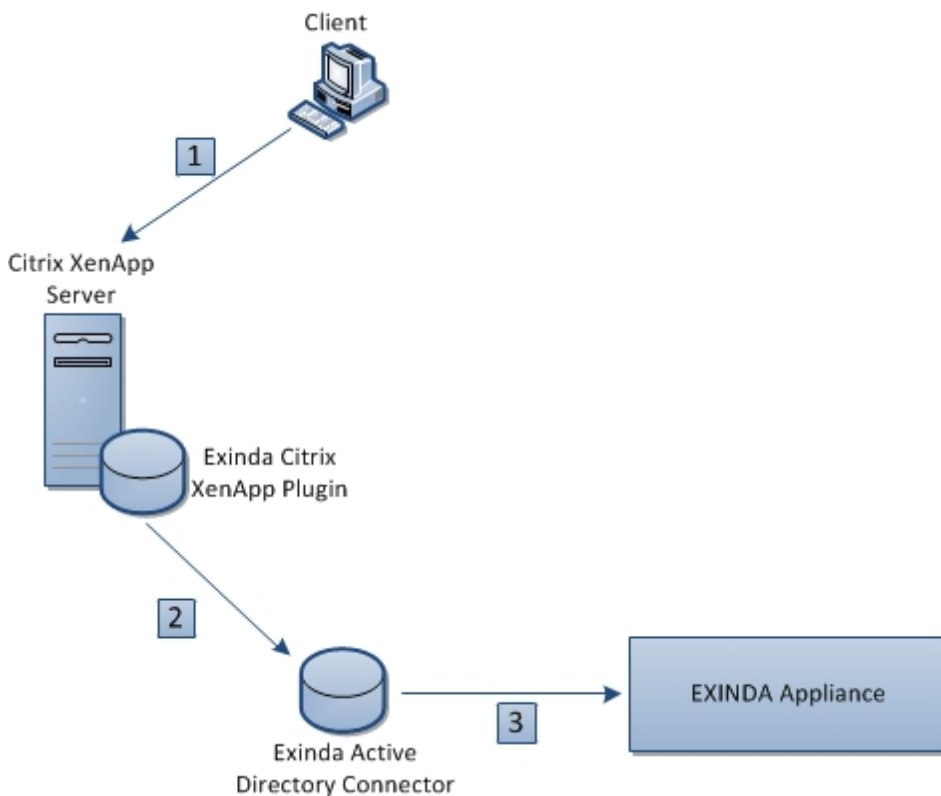
Note Regardless of the case of usernames in Active Directory, the Exinda appliance displays the usernames with the first name capitalized and the surname in lower case; for example Domain/Test.user. Do not use the value in the Exinda appliance when adding a username to the Excluded list.

4. Click **Apply**.
5. ["Request updated user and group information from the Active Directory server"](#) on page 170.
6. Restart the Active Directory service. See ["Change the state of the Exinda Active Directory Connector"](#) on page 170.

Identify users using applications on a Citrix XenApp server

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 whom are accessing the applications on the XenApp server. With the 6.4.1 release, Exinda has created a plug-in for the Citrix XenApp server that sends the IP address of the client computer and the user name used to authenticate with the XenApp server to the Exinda Active Directory Connector when a user accesses an application through 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 Active Directory 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.

1. ["Install the Exinda Citrix XenApp Plugin" on page 173](#)
2. ["Add the Exinda Active Directory Connector to the Exinda Citrix XenApp Plugin" on page 173](#)
3. ["Capture the Exinda Citrix XenApp Plugin activity in a log file" on page 174](#)

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 Active Directory 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 **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. 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 Active Directory Connector to the Exinda Citrix XenApp Plugin

To ensure user activity on the Citrix XenApp server is reported on the Exinda appliance, add the connection details for the Exinda Active Directory Connector to the Exinda Citrix XenApp Plugin.

1. Open the Exinda Citrix XenApp Plugin.
2. On the **Synchronization** tab 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 Active Directory Connector is installed.

Note The port number used to communicate between the Exinda Active Directory Connector and the Exinda Citrix XenApp Plugin cannot be the same as the port number used to communicate between the Exinda Active Directory Connector and the Exinda appliances.

4. In the **Sync Interval** field, identify how frequently the Exinda Active Directory Connector sends XenApp server user information to the Exinda Active Directory 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 Active Directory Connector is communicating to the connected Exinda Citrix XenApp Plugins. The default port number is 8016.

1. Change the port number on the Exinda Citrix XenApp Plugin.
 - a. In the **Start** menu click **All Programs > Exinda Networks > Exinda Citrix XenApp Plugin Configuration**.
 - b. Switch to the **Synchronization** tab.
 - c. Double-click the port number for the appropriate Exinda Active Directory Connector and type the new port number in the field.
 - d. Click **OK**.
2. Change the port number on the Exinda Active Directory Connector.
 - a. In the **Start** menu click **All Programs > Exinda Networks > Exinda AD Configuration Utility**.
 - b. Switch to the **XenApp** tab.
 - c. Type the port number in the field.
 - d. 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 Active Directory Connector is infrequent, trigger the Exinda Citrix XenApp Plugin to send the data to the Exinda Active Directory Connector immediately.

1. In 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 Active Directory Connector.

Report on Network Activity by User

Use the information in reports to determine how the policies on your Exinda can improve the quality of service and the experience of your network users.

The following reports identify user activity on the network:


- ["View All Network Activity for a Specific User" on page 77](#)
- ["Top Users Generating Traffic" on page 75](#)
- ["Top Internal and External Users on the Network" on page 79](#)
- ["Real-time Traffic by Hosts" on page 80](#)
- ["View real-time inbound and outbound conversations" on page 81](#)

Top Users Generating Traffic

The top users generating inbound and outbound traffic are displayed in a graph. The table view shows the amount of packets and Megabytes transferred as well as the average and maximum throughput per user.

1. Click **Monitor > Users**.
2. Select whether the charts display **Internal** or **External** users.
3. ["Set the Time Period Reflected in the Report" on page 76](#).
After the date range is select, the graphs and charts are immediately updated.
4. Hover over the pie slices to view the amount of data transferred in megabytes and percentage.

Top 8 Internal Users Receiving Inbound Traffic



Each table shows the top Users together with the amount of data transferred, number of packets, number of flows and throughput statistics.

- To display additional details for each user such as round trip time (RTT), transaction delay and efficiency statistics, click **Show Details**.
- To display the applications that are generating the inbound or outbound traffic for user, click the user-name. You have the option to view Applications, URLs, Hosts and Conversations in the drill-down.

Top 50 Internal Users Receiving Inbound Traffic					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
[+] Show Details					
MELB\Pforto	256599	328.319	264.20	4034.39	224
themis	224420	142.290	28.87	1656.61	1194
MELB\Kdeikos	76578	101.077	79.85	887.10	433
MELB\Csiakos	90826	89.588	21.56	3961.84	564
rbyrne	55507	75.806	167.84	946.30	84
MELB\Asavant	110339	71.434	7.47	612.92	1124
scott	60506	54.297	12.48	695.97	689
MELB\Matt	33342	27.351	9.67	1854.87	455
salah-pc	34944	26.560	11.65	855.17	498
MELB\Avis	41831	14.614	1.58	172.73	2243

Set the Time Period Reflected in the Report

The data displayed in the reports can be focused on specific periods of time. Date ranges are available on all

reports except the Real Time reports.

1. Select a report from the Monitor list.
2. 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

3. 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.

Range: 12:00AM 25/Oct/2010 - 12:00AM 26/Oct/2010

After the date range is select, the graphs and charts are immediately updated.

Data Granularity

The Exinda appliance stores data for the following amount of time:

- 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

View All Network Activity for a Specific User

The Applications reporting displays the inbound and outbound traffic that has passed through the Exinda broken down by application for a specific user. The applications graphs can be used to determine which applications are currently using the link the most and at what speeds.

The table view shows the amount of packets and Megabytes transferred as well as the average and maximum throughput. The Top Application Objects are shown in a pie chart. You can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie.

1. Click **Monitor > Users**.
2. In the Select Users to View list, select **Internal**.
3. "[Set the Time Period Reflected in the Report](#)" on page 76.

After the date range is select, the graphs and charts are immediately updated.

4. In the details table, click the user name.

The **Traffic Analysis - Users - Applications** report is displayed and lists all the applications that have generated inbound and outbound traffic for that user.

5. On the Traffic Analysis - Users report there are multiple reports available. Select whether to display the user's **Applications**, **Conversations**, **URLs**, or **Hosts** report.


Applications—All applications sending inbound and outbound traffic on the user's computer.

Conversations—All inbound and outbound data sent and received from the user's computer.

URLs—All inbound and outbound URL requests.

Hosts—All internal and external hosts communicating with the user's computer.

6. Hover over the pie slices to view the amount of data transferred in megabytes and percentage.



The table shows the top Application Objects together with the number of packets, number of flows, data transferred and throughput statistics.

Top 50 Inbound Applications					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
[+] Show Details					
HTTP	6817611	5334.422	18.45	8801.78	12960
Shoutcast	1873255	2523.348	258.10	307.57	9
SMTP	1927337	2297.050	297.60	2584.53	291
YouTube	1088274	1512.809	367.74	1377.49	300
HTTPS	2619277	1472.302	7.89	1407.26	2197
IMAP-SSL	1769795	1459.845	26.97	534.73	186
Flash	914997	1199.138	76.21	828.16	1354
FTP	728912	1034.463	5432.26	7970.57	12
MPEG	542279	716.296	203.89	930.83	14
OpenVPN	372946	471.839	96.49	148.52	19

Note If a previously defined Application Object has been deleted, it will appear in these reports as 'Deleted Application'.

- To display additional details in the table for each user such as round trip time (RTT), transaction delay and efficiency statistics, click **Show Details**.


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.

- Click **Monitor > Network**.
- In the Select Graph to Display list, select **Users - Internal** or **Users - External**.
- ["Set the Time Period Reflected in the Report" on page 76](#).
After the date range is select, the graphs and charts are immediately updated.
- Remove specific types of traffic from the graph by deselecting their checkbox in the legend below the graph.
- 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














Real-time Traffic by Hosts

The Real Time Hosts/Users Report shows a breakdown of the Hosts/Users monitored by the Exinda appliance during the last 10 seconds. Hosts/Users are divided into Internal and External Hosts/Users.

Auto-Refresh Rate: | Show Users

1. Click **Monitor > Real Time > Hosts/Users**.

Hosts/Users are sorted by throughput, and display the packet rate and number of flows for each Host/User. The Distribution column shows the percentage of throughput a Host/User consumed relative to all the other Hosts/Users.

Inbound Hosts/Users				
IP Address (User)	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)
Total	138.037	46	117	
172.16.0.246 (Ksiakou)	105.324	10	5	
172.16.0.134 (Pforto)	13.909	3	4	
172.16.1.70 (Selfservice)	6.639	18	3	
172.16.1.240	3.771	6	34	
172.16.0.211	3.554	3	12	
172.16.0.244 (Cniko)	1.295	2	15	
172.16.0.127 (Sshannon)	1.060	2	20	
172.16.1.74	0.684	0	1	
172.16.0.239 (Jbothe)	0.593	1	5	
172.16.0.63 (Lenehan)	0.493	0	1	
Other	0.715	2	9	

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To display the user name associated with an internal IP, select **Show Users**.









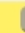

Note Active Directory must be configured on the Exinda appliances before user names can be displayed in reports. See "[Integrate the Exinda Appliance with Active Directory](#)" on page 165.

View real-time inbound and outbound conversations










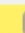
The Real-time Conversations Report shows a breakdown of the Conversations monitored by the Exinda appliance during the last 10 seconds. Conversations are divided into Inbound and Outbound directions.

- Click **Monitor > Real Time > Conversations**.

By default, the Real-time Conversations Report looks like the example below. Conversations are sorted by throughput. You can also see the packet rate and number of flows for each Conversation. Any extra information about a Conversation (a URL for example) will be shown in square brackets next to the Application.

Inbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			1408.428	284	24	
 	192.168.10.1	192.168.10.128	MAPI	570.834	82	1
 	192.168.10.9	192.168.10.128	MAPI	483.247	54	2
 	192.168.10.7	192.168.10.128	MAPI	275.334	92	2
 	192.168.10.10	192.168.10.128	MAPI	65.153	51	2
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	5.496	1	1
	192.168.10.9	192.168.10.128	LDAP	2.939	1	1
	10.20.4.1	239.255.255.250	udp ports 62612 -> 3702	1.097	0	1
	10.20.4.1	239.255.255.250	udp ports 62610 -> 3702	1.069	0	1
	192.168.10.1	192.168.0.1	NetBIOS	0.623	1	1
	192.168.10.10	192.168.10.128	LDAP	0.556	0	2
	192.168.10.132	255.255.255.255	DHCP	0.541	0	1
	192.168.10.9	192.168.0.1	NetBIOS	0.225	0	1
	10.20.3.118	10.20.255.255	NetBIOS	0.225	0	1
	192.168.10.9	192.168.255.255	NetBIOS	0.225	0	1
	10.20.11.100	224.0.0.252	udp ports 58633 -> 5355	0.212	0	1
	10.20.0.14	10.20.255.255	NetBIOS	0.193	0	1
	192.168.10.9	192.168.10.128	LDAP	0.174	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.106	0	1
	192.168.10.9	192.168.0.1	DNS	0.102	0	1
	10.20.0.181	10.20.255.255	NetBIOS	0.075	0	1

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To view only a specific IP address or subnet, type the address in the **IP/Subnet Filter** field.
The report can be filtered by IPv4 or IPv6 addresses.
- To display the optimization policy the conversation falls into, select **Show Policies**.


Outbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			3594.527	412	14	
 	192.168.10.7	192.168.10.128	MAPI	1826.409	196	2
 	192.168.10.10	192.168.10.128	MAPI	1184.445	125	2
 	192.168.10.1	192.168.10.128	MAPI	564.195	72	1
 	192.168.10.9	192.168.10.128	MAPI	12.200	17	2
	192.168.10.9	192.168.10.128	LDAP	3.316	1	1
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	2.902	1	1
	192.168.10.10	192.168.10.128	LDAP	0.565	0	2
	192.168.10.9	192.168.0.1	DNS	0.197	0	1
	192.168.10.9	192.168.10.128	LDAP	0.188	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.109	0	1

- To display the user name associated with an internal IP, select **Show Users**.


6. To group individual connections within a flow as a single line item or show each connection as a separate line item, select **Group**.

Understanding the Conversation Report






When a conversation has been accelerated by the Exinda appliance, the Conversation are highlighted in yellow and the Application Acceleration technologies being 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 is also been process by WAN Memory.







When a conversation has been processed by Edge Cache it is highlighted in blue.



The following legend describes the meaning of each 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 also appear next to each conversation.

	Asymmetric: The traffic is asymmetric, and is not being accelerated.
	Local: The connection is passing through this appliance in the cluster.
	Remote: The connection is passing through another appliance in the cluster.
	Local/Remote: The connection is passing though both this and other appliances in the cluster.

Controlling Traffic based on Users

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. See ["Integrate the Exinda Appliance with Active Directory"](#) on page 165.

1. ["Create Network User Objects"](#) on page 11 and ["Create Network Group Objects"](#) on page 12.
2. ["Optimize Traffic Based on Users and Groups"](#) on page 309

Create Network User Objects

Network Users displays a pre-populated list of Users (and their associated IP addresses) from either the Exinda Active Directory 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, it 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"/>
<input type="checkbox"/>	Dev_user_2 (BRANCH1)	172.1.1.19	<input type="checkbox"/>
<input type="checkbox"/>	Dev_user_3 (BRANCH2)	172.1.1.13	<input type="checkbox"/>
<input type="checkbox"/>	Dev_user_4 (BRANCH2)	172.1.1.14	<input type="checkbox"/>
<input type="checkbox"/>	Dev_user_5 (BRANCH1)	172.1.1.15	<input type="checkbox"/>
<input type="checkbox"/>	Qa_user_7 (BRANCH1)	172.1.1.9	<input type="checkbox"/>
<input type="checkbox"/>	Qa_user_8 (BRANCH1)	172.1.1.10	<input type="checkbox"/>
<input type="checkbox"/>	Qa_user_9 (BRANCH1)	172.1.1.11	<input type="checkbox"/>

To define a user as a Dynamic Network Object


1. In the Exinda WebUI, go to **Objects > Users & Groups > Network Users**.
2. Select the checkbox for the user.

3. Click **Add Network Object**.

The Network Status icon for the user changes to , indicating it is a network object.

To stop identifying a user as a dynamic network object

1. Select the checkbox for the user.
2. Click **Remove Network Object**.

The Network Status icon for the user changes to , indicating it is no longer a network object.

Create Network Group Objects

Network Groups displays a pre-populated list of Groups from either the Exinda Active Directory 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.

To define a group as a Dynamic Network Object


1. In the Exinda WebUI, go to **Objects > Users & Groups > Network Groups**.
2. Locate 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 , indicating it is 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 stop identifying a group as a Dynamic Network Object

1. Locate the group in the list, and click **Delete**.

The Network Status icon for the user changes to , indicating it is no longer a network object.

If the dynamic network object was created from multiple groups, each group is again listed individually in the list.

Optimize Traffic Based on Users and Groups

Create policies that affect the traffic based on the source or destination host.

Note Active Directory must be configured before optimization policies can target specific users and groups. See ["Integrate the Exinda Appliance with Active Directory"](#) on page 165.

1. Click **Optimizer > Policies**.
2. Type a name for the policy.
3. Set the required bandwidth and acceleration parameters.
4. 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.
5. Click **Create New Policy**.
6. Once the desired policies are in place on all Exinda appliances, restart the Optimizer. In the appliance status bar, click **Restart**.

Optimizer Status : On (Restart / Stop)

Troubleshoot issues with Active Directory configuration

If you are experiencing issues with the Active Directory integration, these troubleshooting topics may help resolve the issue.

- ["The Exinda appliance reboots every night" on page 188](#)
- ["WMI Service is not running" on page 188](#)
- ["No communication between the Exinda Active Directory Connector and the Exinda appliance" on page 189](#)
- ["Changes to the Exinda Active Directory Controller have no effect" on page 190](#)
- ["Excluded users still appear on the Exinda appliance" on page 190](#)
- ["Exinda Active Directory Connector stops running" on page 189](#)

The Exinda appliance reboots every night

Problem

When multiple installations of the Exinda Active Directory Connector have the **Send Active Directory user and group information to Exinda appliance(s) at startup** option selected, the Exinda appliance is overwhelmed with duplicate data from the connectors and causes the appliance to shut down.

Resolution

1. On each instance of the Exinda Active Directory 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 Active Directory Connectors.
3. Choose one instance of the Exinda Active Directory Connector, and select the **Send Active Directory user and group information to Exinda appliance(s) at startup** checkbox, and click **OK**.

WMI Service is not running

Problem

When I try to access the Exinda Active Directory Connector, the message "The installer has detected that WMI Service is not running. Consult Windows Help files to find information on how to start WMI Service." is displayed.

Resolution

This message indicates that Windows Management Information (WMI) service is disabled. The Exinda Active Directory 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`

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.

Resolution

When SMB signing is configured and enabled, the SMB signing account is the last user account registered as using an IP address, the Exinda Active Directory 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 Active Directory Connector to ignore the SMB signing account. See "[Exclude specific usernames from reports](#)" on page 171.

No communication between the Exinda Active Directory Connector and the Exinda appliance

Problem

You see one of the following symptoms:

- A connection cannot be established between the Exinda Active Directory Connector and the Exinda appliance.
- The Last Contact status on the **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 Active Directory Connector

Exinda Active Directory Connector stops running

Problem

Even after restarting the Exinda Active Directory Connector or the Exinda AD service the Exinda Active Directory Connector does not continue running, and requires constant restarts.

Resolution

1. The Exinda Active Directory 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 Active Directory Connector.
2. If the Active Directory server is running Windows 2003 R2, ensure the Exinda Active Directory 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.

Excluded users still appear on the Exinda appliance

Problem

Even though a user name has been added to the Excluded list on the Exinda Active Directory Connector, the username continues to appear associated with traffic on the Exinda appliance.

Resolution

1. Verify that the username on the Excluded tab of the Exinda Active Directory Connector matches the username in Active Directory.

The username is 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.

2. If the case matches on the usernames, restart the AD Client Service and reenumerate the Exinda appliance. See ["Change the state of the Exinda Active Directory Connector" on page 170](#) and ["Request updated user and group information from the Active Directory server" on page 170](#).

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.

Resolution

1. To ensure the latest configuration is being used, restart the AD Client Service and reenumerate the Exinda appliance. See ["Change the state of the Exinda Active Directory Connector" on page 170](#) and ["Request updated user and group information from the Active Directory server" on page 170](#).

IPMI Overview

The Intelligent Platform Management Interface (IPMI) is a specification for remote server management. The specification is maintained by a consortium of computer system vendors led by Intel. The current revision is 2.0. Further information on the IPMI specification is available [here](#).

An IPMI enabled server contains a separate, dedicated micro-controller for monitoring and controlling the hardware, usually referred to as the Baseboard Management Controller (BMC). The BMC monitors system components (e.g. power supplies, fans, temperatures) and makes this data available over the LAN using either a shared or dedicated NIC. In either case the IPMI interface must be assigned a dedicated IPv4 address.

The table below shows the Exinda models that support IPMI and their capabilities.

Model	NIC	User Interface
406X	Shared (eth1)	CLI
606X	Shared (eth1)	CLI, Web
806X	Shared (eth1)	CLI, Web
1006X	Dedicated	CLI, Web

Configure IPMI

To configure IPMI connectivity on the Exinda appliance, navigate to System | Network | IPMI.

Use the form to enable the IPMI LAN interface and to configure an IPv4 address and gateway. Note that although the IPMI LAN interface may be the same as the appliances eth1 interface, a separate IPv4 address is required. Alternatively you may use DHCP to configure the IPMI interface, although this is not recommended.

Example

Enable IPMI with an IPv4 address of 172.16.0.71, subnet 255.255.254 and gateway 172.16.1.254.

IPMI Network Settings	
Enable	<input checked="" type="checkbox"/>
DHCPv4	<input type="checkbox"/>
IPv4 Address	<input type="text" value="172.16.0.71"/> / <input type="text" value="23"/>
IPv4 Gateway	<input type="text" value="172.16.1.254"/>
Admin User	admin

This will allow the appliance to be managed using IPMI on the 172.16.0.71 address, either from another Exinda appliance or by connecting to <http://172.16.0.71> (where a Web interface is supported). When a Web interface is not supported, a command line tool such as `ipmitool` may be used.

Use the form below to change the authentication details. The default is username: admin, password: exinda.

Change IPMI Administrator Details	
Administrator User Name	<input type="text" value="admin"/>
New Password	<input type="password" value="••••••"/>
Confirm Password	<input type="text"/>

The equivalent CLI commands are:

```
> en
# configure terminal
(config) # ipmi enable
(config) # ipmi ip address 172.16.0.71 255.255.254.0
(config) # ipmi ip default-gateway 172.16.1.254
(config) # ipmi username admin password exinda
```

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 **System > Tools > IPMI**.

Power Control Options	
Command	<input type="text" value="Get Status"/> <input type="button" value="v"/>

Remote IPMI Login Details	
IPv4 Address	<input type="text"/>
Username	<input type="text"/>
Password	<input type="text"/>

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

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
```

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
```

System Setup

The System Setup section of the Exinda appliance allows you to configure basic and advanced system settings. The various configuration pages include:

- ["Date and Time Configuration" on page 193](#): Configure Date and Time on the Exinda appliance.
- ["UI Configuration" on page 195](#): Configure Web UI and CLI settings.
- ["SDP Configuration" on page 197](#): Enable SDP.
- ["Configure SQL Access" on page 197](#): Configure remote SQL access.
- ["Monitoring Configuration" on page 216](#): Fine-tune monitoring settings.
- ["Netflow Configuration" on page 221](#): Configure netflow parameters and items to export.
- ["Create a Scheduled Job" on page 271](#): View and remove scheduled jobs.
- ["Notify administrators of system issues" on page 279](#): Enable/Disable system alerts.
- ["License" on page 226](#): Install a new license.
- ["Control Configuration" on page 229](#): Change Optimizer mode
- ["Allocate Disk Storage for System Services" on page 229](#): Configure dynamic partitions.

Date and Time Configuration

Use the form below to set the time, date and time zone on the Exinda appliance.

Note If NTP time synchronization is enabled, the date and time cannot be manually configured.

- ["Set the date and time of the appliance" on page 194](#)
- ["Add an NTP server" on page 194](#)
- ["Disable an NTP server" on page 194](#)
- ["Remove an NTP server" on page 194](#)

Set the date and time of the appliance

If you are not using an NTP server to manage the date and time information on the Exinda appliance, specify the date, time, and time zone of the Exinda appliance.

Note If NTP time synchronization is enabled, the date and time cannot be manually configured.

1. Click **System > Setup** and switch to the **Date and Time** tab.
2. Click **Apply Changes**.
3. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Add an NTP server

Configure NTP servers that the Exinda appliance will use to set its time and date.

1. Click **System > Setup** and switch to the **Date and Time** tab.
2. In the Date and Time area, select the **NTP Time Synchronization** checkbox.
3. In the Add New NTP Server area, type the IP address or hostname of the NTP server.
Only hostnames and IPv4 addresses are supported.
4. Select the version of NTP supported by the server.
5. To enable the NTP server, select the checkbox.
6. Click **Add New NTP Server**.
7. Click **Apply Changes**.
8. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Disable an NTP server

Stop an existing NTP server from setting the date and time of the Exinda appliance.

1. Click **System > Setup** and switch to the **Date and Time** tab.
2. To disable the NTP server, select the server from the list and click **Disable Server**.
3. If there are no NTP servers enabled, in the Date and Time area, clear the **NTP Time Synchronization** checkbox.
4. Click **Apply Changes**.
5. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Remove an NTP server

Remove any servers that are no longer used for setting the date and time of the Exinda appliance.

1. Click **System > Setup** and switch to the **Date and Time** tab.
2. Select the server from the list and click **Remove Server**.
3. If there are no NTP servers enabled, in the Date and Time area, clear the **NTP Time Synchronization** checkbox.
4. Click **Apply Changes**.
5. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

UI Configuration

Use the form below to enable/disable the Web User Interface as well as configure Web UI options.

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

Web UI	Enable web access to the Exinda appliance.
Auto Logout Timeout	The time (in minutes) before the web session is automatically logs out. A value of '0' means the Web UI will never auto-logout.
HTTP Access	Enable HTTP access to the Exinda appliance. This is disabled by default.
HTTP Port	Configure the HTTP port. The default is 80.
HTTPS Access	Enable HTTPS access to the Exinda appliance. This is enabled by default.
HTTPS port	Configure the HTTPS port. The default is 443.

Web Session Renewal	Web session renewal time in minutes. This cannot be more than the Web Session Timeout.
Web Session Timeout	Web session timeout in minutes. This cannot be less than the Web session renewal.

Use the form below to enable/disable the CLI as well as configure CLI options.

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"/> ▼

Auto Logout Timeout	The time (in seconds) before the CLI session automatically logs out.
Telnet Access	Enable telnet access to the Exinda appliance. This is disabled by default.
SSH Access	Enable SSH. This is enabled by default.
SSH version	Configure a SSH version. This can be set to 'SSH v2 or v1' or 'SSH v2 only'

Configure when the Exinda Web UI logs out

Set how long the Exinda appliance should remain logged in before the system is automatically logged out due to inactivity.

1. Click **System > Setup** and switch to the **Access** tab.
2. In Web UI Options area, specify how many minutes of inactivity pass before the user is automatically logged out in the **Auto Logout Timeout** field.
To configure the system to never automatically log out, set the field to **0** minutes.
3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Configure when the CLI console times out

Set how long the CLI console should remain logged in before the system is automatically logged out due to inactivity.

1. Click **System > Setup** and switch to the **Access** tab.
2. In CLI Options area, specify how many minutes of inactivity pass before the user is automatically logged out in the **Auto Logout Timeout** field.
To configure the system to never automatically log out, set the field to **0** minutes.
3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

SDP Configuration

Use the form below for 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 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.

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 XP and Windows Vista/7 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

In order to allow the Exinda appliance to accept remote SQL connections from an external ODBC connector, you must configure the following settings.

On the Exinda appliance, using the Web User Interface, navigate to System | Setup | SQL Access. You will be presented with the following form.

Remote SQL Options	
Remote SQL	<input type="checkbox"/> Enable
Allow access from (Hostname or IP)	<input type="text"/> (% = 'any')
Username	<input type="text"/>
Password	<input type="text"/>
Confirm Password	<input type="text"/>

[Apply Changes](#)

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 enter 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:

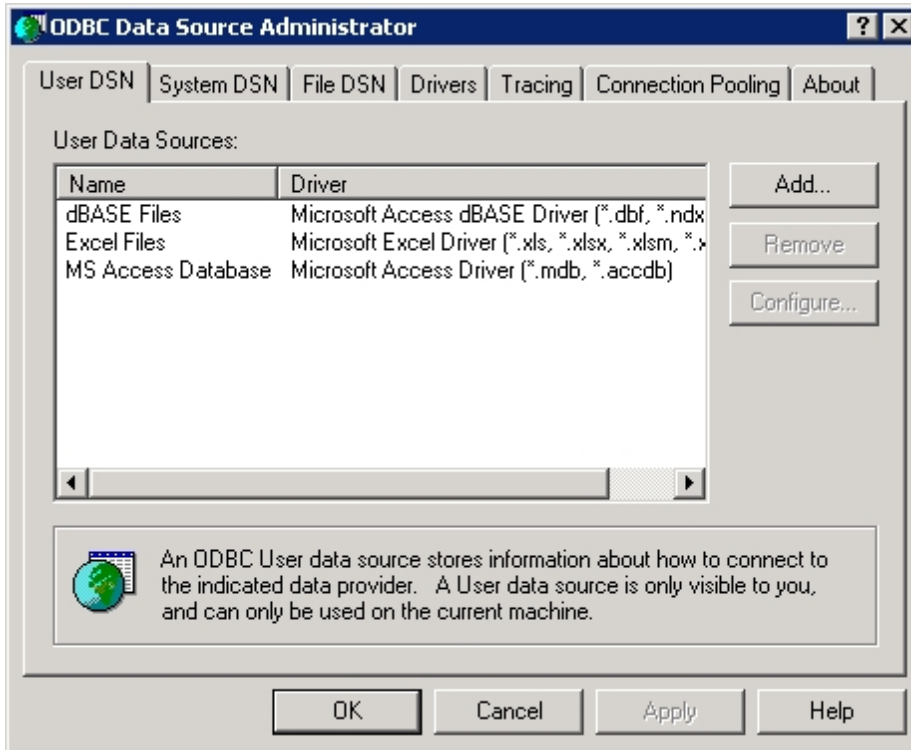
Remote SQL Options	
Remote SQL	<input checked="" type="checkbox"/> Enable
Allow access from (Hostname or IP)	<input type="text" value="%"/> (% = 'any')
Username	<input type="text" value="database"/>
Password	<input type="text" value="•••••"/>
Confirm Password	<input type="text" value="•••••"/>

[Apply Changes](#)

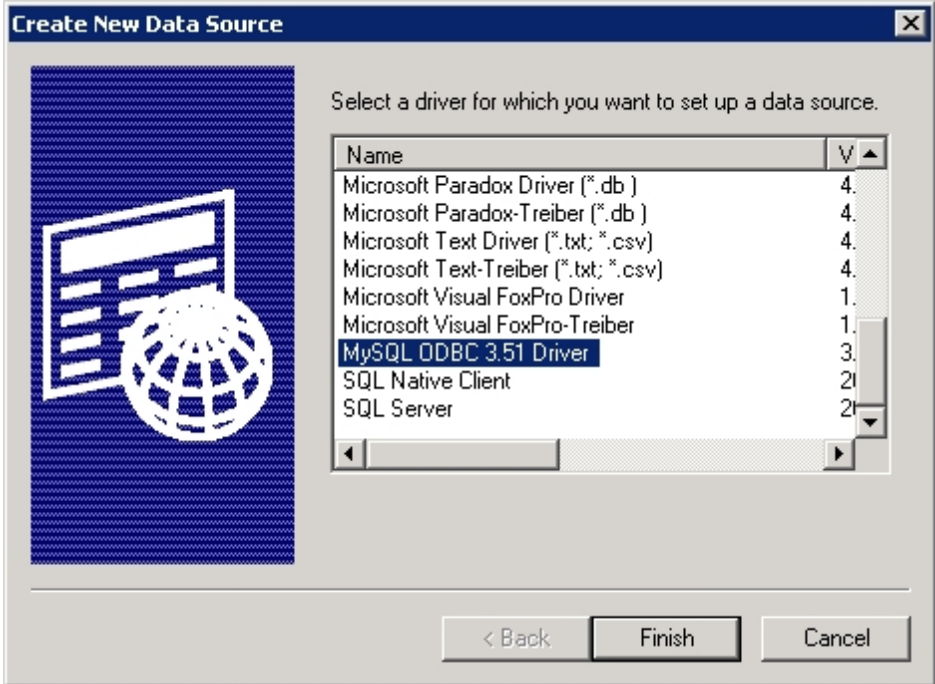
Once remote SQL access has been configured on the Exinda appliance, the next step is to create an ODBC data source on the client.

Create ODBC Data Source on Windows XP

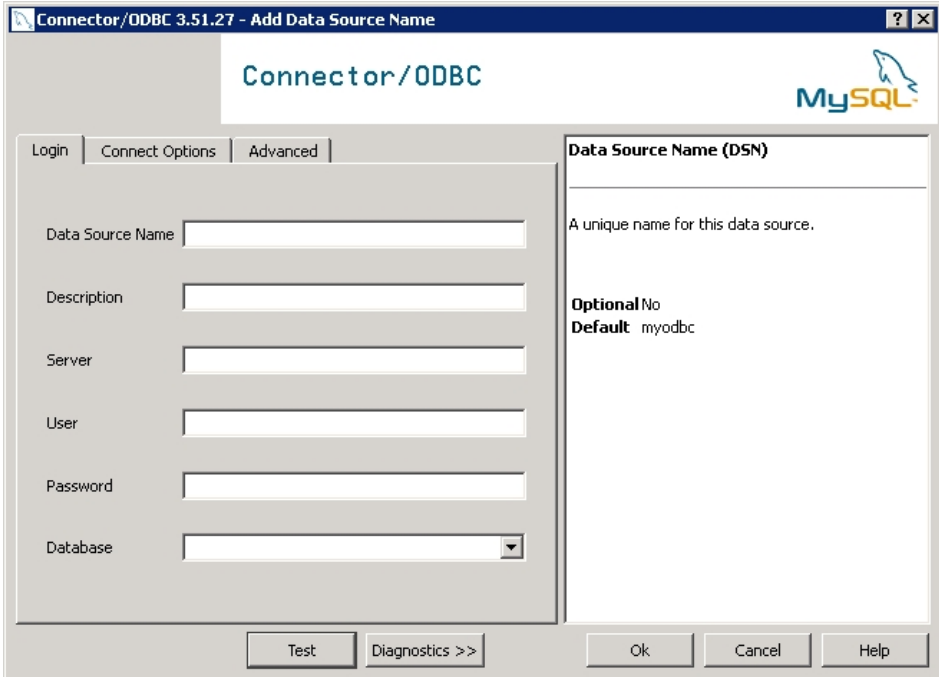
Open **Administrative Tools** and select **Data Sources (ODBC)**. You should be presented with the following dialog.



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.



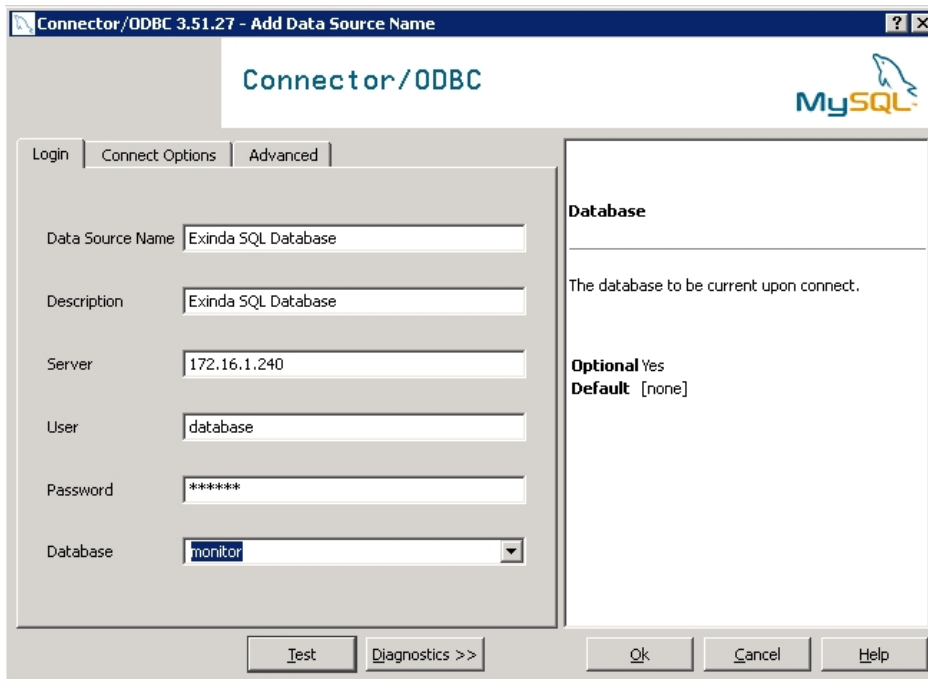
Select **MySQL ODBC Driver** and click **Finish**. You will be prompted to enter details about the SQL access using the form below:



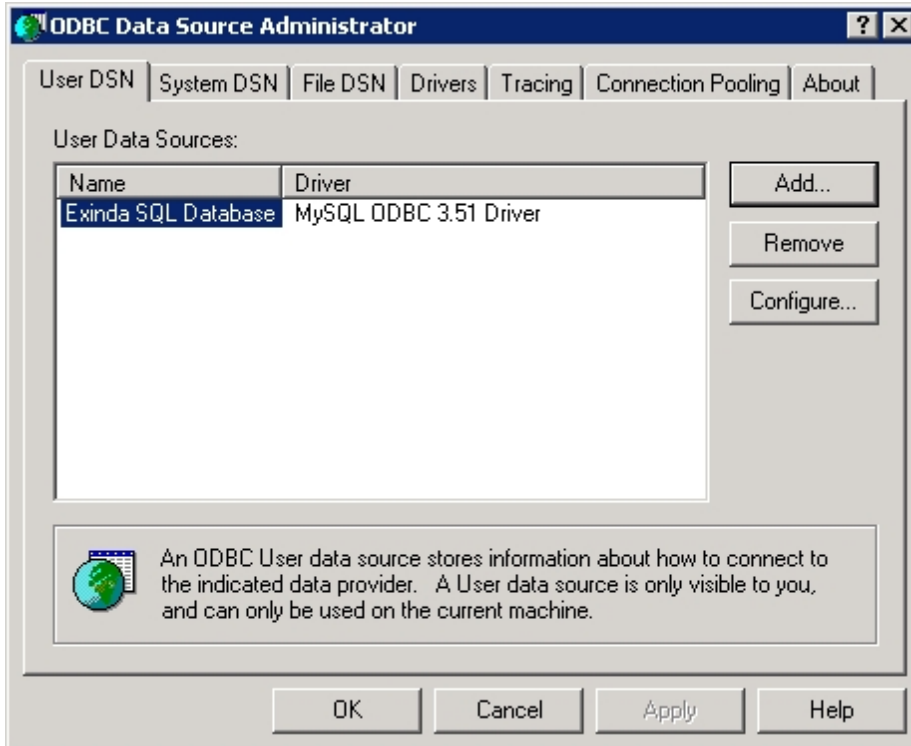
Data Source Name / Description	Enter a descriptive name for the DSN. E.g. 'Exinda SQL Database'.
--------------------------------	---

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:

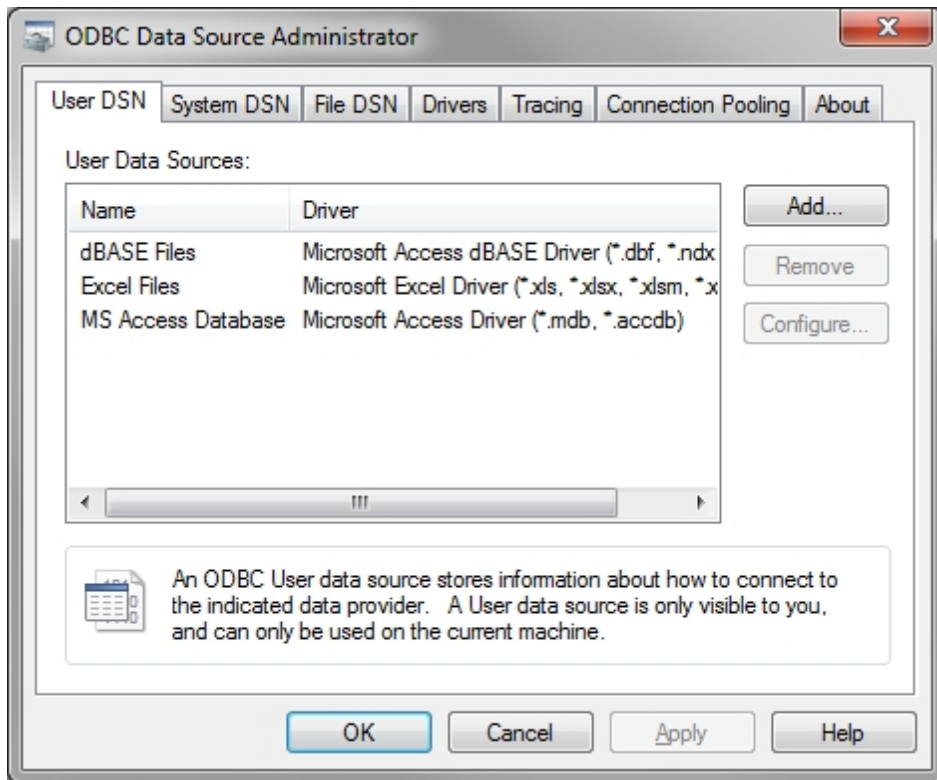


Click **OK**. This will add the 'Exinda SQL Database' to the list of available data sources that can be used by 3rd party applications on this client.

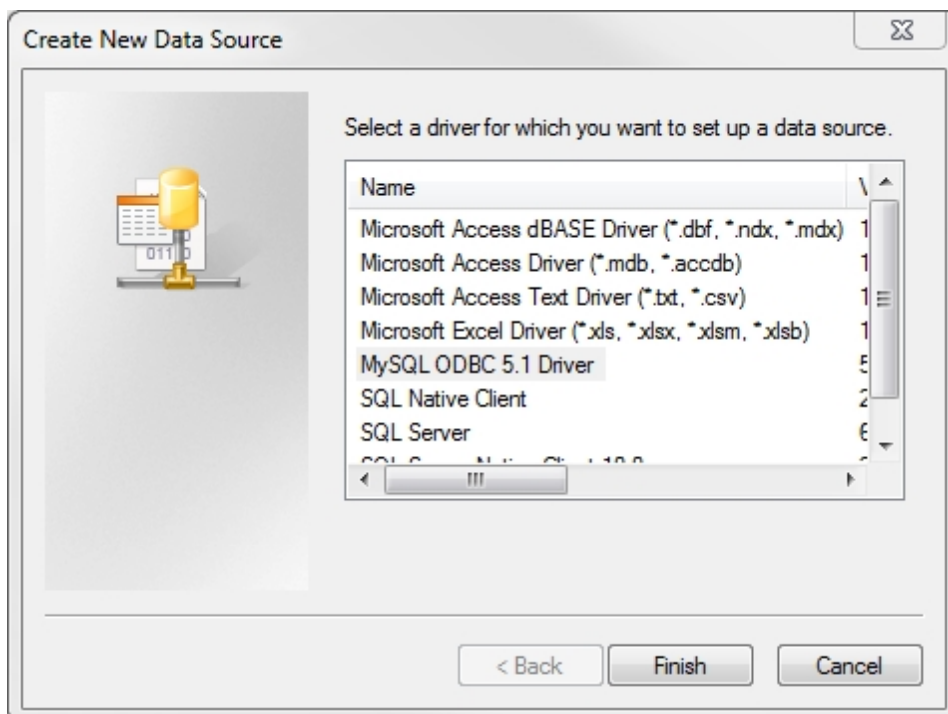


Create ODBC Data Source on Windows 7

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.



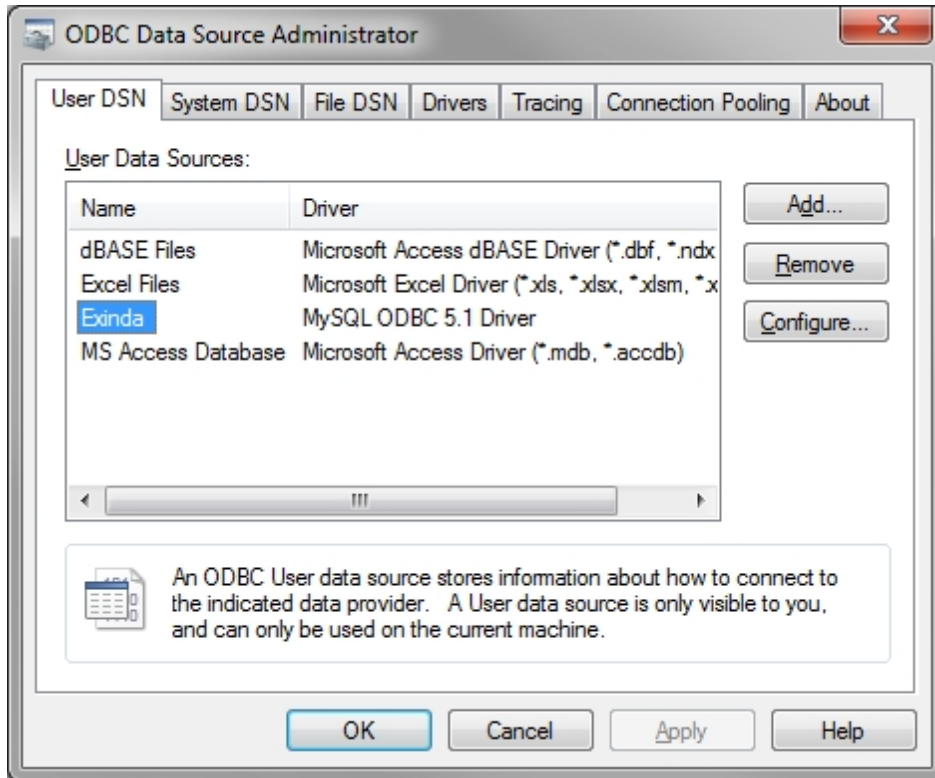
Select **MySQL ODBC Driver** and click **Finish**. You will be prompted to enter details about the SQL access using the form below:

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:



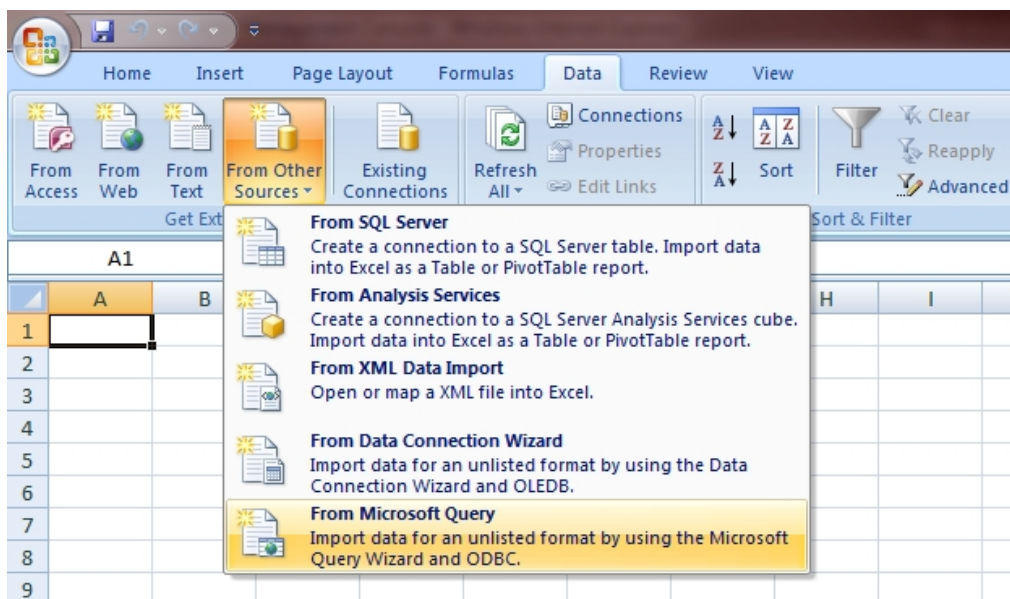
Click **OK**. This will add the 'Exinda SQL Database' to the list of available data sources that can be used by 3rd party applications on this client.



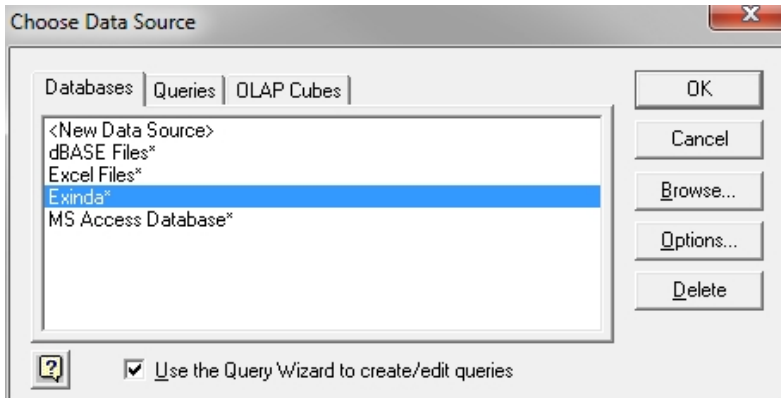
View SQL Access data in Microsoft Excel

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.

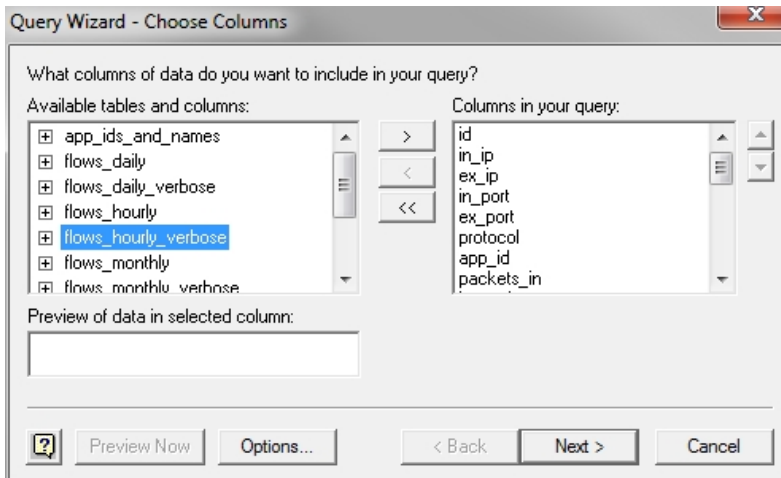
From the **Data** tab in Excel, select **From Other Sources > From Microsoft Query**.



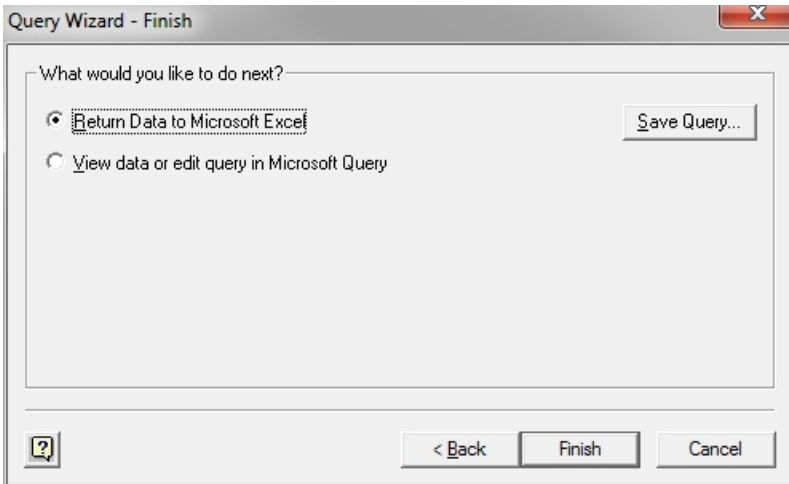
You will be presented with a dialog box that allows you to select the DSN you created in the previous chapter.



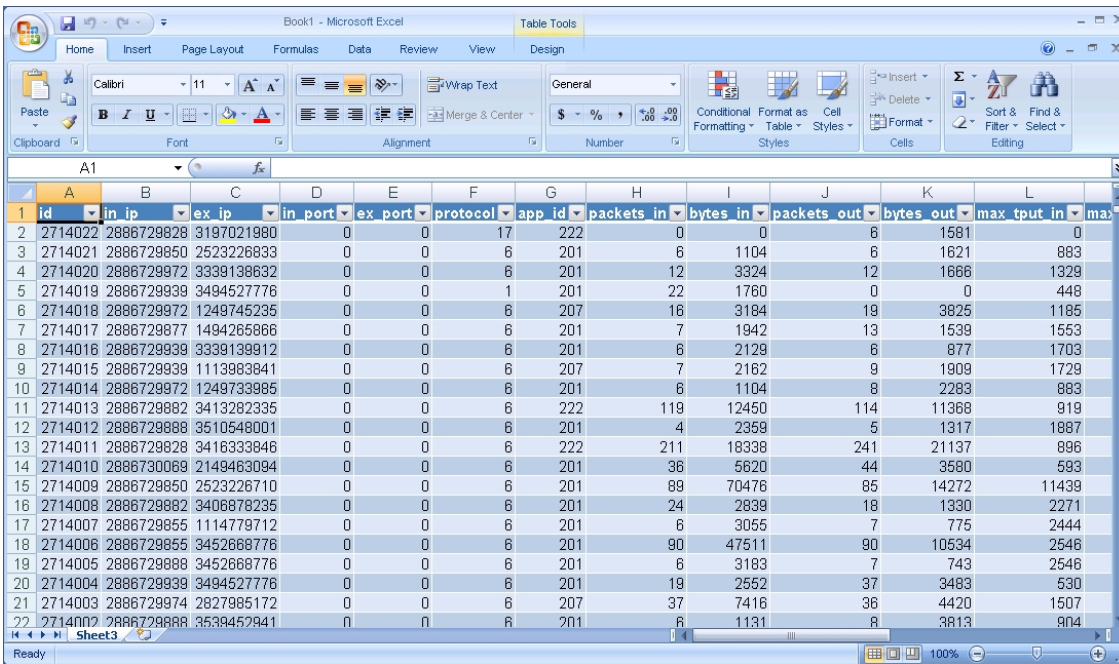
Select the **Exinda SQL Database** DSN. This will allow you to choose from the available tables and select the columns to query. Select a table and click the > button to move that table's fields into the list of columns to query.



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.



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.
flows_	Flow records at monthly resolution, that is, information for each flow is stored monthly,

Name	Description
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. ¹
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. ¹
protocol	unsigned	The IANA assigned IP protocol number of the flow. See

Field	Type	Description
	24-bit integer	http://www.iana.org/assignments/protocol-numbers/ 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.
bytes_in	unsigned 64-bit integer	The number of inbound (WAN -> LAN) bytes 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.
bytes_out	unsigned 64-bit integer	The number of outbound (LAN -> WAN) bytes recorded for this flow over 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.
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.
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). ¹
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. ²

Field	Type	Description
network_delay	unsigned 32-bit integer	A normalized measure of the time taken for transaction data to traverse the network. ²
network_jitter	unsigned 32-bit integer	A normalized measure of the network_delay variability. ²
server_delay	unsigned 32-bit integer	A normalized measure of the time taken for a server to respond to a transaction request. ²
bytes_lost_in	unsigned 64-bit integer	The number of bytes lost due to retransmissions (WAN -> LAN). ²
bytes_lost_out	unsigned 64-bit integer	The number of bytes lost due to retransmissions (LAN -> WAN). ²
aps	unsigned 64-bit integer	Application Performance Score. A measure of an applications performance on the network. ²

¹ 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).

² See the APS HowTO Guide for further information.

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
in_ip	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.
ex_ip	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
app_id	unsigned 24-bit	A unique id that defines the Application. This is the primary key.

Field	Type	Description
	integer	
app_name	string	The Application name (e.g HTTP, Hotmail)
deleted_flag	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
id	unsigned 32-bit integer	This id references an id in the corresponding parent flows_* table. There can be multiple url records referencing the same flow id, so this field is not unique.
url	string	The URL (host) extracted from the HTTP header of the parent flow.
packets_in	unsigned 64-bit integer	The number of inbound (WAN -> LAN) packets recorded for this URL over the sample period.
bytes_in	unsigned 64-bit integer	The number of inbound (WAN -> LAN) bytes recorded for this URL over the sample period.
packets_out	unsigned 64-bit integer	The number of outbound (LAN -> WAN) packets recorded for this URL over the sample period.
bytes_out	unsigned 64-bit integer	The number of outbound (LAN -> WAN) bytes recorded for this URL over the sample period.
max_tput_in	unsigned 64-bit integer	The maximum inbound (WAN -> LAN) throughput observed for this URL during the sample period.
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 id's are only consistent across the same sample periods. For example, id's in the urls_hourly table only reference id's 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) ¹
ex_port	unsigned 24-bit integer	The TCP or UDP port number on the external side (the WAN side of the Exinda appliance) ¹
protocol	unsigned 24-bit integer	The IANA assigned IP protocol number of the flow. See http://www.iana.org/assignments/protocol-numbers/ 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	A UNIX timestamp (number of seconds since epoch - 1st Jan 1970) that represents the start of the sample period.

Field	Type	Description
	integer	
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).
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. ²
network_delay	unsigned 32-bit integer	A normalized measure of the time taken for transaction data to traverse the network. ²
network_jitter	unsigned 32-bit integer	A normalized measure of the network_delay variability. ²
server_delay	unsigned 32-bit integer	A normalized measure of the time taken for a server to respond to a transaction request. ²
bytes_lost_in	unsigned 64-bit integer	The number of bytes lost due to retransmissions (WAN -> LAN). ²
bytes_lost_out	unsigned 64-bit integer	The number of bytes lost due to retransmissions (LAN -> WAN). ²

¹ 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).

² See the APS How To Guide for further information.

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.

Field	Type	Description
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.
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

Field	Type	Description
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

¹ See the APS How To Guide for further information.

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, whether data is collected, and whether collected data is deleted.

- ["Configure monitoring settings" on page 216](#)
- ["Control the order that IP addresses are resolved" on page 217](#)
- ["Enable or disable Application Specific Analysis Modules" on page 218](#)
- ["Identify the statistics to collect" on page 219](#)
- ["Clear saved monitor statistics" on page 220](#)

Configure monitoring settings

Configure the parameters for what is displayed in the monitoring reports and when information is displayed in the reports.

1. Click **System > Setup** and switch to the **Monitoring** tab.
2. In the Monitoring Options area, select the parameters for how reports are displayed.
 - a. Specify the maximum number of top items displayed in the monitoring tables in the **Table Items** field. Acceptable values are 1-1000.
 - b. Select the maximum number of top items displayed in the pie chart graphs from the **Chart Items** list. Acceptable values are 1-10.
 - c. In the **Maximum URL Size** field, specify the maximum length of the URLs displayed on the Real Time report tables.
 - d. In the Graph Display Options list, select whether the graphs display in **Flash** or **non-Flash** format. The default is Flash.
 - e. From the **Display for applications details per subnet** list, select whether the charts display as a **Time series chart** (line chart), or as a **Pie graph**. When this option is selected, the Applications per subnet chart displays as a line chart. All other charts continue to display as a pie graph.
 - f. To sort the list of subnets within reports by name, select the **Enable** checkbox.

- g. To reduce CPU usage, disable **Detailed Record Retention** if there are excessive traffic flows through the appliance.

When this is disabled, detailed monitoring records for Applications, Hosts, URLs, Users, Conversations, and Subnets are not stored, and the drill-down options are not available.

3. Specify what details to include in the reports.

- a. To analyze the application signatures within a packet to further classify the traffic within the reports, select the **Layer 7 Inspection** checkbox. For example, when analyzing HTTP or 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.

- b. 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 reports.

When this option is enabled, traffic between local subnets is also ignored when the router is on the WAN side of the Exinda appliance and has been marked as internal.

- c. Set the **Bittorent Sensitivity** level.

Setting this to *high* is recommended for most service provider environments. Setting it to *low* is recommended in cases of high false positives.

- d. Set the **EDonkey Sensitivity** level.

Setting this to *high* is recommended for most service provider environments. Setting it to *low* is recommended in cases of high false positives.

- e. Set the **Skype Sensitivity** level.

Setting this to *high* is recommended for most service provider environments.

- f. Specify the minimum number of packets in a flow before the Exinda appliance records the flow in the database with the **Reporting Sensitivity** option. 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.

4. Click **Apply Changes**.

5. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Control the order that IP addresses are resolved

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.

1. Click **System > Setup** and switch to the **Monitoring** tab.
2. Set the order of resolution methods tried when resolving IP addresses to hostnames.

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.

3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Enable or disable Application Specific Analysis Modules

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.

1. Click **System > Setup** and switch to the **Monitoring** tab.
2. To enable an Application Specific Analysis Module (ASAM), select the appropriate checkbox.

The following ASAM modules are available:

- **Anonymous Proxy**

When enabled, the Exinda appliance attempts to match 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 watches for client requests for Microsoft services such as MAPI and SMB.

- **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. See the "[RAM Usage Report](#)" on page 67 or "[CPU Usage Report](#)" on page 66.

- **SSL**

When enabled, this module extracts public certificates from connections identified as SSL and decodes the information from those certificates.

- **VoIP**

When enabled, this module extracts VoIP related information such as code type and call quality information 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.

3. Click **Apply Changes**.

4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Identify the statistics to collect

Collect various types of statistics data for traffic passing over the network. Statistics collection can be disabled if the appliance is not performing as expected, and these statistics are not necessary.

1. Click **System > Setup** and switch to the **Monitoring** tab.

2. Select the statistics to be collected on the appliance.

- **Network Object/Subnet Application Stats** - Applications by network objects and subnets are collected. By default the collection is enabled.

When disabled, the Subnets report will not include application data for the time period the collection was disabled.

3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Clear saved monitor statistics

Remove all the monitoring statistics from the database, or remove statistics for selected record types.

1. Click **System > Setup** and switch to the **Monitoring** tab.
2. Select the records that should be cleared from the monitoring database. To select all the listed records, select the top checkbox.
 - **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.

Caution This will permanently delete the selected records from the monitoring database.

3. Click **Clear Records**.

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	<input type="text"/>
Port	<input type="text" value="2055"/>
Version	<input type="text" value="9"/>

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	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	1440 bytes
Template Refresh Rate	100 packets
Template Timeout Rate	600 seconds
General Options Refresh Rate	10000 packets
General Options Timeout Rate	600 seconds
Username Options Timeout Rate	1440 minutes
Inactive Username Expiry Rate	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

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	
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	
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	
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.

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.

1. Click **System > Setup** and click the **Scheduled Jobs** tab.
2. In the Add New Job area, type a unique **ID** for the job.
3. Type a name for the job.
4. [Optional] In the **Comment** field, type a description for the job.
5. To run the job immediately, **Enable** the job.
6. If the job should be completed, even if one or more commands fail to execute, set **Fail-Continue** to **Yes**.
7. Set the frequency of the scheduled job. Jobs can be set to run Once, Daily, Weekly, Monthly, or Periodically.
8. After selecting the frequency 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.
9. In the **Commands** field, type the necessary commands for the job you want run. Each command must be on a new line.

- Click **Add Job**.

The job is added to the list, and is now available for selection in the Pre-population

Notify administrators of system issues

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. SLA and APS email alerts are sent when the set threshold limits are exceeded. Use the form below to disable alerts that you do not wish to trigger or receive emails and SNMP traps for.

Note You must configure valid SMTP and DNS settings prior to receiving email alerts. See "Add an SMTP server for sending email notifications" on page 281 and "Configure DNS and Domain Names" on page 159.

- Click **System > Setup**, and switch to the **Alerts** tab.
- To receive alert notifications, select the **Enable** checkbox for the appropriate alert.
For a description of when each alert is triggered, refer to the table below.
- Select what type of notification to receive: an **Email**, an **SNMP Trap**, or both.
- For CPU Utilization, Disk Usage, or NIC Collisions alerts, specify the **Trigger Threshold** and **Clear Threshold** levels that cause the notifications to be sent.

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.

- Click **Apply Changes**.

Alert Name	Description
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.

Alert Name	Description
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.
System Startup	Alert raised when the Exinda appliance boots up.
SMB signed connections	Alert raised when SMB signed connections are present.
SLA Latency	Alert raised when the set latency for an SLA object is exceeded.
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.
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).
Connection Limiting	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.
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.

License

Licensing Exinda appliances is simple. A single License Key is required to enable features. Multiple License Keys on the same appliance are also supported. The appliance will use the license that provides the highest specification limits. The license is automatically fetched and installed on first bootup. The auto license service checks for new licenses every 24 hours, if a new license is found it is automatically installed. The

table below shows the time since the Auto License Service checked for a new license and the time since a new license was found. The Exinda appliance allows you to Stop, Restart or Disable the Auto License Service.

Auto License Service: Running <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)

The Current System License Status displays the "effective" license limits and enabled features. These are the limits that are currently effective on the appliance. These effective limits can change depending on the license key or combination of keys installed.

Licensed	Host ID	Model	SS Expiry
<input checked="" type="checkbox"/>	002219d48dc4	Exinda 4860 20Mbps	Jun 19, 2012
	Max Bandwidth:	20480 kbps	
	Optimizer:	<input checked="" type="checkbox"/>	
	Max AA Bandwidth:	20480 kbps	
	Max Connections:	384000	
	Max Connection Rate:	300 / sec	
	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:	<input checked="" type="checkbox"/>	

Licensed	License status.
Host ID	The Host ID is unique to each Exinda appliance.
Model	Exinda appliance model.
End Date	Expiry date of a temporary key.
SS Expiry	Expiry date of Exinda Software Subscription. After this date no updates can be installed on the appliance.
Max Bandwidth	Maximum monitoring and QoS bandwidth.
Optimizer	QoS and Acceleration module status.
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 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.
SSL Acceleration	SSL Acceleration license feature status.
Virtualization	Virtualization license feature status.
Edge Cache	Edge Cache Acceleration license feature status.

The available license keys are listed along with their respective 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 Tied to hex host ID: 002219d48do4 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"/>	EXINDA	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

The Exinda appliance allows you to manually check the Exinda license server, at any time, for updated licenses by clicking 'Check For License Online'. If you already have a license, you can copy and paste it in the text box. Click 'Add Licenses' to add the newly found or manually entered license key.

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 **System > System Setup > Control Configuration** on the Web UI, advanced mode.

The form below is used to enable/Disable Global Control:

Control Options	
Global Control	<input type="checkbox"/> Enable

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.

Allocate Disk Storage for System Services

The Exinda appliance has the capability to 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.

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](#)" on page 152.

Disk Storage Map.



Storage Configuration							Operation		
Service	Status	Free	Size	Minimum	Encrypted				
cifs	available	127.45G 98%	<input type="text" value="129.67G"/>	1024.00M	✘	<input type="button" value="Resize"/>	<input type="button" value="Format"/>	<input type="button" value="Encrypt"/>	
edge-cache	available	127.23G 98%	<input type="text" value="129.45G"/>	1024.00M	✘	<input type="button" value="Resize"/>	<input type="button" value="Format"/>	<input type="button" value="Encrypt"/>	
monitor	available	126.97G 98%	<input type="text" value="129.45G"/>	10.00G		<input type="button" value="Resize"/>	<input type="button" value="Format"/>		
users	available	974.62M 95%	<input type="text" value="1024.00M"/>	512.00M		<input type="button" value="Resize"/>	<input type="button" value="Format"/>		
virt	available	49.04G 98%	<input type="text" value="50.00G"/>	512.00M		<input type="button" value="Resize"/>	<input type="button" value="Format"/>		
wan-memory	available	467.01G 98%	<input type="text" value="474.65G"/>	5120.00M	✘	<input type="button" value="Resize"/>	<input type="button" value="Format"/>	<input type="button" value="Encrypt"/>	
unallocated storage			<input type="text" value="0.00"/>						
Total Available Storage:			914.22G						

Service	The system service using disk storage.
Status	<p>The disk storage may be in one of several states, depending on which operation has been selected.</p> <ul style="list-style-type: none"> ▪ available—The storage is online and available to the service. ▪ growing —The storage size was increased, and the filesystem is being reconfigured to use the newly created space. ▪ shrinking—The storage size was decreased, and the filesystem is being reconfigured to use the decreased amount of storage available. ▪ formatting—The storage is being formatted. ▪ checking—The storage filesystem is being checked for consistency. ▪ error—The 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—The storage is not available.
Free	The amount of free storage available, shown as the number of bytes as well as a percentage of available space.
Size	The total amount of storage available for this service.
Minimum	The minimum amount of storage required for this service.
Encrypted	Identifies whether the storage for the service is encrypted or not.
Operation	Selects an operation to perform on the storage - either resize, format, or encrypt.

1. Click **System > Setup** and switch to the **Storage** tab.
2. To increase or decrease the amount of storage available to a service, enter the new amount in the Size field, and click **Resize**.

The storage size can be specified in terms of kilobyte (KB), megabytes (MB), gigabytes (GB), or percentages (%). 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.

Note When decreasing the amount of storage available to a service, the service may stopped until the storage operation has completed.

The form below shows a summary of storage by disk partition.

Disk Configuration			
Disk	Status	Size	Operation
sda9	in-use	914.22 GB	

Refresh Disk Information

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](#)" on page 152.

Formula:	6062:
HDD size: M GB, X GiB1	HDD size: 1000GB (928GiB)
Base OS: 14GiB	Base OS: 14GiB
Data Storage: X – 14	Data Storage: 928 – 14 = 914GiB
By default, the data storage is divided up as follows:	
CIFS: 15%	CIFS: 15% = 58GiB
Monitor: 15% or 10GiB, whichever is larger	Monitor: 15% = 129GiB
User DB: 1GiB	User DB: 1GiB
Virt: 50GiB (not available on 2060)	Virt: 50GiB (supports virt)
WM (wan-memory): 55%	WM (wan-memory): 55% = 474GiB
Edge Cache: 15%	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 amount of storage allocated to system services, use the CLI `storage` command. For a complete list of `storage` commands, see [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 Virt, 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
cifs	1GB
edge cache	1GB
monitor	10 GiB or current usage, whichever is larger.
users	500 MiB or current usage, whichever is larger
virt	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 virt and CIFS services:

```
(config) # storage service virt 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
  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:          914.22G
Unallocated: 647.82G
```

There is now 647.82G of storage to be allocated to the edge-cache and monitor services. Increase the monitor space by 73.03G to a total of 200G.

```
(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 574.79G) to increase edge-cache to 706.71G (or 723671.04M).

```
(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

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.

Remove all data from a service's disk storage

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. Click **System > Setup > Storage**.
2. To format a services storage, locate the service and click **Format**.

Optimization

Exinda's Optimization technology provides protocol optimization and data reduction, enabling applications to run faster over the WAN. This technology is provided by the following services:

Exinda Community	Provides appliance auto-discovery and acceleration capability services between all Exinda appliances in the WAN.
TCP Acceleration	Provides layer 4 (TCP) protocol optimization.
WAN Memory	Provides data reduction using de-duplication and compression technology.
SMB Acceleration	Provides layer 7 SMB1 and SMB2 (Windows File Sharing) protocol optimization.
NCP Acceleration	Provides layer 7 NCP (Netware Core Protocol over TCP port 524) protocol optimization.
SSL Acceleration	Provides acceleration for SSL encrypted connections. The SSL Acceleration feature is a separately licensed component, please contact your local Exinda representative if you wish to enable this feature.
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.

The Application Acceleration section of the Exinda appliance System Setup allows you to configure and fine-tune various Acceleration related parameters. The configuration pages include:

- [Services](#): Start/Stop/Disable Application Acceleration services.
- ["Group appliances into a community" on page 238](#): Configure Exinda Community settings.
- [TCP Acceleration](#): Configure and fine-tune TCP Acceleration settings.
- [WAN Memory](#): Configure and fine-tune WAN Memory settings.
- ["Accelerate file transfers" on page 240](#): Configure and fine-tune SMB1 and SMB2 acceleration settings.
- ["Add servers for SSL acceleration" on page 255](#): Configure SSL Acceleration settings.
- [Edge Cache](#): Configure Edge Cache settings.
- [Pre-population](#): Configure SMB1 and Edge Cache pre-population.

Auto Discovery

The Exinda auto-discovery process is used for two purposes:

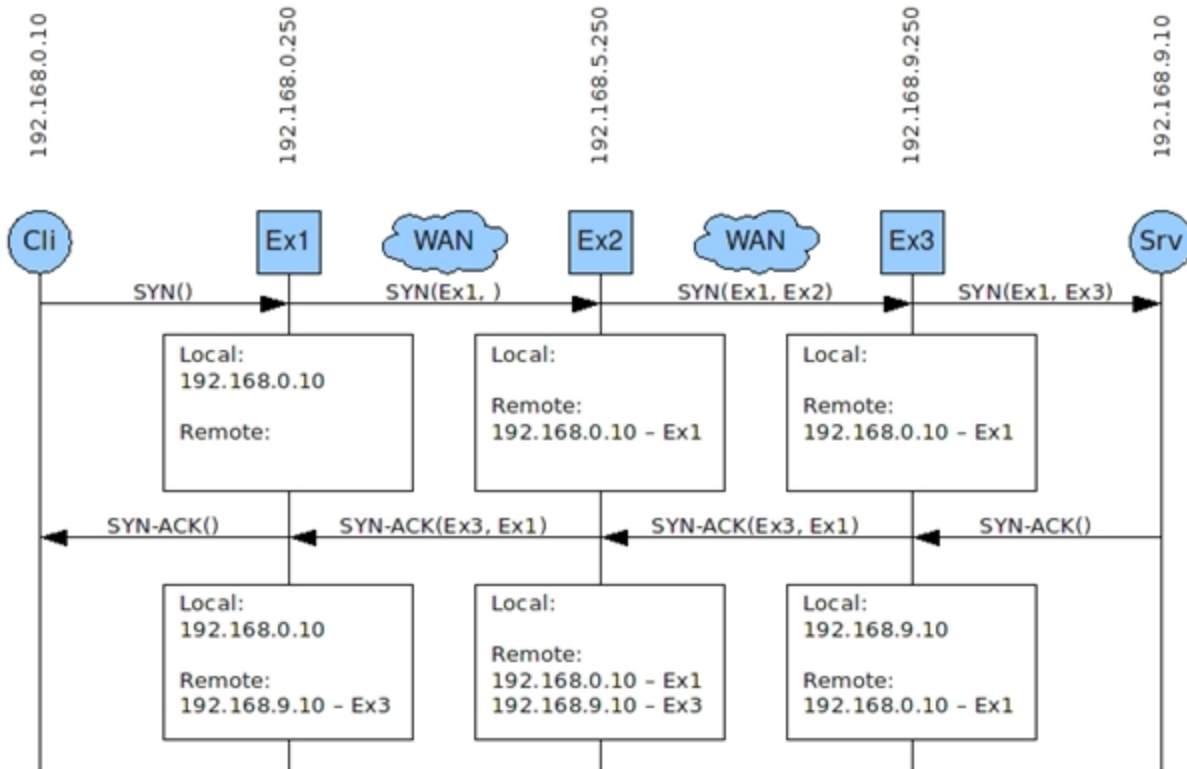
1. The discovery of which connections can be accelerated.
2. The discovery of new Exinda appliances on the network.

To achieve this, some extra information is included in the SYN, SYN-ACK and first ACK packets of each new connection. This information is in the form of a TCP option. The required information is the:

1. Source Appliance ID
2. Destination Appliance ID
3. Acceleration Module Map

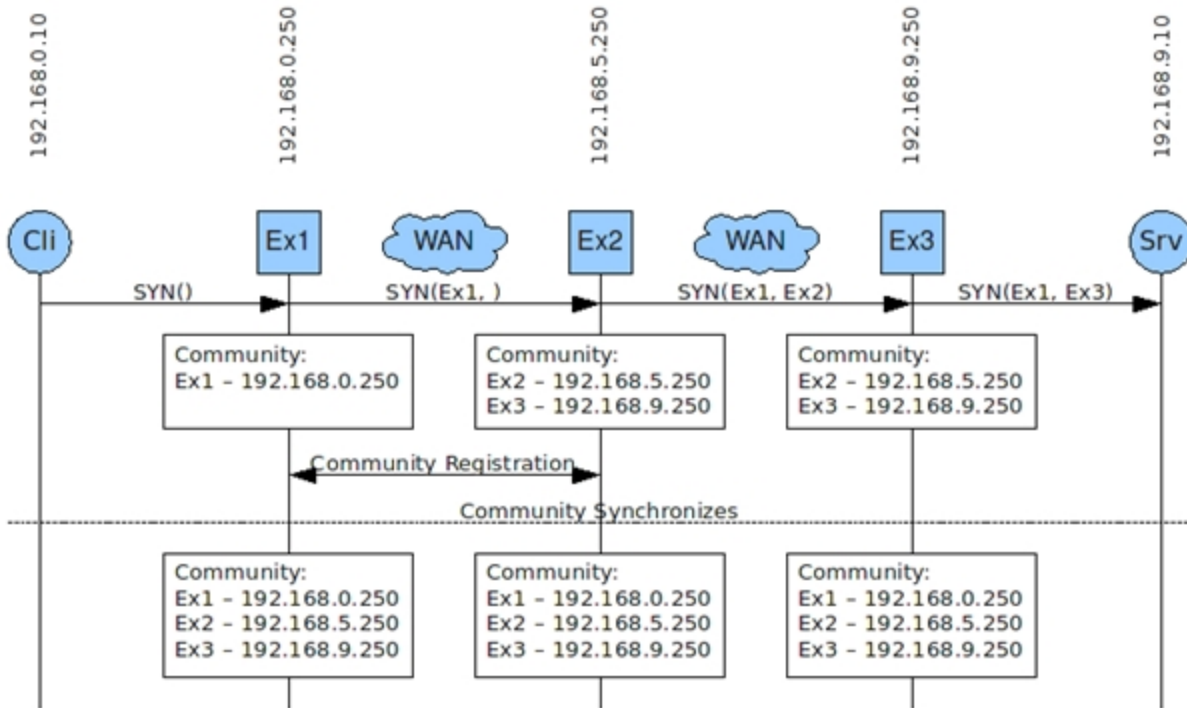
Optionally an IP address corresponding to one of the appliances will also be sent.

In addition to this, each appliance must keep a list of host/appliance id pairs.



The connection discovery process is as follows:

1. When an appliance receives a packet SYN from a client, it adds the client's IP to its local list. It adds the auto-discovery option to the packet, filling out the source details. If the server exists in the appliances remote list, then the destination field is filled out with that appliances details, otherwise the destination is left blank.
2. When an appliance receives a SYN packet containing the auto-discovery option, it will record the client IP address and source appliance id in its remote list. It will then fill out the destination details and forward the packet on.
3. The end appliance can be determined by receiving a SYN-ACK from the server without any auto-discovery option. This appliance will add an auto-discovery option with both the source and destination details filled out.
4. When another appliance receives the SYN-ACK, it will add the server IP address and source appliance id will be added to the remote list. If it finds that the destination does not refer to itself, then it will ignore all packets further packets that are part of that connection.
5. After the SYN-ACK has passed through, both end devices know who they are accelerating between and any appliances have a record of who to accelerate to if they have a connection to either the client or the server.



This process also allows for the discovery of new Exinda appliances on the network. When an appliance receives an auto-discovery option from a source that the Exinda community doesn't know about, it can notify the community which will establish a connection to that appliance, and add it to the community. This may also cause two existing communities to join together.

The Auto Discovery process is very lightweight - it adds negligible latency/delay to packets as they pass through the Exinda appliance.

Manage optimization services

The Optimization Services page allows you to start, stop, and disable the Optimization Services running on the Exinda appliance.

Note To control Optimization Services, navigate to **System > Optimization > Services** on the Web UI, advanced mode.

The table below lists all the Optimization Services running on the Exinda appliance and allows you to Start, Restart (if already running), Stop and Disable them.

Manage Optimization Services			
NCP Acceleration: Running	Restart	Stop	Disable
SMB Acceleration: Running	Restart	Stop	Disable
TCP Acceleration: Running	Restart	Stop	
WAN Memory: Running	Restart	Stop	Disable
Exinda Community: Running	Restart	Stop	
SSL Acceleration Running	Restart	Stop	Disable
Edge Cache: Running	Restart	Stop	

Note Stopping or restarting running Optimization Services may impact connections that are currently accelerated.

Group appliances into a community

A group of Exinda appliances in a user's network is referred to as a community. Exinda appliances that are part of the same community can accelerate to and from each other.

In this area of the Exinda Web UI you can:

- "Add an Exinda appliance to the community" on page 238
- "Change the IP address of an appliance in the community" on page 238
- "Remove an Exinda appliance from the community" on page 239
- "Specify the community groups an Exinda appliance can join" on page 239
- "Remove the community from a community group" on page 240

Add an Exinda appliance to the community

Generally, Exinda appliances automatically discover each other when attempting application acceleration, however, if an appliance is not automatically discovered manually add the Exinda appliance to the community.

1. Click **System > Optimization > Community**.
2. In the Manually Add New Community Node area, type a name for the Exinda appliance.
3. Type the **IP Address** of the Exinda appliance.
4. Click **Apply Changes**.

The appliance is added to the list of manually added community nodes.

Change the IP address of an appliance in the community

When the IP address of an Exinda appliance changes, the community node must be updated as well.

1. Click **System > Optimization > Community**.
2. Click **Edit** for the Exinda appliance.
3. Modify the name or IP address of the Exinda appliance.
4. Click **Apply Changes**.

Remove an Exinda appliance from the community

Remove any manually added appliances from the community as they are removed from the network.

1. Click **System > Optimization > Community**.
2. To remove an Exinda appliance from the community, click **Delete** beside the appliance.
3. To remove the community, click **Remove all community peers from system**.

Specify the community groups an Exinda appliance can join

Community Groups allow you to create multiple, separate Exinda Communities in the same network. You may wish to isolate Application Acceleration between several Exinda appliances in your network. An Exinda appliance can belong to multiple Community Groups. By default, all appliances belong to the community group with Group ID 0.

1. Click **System > Optimization > Community**.
2. Select **Support version (pre v6.4.0) Enabled**.
3. Click **Apply Changes**.

The Community Groups areas are displayed.

4. To identify the community groups that the Exinda appliance can join, type the number in the **Group ID** field.
5. Click **Add New Community Group ID**.

All appliances in the network with the same community group are auto-detected and listed in the Community Peers list. Traffic between the appliances in the community group is accelerated.

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.

Example

In a network, there are two WANs: a VPN-based WAN, and a MPLS-based WAN. I want all the Exinda appliances in the VPN-based WAN to accelerate to each other, and I want all Exinda appliances in the MPLS-based WAN to accelerate to each other; but I don't want the VPN-based Exinda appliances accelerating with the MPLS-based Exinda appliances. The two WANs also both access a Datacenter where there is a single Exinda appliance.

For this scenario, configure the community groups as follows:

Exinda appliances at VPN-based sites: Community Group ID: 10
Exinda appliances at MPLS-based sites: Community Group ID: 20
Exinda appliance at Datacenter Community Group ID: 10 and 20

Remove the community from a community group

If the organization of the community groups changes, or a community group becomes obsolete, remove the group ID from the list of community groups that will auto-detect the Exinda appliance.

Note If all group IDs are removed, the Exinda appliance cannot be added to any community groups.

1. Click **System > Optimization > Community**.

2. Select the community group to remove.

To select all community groups, select the Group ID checkbox.

3. Click **Remove Community Group**.

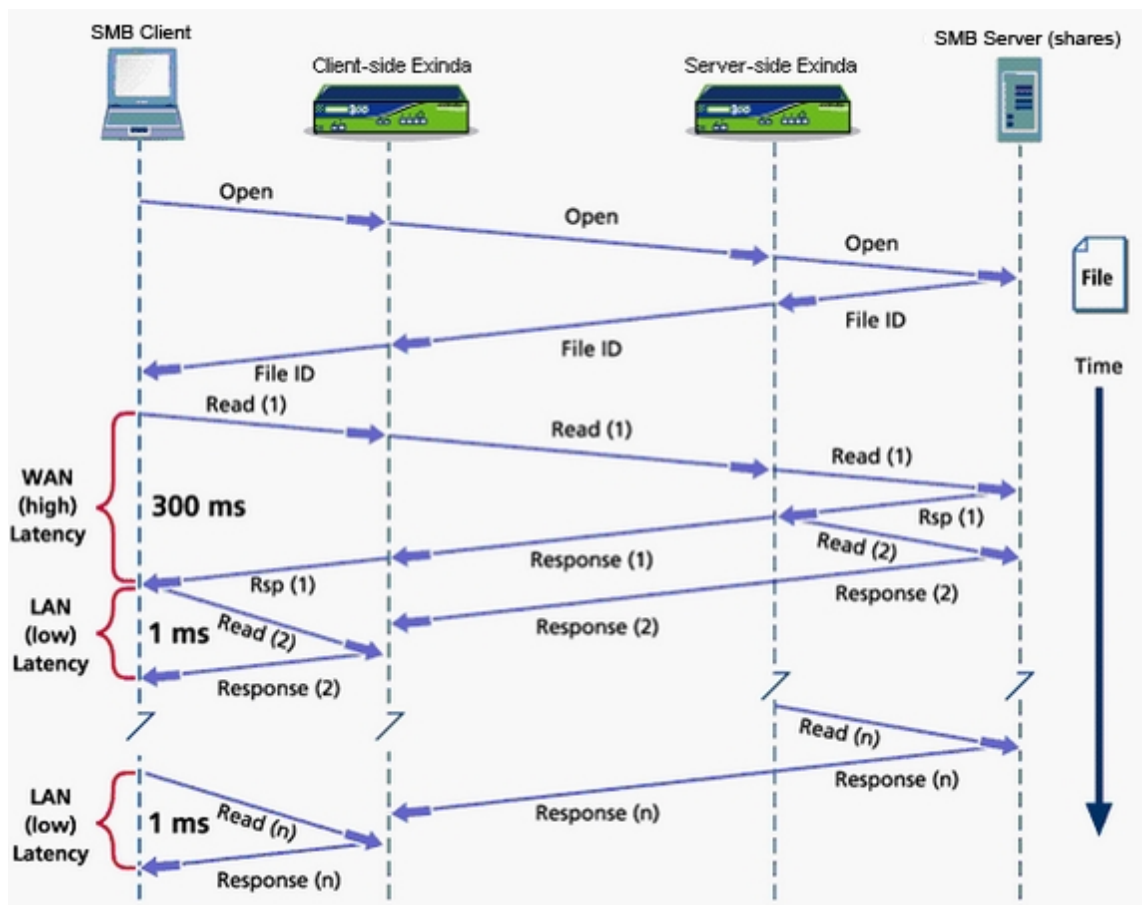
The Exinda appliance is no longer auto-detected by the other appliances in the community group.

Accelerate file transfers

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, the SMB protocols are used to transport information (files or directory information) back and forth between your computer and the server you are accessing.

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.

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 appliances memory for future reference. Once the pre-fetched data is referenced, the data is deleted from the memory.



The above graphic illustrates the primary goal of Exinda SMB acceleration: reduce the overall accumulation latency introduced by the chattiness of the SMB protocol. SMB acceleration works seamlessly with Exinda's TCP Acceleration, WAN Memory, and Compression and benefits from WAN memory's ability to reduce data traversing the WAN just as other applications such as FTP, HTTP or email do.

Configure the Exinda appliance to accelerate file transfers:

1. Create a user account in the domain to sign SMB connections.

An existing account can be used for this purpose. The account used for signing the connections must be able to authenticate against the server, but Exinda recommends that the account should be a highly restricted account that does not have access to the files being accessed by the client computers, or administrator access to the domain.

2. ["Accelerate digitally signed SMB connections" on page 245](#)
 - a. ["Configure file acceleration with SMB1" on page 244](#)
 - b. ["Enable file acceleration with SMB2" on page 245](#)
 - c. ["Allow file transfer acceleration with older versions of Exinda OS" on page 246](#)
3. Restart the SMB Acceleration service. See ["Manage optimization services" on page 237](#).

Restarting the service ensures that any cached information about hosts is cleared. This will also drop any previously established optimization connections to servers allowing any subsequent connections to the signed server to attempt to sign the connections on the next attempt.

4. View the chart of SMB connections. See "[Report: Accelerated Connections](#)" on page 65.

SMB optimizations

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 de-duplicated.

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 saving the contents 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 respond 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 De-duplication

Aside from the protocol specific optimizations that are provided by the appliance, Exinda's SMB acceleration framework also provides some significant downstream optimization benefits. Primarily in the areas of compression and de-duplication. 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, we are passing off these large blocks of data to our WAN memory framework. This allows the WAN memory framework to heavily optimize for compression and de-duplication.

Common SMB Use Cases

A file download was used to illustrate how Exinda performs SMB acceleration. However, SMB acceleration uses similar mechanisms to achieve greatly improved performance for many other scenarios. Below are a few examples:

File Upload (Write)

This is conceptually very similar to a file download with the obvious difference being that a SMB client is writing a file to a SMB server instead of reading it. In this case, the client side Exinda responds locally to the SMB client's write requests 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 (e.g. MS Word, PowerPoint, Excel, etc.) which reside on a remote SMB server are often opened from a SMB client. This action suffers from all of the SMB related problems that are described

in this paper because the file data is retrieved serially, 61k bytes at a time. The result is a long wait time to open the file, browse or perform any type of action (e.g. save). Exinda's SMB Acceleration addresses these problems 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.

Configure file acceleration with SMB1

To reduce the severe effect high network latency has on the SMB protocol, configure the Exinda appliance to pre-fetch and caches data for SMB1 requests.

1. Click **System > Optimization** and switch to the **SMB** tab.
2. In the SMB Acceleration Options area, select **Enabled**.
3. To prefetch data from the SMB server in anticipation of subsequent client requests, select **Read Ahead**.
4. To update the SMB cache, aggregating requests to the SMB server, select **Write Behind**.
5. To enable caching of SMB file attributes such as file access times and size, select **Meta-Data Caching**.
6. Specify the amount of SMB data to prefetch when performing read-ahead or write-behind in the **Data to Prefetch** field.
This value should only be increased if network latency is very large, for example latency is greater than 500 milliseconds. The default value is 1024 kb.
7. Click **Apply Changes**.
8. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Change the amount of data prefetched for the SMB1 cache

If you are experiencing network latency with file transfers, for example latency is greater than 500 milliseconds, increase the amount of data being prefetched for the SMB1 cache.

1. Click **System > Optimization** and switch to the **SMB** tab.
2. Modify the amount of SMB data to prefetch when performing read-ahead or write-behind in the **Data to Prefetch** field.
3. Click **Apply Changes**.

4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Enable file acceleration with SMB2

To reduce the severe effect high network latency has on the SMB protocol, enable the Exinda appliance to cache data for SMB2 requests.

1. Click **System > Optimization** and switch to the **SMB** tab.
2. In the SMB2 Acceleration Options area, select **Enabled**.
3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Accelerate digitally signed SMB connections

Ensure the security of SMB sessions by digitally signing SMB connections. With signing, the server can verify that each message was sent by the client that initially requested the connection, and the client can verify that the response it is receiving is from the server.

1. Click **System > Optimization** and switch to the **SMB** tab.
2. To accelerate digitally signed connections, in the **SMB Signing** area select **Enabled**.
3. In the SMB Signing Options area, type the NetBIOS name of a domain or computer where shared files are located. Click **Add**. Do not use the fully qualified domain name.
The domain is added to the SMB Signing Credentials area.
4. In the SMB Signing Credentials area, type the **Username** and **Password** to use when generating the signing key.

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.

5. Select whether the authentication credentials are enabled or disabled.

The Exinda appliance will use the recorded credentials for the requested domain. If it is unable to connect to that domain because of the server is unavailable or incorrect credentials, the status of the signed connection is reported as *Bypassed* or *Unhandled*. See "[View the SMB configuration and connections](#)" on page 334.

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*.

6. Click **Apply Changes**.
7. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Allow file transfer acceleration with older versions of Exinda OS

If your network contains Exinda appliances running ExOS versions older than v6.3.0, for example 6.1, as well as appliances running ExOS version 6.3.0 or later, you must enable version compatibility to accelerate file transfer traffic successfully.

1. Click **System > Optimization** and switch to the **SMB** tab.
2. In the Compatibility Options area, select **Support versions (pre v6.3.0) Enabled**.
3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status *Unsaved changes (Save)*

Verify acceleration configuration

Acceleration diagnostics aid in troubleshooting TCP Acceleration, SMB Acceleration and WAN Memory issues by displaying the current configuration for those areas.

In this area of the Exinda Web UI you can:

- ["View the SMB configuration and connections" on page 334](#)
- ["View the TCP acceleration configuration and connections" on page 335](#)
- ["View the WAN memory configuration and reduction statistics" on page 336](#)

View the SMB configuration and connections

The SMB Acceleration diagnostics display the current configuration settings as well as the number of new and concurrent accelerated connections. If SMB signed connections are present, the total number of signed connections is also displayed.

Note To configure CIFS acceleration settings, navigate to **System > Optimization > SMB** on the Web UI, advanced mode.

1. Click **System > Diagnostics** and switch to the **Acceleration** tab.
2. From the Module drop-down, select **SMB Acceleration**.

The configuration settings for SMB and SMB2 are displayed. The connections statistics are broken down 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 first time an attempt to validate the domain credentials fails, the connection is identified as being signed, but is not accelerated. This attempt and 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 were accelerated.
- **Unhandled** — After a signed connection has failed validating the domain credentials the first time, and the connection is marked as **Bypassed**, all subsequent attempts to validate credentials of a signed connection against the IP address of the server are marked as **Unhandled**.

Note The statistics reported on this page are reset each time the SMB Acceleration service is restarted.


3. ["View System Log Files"](#) on page 327.

Report: Accelerated Connections

This report shows the number of concurrent accelerated connections as well as the accelerated connection establishment rate through the Exinda appliance over time.

1. Click **Monitor > System** and switch to the **Accelerated Connections** tab.


The reports are displayed.



The second graph shows the number of connections through each of the specific application acceleration modules, such as SMB or SMB2 Acceleration, SSL Acceleration, and NCP Acceleration.

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.



By transparently proxying TCP connections like this, Exinda has full control of the TCP connection over the WAN. It is this WAN connection that TCP Acceleration technologies are applied to.

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.

- **Window 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 high speed high latency links, something like High Speed TCP should be used. For Satellite links, or other high-latency links, something like Hybla should be used. This allows for better TCP performance over different WAN technologies.
- **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.
- **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 when TCP uses a combination of the slow start and congestion avoidance algorithms to reset the send window size temporarily, to avoid congestion.

Modes of Operation

The Exinda appliance's default mode of acceleration is called Transparent Mode. In this mode the source and destination addresses and ports of the client and server are maintained to provide the least intrusive network

implementation.

A secondary supported mode of acceleration is called Protocol Mode . In Protocol Mode the Exinda solution does not use traditional TCP connections to send traffic and dynamically sends the TCP connections over a different protocol. Once a connection is setup via the protocol, subsequent connections can use the protocol connection and avoid the 3-way TCP handshake. The benefit of Protocol Mode is that it reduces the number of TCP connections traversing the WAN and also reduces the TCP connection setup time. Protocol Mode is also used to bypass connectivity issues that can be caused by firewalls and intrusion detection systems.

Configure TCP Acceleration

TCP Acceleration is the heart of Exinda's Application Acceleration technology. All accelerated connections will be passed through TCP Acceleration.

Note To configure the TCP Acceleration settings, navigate to **System > Optimization > TCP** on the Web UI, advanced mode.

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	auto
Transport Type	Transparent TCP
Window Scaling Factor	5 / 2M
Congestion Control	<input checked="" type="radio"/> [cubic] General Purpose <input type="radio"/> [hybla] Satellite (High speed, high round-trip-time) <input type="radio"/> [highspeed] High speed <input type="radio"/> [veno] Wireless (Loss handling) <input type="radio"/> Other: reno
TCP Keep-Alive Enabled	<input checked="" type="checkbox"/>
TCP Keep-Alive Timeout	3600 seconds

Apply Changes

Auto Discovery	<p>If enabled, the Exinda appliance will attempt to automatically discover other Exinda appliances on the network when making acceleration attempts. 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.</p> <p>If this setting is 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</p>
----------------	--

	network on the System > Acceleration > Community page.
Auto Discovery IP Address	This is the IP address the Exinda appliance will use when notifying other Exinda appliances about how to connect back to itself. Usually this is the management IP address or the IP address to which the default route is associated.
Transport Type	<p>There are two acceleration Transport Types available. The default, 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 2 accelerating Exinda appliances can still see correct IP and TCP headers.</p> <p>The other option, Protocol 139, still preserves the original IP header of the accelerated connection, but tunnels the connection over IP protocol 139, so it no longer appear as TCP. This mode is useful if you have equipment in between 2 accelerating Exinda appliances the strips or blocks TCP option 30.</p>
Window Scaling Factor	<p>The Window Scaling Factor determines how large the TCP window is allow to grow per connection. The TCP window size is calculated using the following formula: TCP Window Size (kbytes) = 64 kbytes x 2 ^ Window Scaling Factor. Both the Window Scaling Factor and the TCP Window size are displayed in the drop-down list.</p> <p>The default Window Scaling Factor is 5, which equates to a TCP window of 2Mbytes. Larger window sizes result in more potential memory usage, however, this value may need to be increased in high-bandwidth, high-latency environments.</p>
Congestion Control	There are various types of congestion control algorithms that can be used depending on the type of WAN. 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. The timeout specifies when to activate the keep-alives if enabled.
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 five minutes passes. If five minutes pass without a response the connection is terminated.

Note The Exinda appliance uses TCP Option 30 to facilitate transparent TCP Acceleration and Appliance Auto Discovery. Any equipment between the Exinda appliances must not block or strip TCP option 30 otherwise transparent TCP Acceleration and Appliance Auto Discovery will not work.

Caution These settings are exposed for advanced users and should only be changed after consultation with an Exinda representative.

There is one additional TCP Acceleration setting that is only available via the CLI. This is an advanced setting that controls how TCP Acceleration behaves when more than one bridge is accelerating traffic and it is enabled by default.

```
[no] acceleration tcp dual-bridge-bypass
```

The Dual Bridge Bypass option should be enabled when an accelerated connection needs to pass through multiple bridges on the same appliance. For example, if the accelerated traffic arrives on br0, gets decelerated and then routed back through br1 of the same appliance, this option needs to be enabled.

The Dual Bridge Bypass option should be disabled if accelerated packets from a single connection could arrive on more than one bridge. This behaviour is typically exhibited when packet-based load balancing is used.

For more information on TCP Acceleration, refer to [Appendix A](#).

WAN Memory

WAN Memory is the data de-duplication 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 2 or more Exinda appliances.

Note To configure the WAN Memory settings, navigate to **System | Optimization | WAN Memory** on the Web UI, advanced mode.

Use the form below to configure WAN Memory settings:

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

LZ Compression	If selected, in addition to data de-duplication, WAN Memory will also attempt to compress accelerated traffic with a standard LZ-based compression algorithm.
Persistent Cache	If selected, the WAN Memory patterns stored on the Exinda appliance's hard disk will survive a system reboot.
HA cache	When in cluster mode, WAN Memory caches are mirrored to the WAN Memory caches on

sync	the other appliances in the cluster.
Force Data Expiration	Use this button to expire the entire WAN Memory cache, thereby removing any patterns stored on the Exinda appliance's hard disk drive. This may take several minutes depending on the amount of data.
Reset Persistent Data	Use this button to tell WAN Memory NOT to load any persistent data from the hard disk next time it starts. Using this function and restarting the WAN Memory service is a quick way to clear the WAN Memory cache.

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.


Add SSL certificates and keys on the Exinda appliances

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's desktop. The SSL is used protocol to secure each session link by automatically establishing connections on-demand using standards-based protocols, encryption techniques, and certificate exchange.

SSL encryption requires a server certificate to authenticate the identity of a server. A certificate is an electronic confirmation that the owner of a public key is who he or she really claims to be, and holds the private key corresponding to the public key in the certificate.

A public key and a particular server can communicate securely by generating a certificate signing request using the server's public key. The identity of the server is verified by a Certificate Authority (CA), and generates a signed certificate. The server offers the generated certificate to clients that recognize the CA's signature, and verifies the identity of the server.

The process of setting up the private connection is automatically initiated by the server communicating directly with the browser. The result is a private, encrypted tunnel used to move information between the server and client desktop. When the session is over, the connection is automatically terminated.



If the connection between the client and the server uses SSL to encrypt the sessions, the benefits that can be gained by application acceleration are limited. For example, Exinda's WAN Memory technology achieves higher reduction on clear text rather than encrypted data.

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.

Note The SSL Acceleration settings on each Exinda appliance be configured exactly the same.

SSL certificates and keys can be added to the Exinda appliance in one of the following ways:

- "Import a certificate onto an appliance" on page 253 that has been [exported from another appliance](#).
- "Generate a self-signed certificate" on page 254

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.

1. Click **System > Certificates > All Certificates**.

Export an SSL certificate

If an SSL certificate is only available on once appliance, export the certificate so it can be imported onto the other Exinda appliances on the network.

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.

1. Click **System > Certificates**.
2. On the Certificates, CA Certificates, or All Certificates tab, locate the certificate in the list, and click **Export**.
3. Select the format for the exported certificate.
4. Click **Save**.

The certificate is downloaded onto the computer accessing the Exinda Web UI.

Import a certificate onto an appliance

There are two types of certificates that can be imported onto the Exinda appliances:

- A Certificate Authority (CA) generated certificate and private key from the server that is hosting the accelerated SSL application.

Advantage: The end user will see the real certificate when they access the accelerated SSL application.

Disadvantage: The private key of the accelerated SSL application needs to be loaded on ALL Exinda appliances.

- Load different or self-signed certificate and private key.

Advantage: Private keys never leave the server, a different set of private keys are installed on the Exinda appliances.

Disadvantage: End users will receive a warning notifying them that the certificate is invalid unless the common name (CN) in the certificate loaded on the Exinda appliances matches the common name of the accelerated SSL application AND the end-user indicates they trust the certificate.

Regardless of the type of certificate used, they are imported in the same way.

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.

1. Click **System > Certificates**.
2. On the **Certificates** tab, select **Import Certificate**.
3. (Optional) In the **Import Certificate and Key Details** area, 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/private 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 Password** field type the password.
6. In the **Certificate File** field click **Choose File** and navigate to the file to be uploaded to the Exinda appliance.
7. If the PEM format is selected, the private key must be uploaded to the Exinda appliance. In the **Private Key File** field, click **Choose File** and navigate to the private key file.
8. Click **Import**.

The certificates are displayed in a table at the top of the page where the contents can be viewed, or the certificate can be deleted.

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.

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

1. Click **System > Certificates**.
2. On the **Certificates** tab, select **Generate Certificate**.

3. In the Generate Certificate and key Details area, type a name for the certificate.
4. In the **Key Size** field, specify the number of bits to use when encrypting the contents of the certificate.
5. Specify how many days the certificate is valid for.
6. Type the name of the organization and the name of the area that will be using this certificate.
7. In the **Common Name** field, type the name of the person issuing the certificate.
8. Click **Generate**.

After the certificate has been created, it appears in the list of certificates on the Certificates tab.

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.

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.

1. Click **System > Certificates**.
2. On the Certificates, CA Certificates, or All Certificates tab, locate the certificate in the list, and click **Show**.
3. To return to the list of certificates, click **Back**.

Delete an SSL certificate

Delete an SSL certificate from the Exinda appliance when it expires, or becomes invalid.

1. Click **System > Certificates**.
2. On the Certificates, CA Certificates, or All Certificates tab, locate the certificate in the list, and click **Delete**.
3. In the confirmation dialog, click **OK**.

The certificate is deleted.

Add servers for SSL acceleration

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 servers that are explicitly configured are SSL accelerated. Any SSL traffic that the Exinda appliance sees that does not belong to a configured server is ignored.

1. Click **System > Optimization > SSL**.
2. In the **Add SSL Acceleration Server** area, type a name for the server or application you wish to enable for SSL Acceleration.

3. Type the IPv4 address of the server running the SSL enabled application.
4. Type the port number running the SSL enabled application on the server.
5. If the server has multiple SSL certificates with a Server Name Indication (**SNI**) specified, type the SNI extension in the field.

Caution 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. The server without the SNI extension is used as a fallback in case 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.

6. 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.

Note If there are any problems with the certificate or key associated with a configured SSL server (E.g. missing key, expired certificate), then SSL Acceleration will ignore that traffic until the issue is resolved.

7. Select the **Client Auth Certificate** to authenticate sessions on the SSL server.
8. 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 if the server's SSL certificate is invalid or expired. The connection is still accelerated, but not SSL accelerated.
 - **Certificate**—SSL Acceleration accepts and process the connection only if the server's certificate matches the validation certificate. Otherwise, the connection is not processed.
9. If **Certificate** is selected as the Validation type, select the certificate to validate against.
10. Click **Add SSL Server**.

The servers are displayed at the top of the page, where they can be edited or deleted.

Edit the settings of an SSL Acceleration server

Update the settings of the SSL acceleration server when new certificates are available.

1. Click **System > Optimization > SSL**.
2. Locate the server in the list, and click **Edit**.
3. Modify the settings for the server, and click **Apply Changes**.

The settings for the server are changed.

Delete an SSL Acceleration server

When a server is no longer active, or if a server no longer requires SSL Acceleration, remove it from the list of optimized servers.

Note A server cannot be deleted if another server with the same IP address and port number and an Server Name Indication (SNI) extension has been configured on the Exinda appliance. Servers with SNI extensions must be deleted before the server can be deleted.

1. Click **System > Optimization > SSL**.
2. Locate the server in the list, and click **Delete**.
3. In the confirmation dialog, click **OK**.

The server is deleted.

Create policies to accelerate SSL traffic

By default, SSL traffic is captured by a QoS only policy, meaning no attempt is made to accelerate any SSL traffic by default. 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.

1. Click **Optimizer > Policies**.
2. Create the policy for accelerating an SSL application.
3. Click **Create New Policy**.
4. Add the policy on all Exinda appliance.
5. Once the desired policies are in place on all Exinda appliances, restart the Optimizer. In the appliance status bar, click **Restart**.

Optimizer Status : On (Restart / Stop)

Example

The following is an example of a policy that accelerates an SSL application. This Policy is placed above any other Policy that captures SSL traffic in the policy tree.

The screenshot shows the 'Edit Policy' configuration window. The policy name is 'SSL Accel'. The schedule is set to 'ALWAYS'. The action is 'Optimize'. The policy is enabled. The guaranteed bandwidth is 5%, burst (max) bandwidth is 100%, and burst priority is 4. The acceleration is enabled, and the WM reduction type is 'Disk'. The ToS/DSCP mark is empty. The filter rules table is as follows:

VLAN	Host	Direction	Host	ToS/DSCP	Application
ALL	ALL	< - >	SSL Server	ALL	HTTPS
		< - >			
		< - >			
		< - >			
		< - >			

Ciphers supported in SSL acceleration

SSL Acceleration supports the following ciphers.

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.

Configure the websites served up from Windows Server 2012 to include the SNI extension in the connection requests.

Note Only Windows Server 2012 has support for Server Name Indication (SNI).
SNI is not supported on Internet Explorer running on Windows XP.

1. ["Install IIS 8.0 on Windows Server 2012" on page 259](#)
2. ["Add sites to the web server" on page 259.](#)
3. Ensure the certificates required for the sites are available on the server.

Depending on how your company 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](#)" on page 260.
5. "[Identify the certificate to be used by each website](#)" on page 260.
6. "[Export SSL certificates from Windows Server 2012](#)" on page 261.
7. "[Import a certificate onto an appliance](#)" on page 253.
8. "[Add servers for SSL acceleration](#)" on page 255.

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:

```
selfssl.exe /N:CN=TEST.SITE.3 /K:1024 /V:<days-valid> /S:<site-ID> /P:443
```

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.

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 ["Identify the certificate to be used by each website" on page 260](#).

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" on page 259](#).
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" on page 261](#)
4. ["Specify the ports referenced by the virtual hosts" on page 262](#)
5. ["Add a <VirtualHost> block for each secure site on the server" on page 262](#)
6. ["Verify the secure server configuration" on page 263](#)
7. ["Import a certificate onto an appliance" on page 253](#).
8. ["Add servers for SSL acceleration" on page 255](#).

Enable SSL on Apache

To use SSL on Apache, the `mod_ssl` module must be enabled.

1. 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.

Accelerate web applications through caching

Note The Edge Cache Acceleration feature is a separately licensed component. See "[Check the Features Included in your License](#)" on page 139. Please contact your local Exinda representative if you wish to enable this feature.

Edge Cache provides acceleration of HTTP based applications by caching objects in memory and on disk. Edge Cache acceleration is single-sided, and acceleration requires only one Exinda appliance.

There are two important requirements for Edge Cache to function correctly:

- A gateway – this must be a gateway (firewall or internet router) on the WAN side of the Exinda appliance.
- A valid DNS server.

Edge Cache is a fully transparent solution, however these requirements must be met in order for Edge Cache to function correctly.

To accelerate applications using Edge Cache, perform the following tasks:

1. "[Create the Edge Cache Policy](#)" on page 264
2. "[Add the Edge Cache Policy to a Virtual Circuit](#)" on page 264
3. "[Configure the Edge Cache Default Settings](#)" on page 264
4. "[Exclude URLs from the Edge Cache](#)" on page 265
 - a. "[Remove URLs from the Edge Cache Exclusion List](#)" on page 265

5. ["Add an Edge Cache Peer" on page 265](#)
 - a. ["Edit an Edge Cache Peer" on page 265](#)
 - b. ["Delete an Edge Cache Peer" on page 265](#)
6. ["Remove All Objects from the Edge Cache" on page 266](#)
7. ["View Edge Cache Statistics" on page 55](#)

Create the Edge Cache Policy

Edge Cache works on outbound, HTTP based conversations. To enable Edge Cache, create a policy with an application or application group that will capture the HTTP application traffic that you wish to cache.

1. Click **Optimizer > Policies > Create New Policy**.
2. Type a name for the policy.
3. Set the required bandwidth parameters.
4. Select the **Acceleration** checkbox and select **Edge Cache** from the list.
5. Create the rules for the policy, ensuring that HTTP or an HTTP based custom application is selected from the application list.
6. Click **Add New Policy**.

Add the Edge Cache Policy to a Virtual Circuit

Add the Edge Cache policy to an outbound Virtual Circuit.

1. Click **Optimizer > Optimizer**.
2. In the appropriate Virtual Circuit area, select the Edge Cache policy from the list.
3. Type the order number for the policy.

As the traffic is passed through the Exinda appliance, it is first matched to a Virtual Circuit in order of Virtual Circuit number, and then by a policy within that Virtual Circuit in order of Policy number.

4. Click **Add to Virtual Circuit**.

Configure the Edge Cache Default Settings

Edge Cache only stores objects within configured limits. Use the form below to change the default minimum and maximum object sizes.

1. Click **System > Optimization > Edge Cache**.
2. In the Memory Object Options area, type the minimum and maximum size of the objects to be cached.
3. In the **Connection Timeout** field, type the maximum time 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

4. Click **Apply Changes**.

Exclude URLs from the Edge Cache

Create a list of URLs or domains that will never be cached. URL's matching the exclusion list still pass through Edge Cache, and are highlighted in blue on the Real Time conversations list.

Use the Clear button to remove all objects from the cache.

1. Click **System > Optimization > Edge Cache**.
2. In the **Add URL/Domain** area, type the URL or domain that will be excluded from the Edge Cache.
3. Click **Add URL**.

Remove URLs from the Edge Cache Exclusion List

Create a list of URLs or domains that will never be cached. URL's matching the exclusion list still pass through Edge Cache, and are highlighted in blue on the Real Time conversations list.

1. Click **System > Optimization > Edge Cache**.
2. In the **Add URL/Domain** area, locate the URL or domain to be deleted, and click **Delete**.

Add an Edge Cache Peer

Exinda appliances with Edge Cache can establish a community of peers in order to share objects between caches.

1. Click **System > Optimization > Edge Cache**.
2. Click **Add New Peer**.
3. Type the **Host Name** of the peer.
4. Type the **HTTP Port** of the new peer.
5. Type the **ICP Port** of the new peer.
6. Click **Add New Peer**.

The peer appears in the list of configured peers.

Edit an Edge Cache Peer

Update the settings of the Edge Cache peers when server settings change.

1. Click **System > Optimization > Edge Cache**.
2. Locate the peer to be modified in the list, and click **Edit**.
3. Change the peer settings.
4. Click **Apply Changes**.

Delete an Edge Cache Peer

Remove peers that are no longer active, or peers that no longer need to share objects between the caches.

1. Click **System > Optimization > Edge Cache**.
2. Locate the peer to be modified in the list, and click **Delete**.

The peer is deleted from the appliance.

Remove All Objects from the Edge Cache


Remove peers that are no longer active, or peers that no longer need to share objects between the caches.

1. Click **System > Optimization > Edge Cache**.
2. To remove all objects from the Edge Cache, click **Clear**.

View Edge Cache Statistics


The Edge Cache report shows the amount of data reduced by the Edge Cache.

1. Click **Monitor > Optimization > Edge Cache**.
2. Select the type of report to view.
 - **Edge Cache Throughput**—the amount of WAN and LAN traffic that was sent from the Edge Cache instead of from the HTTP application.




- **Edge Cache Reduction Ratio (%)**—the percentage of requested data that was sent from the

Edge Cache instead of from the HTTP application.



3. To display the number of requests per second and the number of hits per second, review the **Edge Cache Statistics** graph. A hit occurs when a request is made for an object stored in the Edge Cache.



4. Review the table to see 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: #6aa84f; height: 10px;"></div> 9.71	204030	120030	<div style="width: 58.83%; background-color: #6aa84f; height: 10px;"></div> 58.83

Pre-populate the cache

The Pre-Population feature allows you to warm the Exinda's SMB, Edge, or WAN memory cache 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.

1. Create a pre-population job
 - a. ["Create a pre-population job with the Exinda Web UI" on page 268.](#)
 - b. ["Create a pre-population job with the CLI" on page 270.](#)
2. ["Create a Scheduled Job" on page 271](#)

Prerequisites for Pre-population

Before configuring pre-population, you must have the following:

- Exinda 6.3.x or later
- One of the following licenses:
 - X800 License
 - X800 License with Edge Cache
 - X700 License with Edge Cache

Note In an X800 environment you need an active Community and Wan Memory for SMB pre-population to work.

Create a pre-population job with 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. Click **System > Optimization** and click the **Pre Population** tab.
2. In the Add New Pre-Population area, type a name for the pre-population job.
3. Select the type of cache to pre-populate.
4. Type the hostname or IP address of the server containing the cache.
5. Type the path to the file or cache directory on the server.
6. To recursively fetch all the files in the specified directory, as well as those in sub-directories, select **Yes**.
7. Select the appropriate schedule job for how frequently the cache should be updated.
If no jobs are available, [create a job schedule](#).
8. If required, type the username and password for the server.

9. Click **Add New**.

The job appears in the list of configured Pre-population jobs.

10. 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.

11. 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.

Add New Pre-Population	
Name	Documentation
Protocol	<input checked="" type="radio"/> SMB <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.

Add New Pre-Population	
Name	Intranet Webserver
Protocol	<input type="radio"/> SMB <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 with the CLI

Add and configure pre-population jobs using CLI commands.

1. Click **System > Tools** and click **Console**.

2. Log into the console.

The prompt `exinda-appliance >` is displayed.

3. To enter configuration mode, type `enable` then `configure terminal`.

The prompt `exinda-appliance (config) #` is displayed.

4. 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.

5. To view the pre-population rules and the status of those rules, run the following command:

```
show acceleration prepopulate
```

Examples

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
```

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.

1. Click **System > Setup** and click the **Scheduled Jobs** tab.
2. In the Add New Job area, type a unique **ID** for the job.
3. Type a name for the job.
4. [Optional] In the **Comment** field, type a description for the job.
5. To run the job immediately, **Enable** the job.
6. If the job should be completed, even if one or more commands fail to execute, set **Fail-Continue** to **Yes**.
7. Set the frequency of the scheduled job. Jobs can be set to run Once, Daily, Weekly, Monthly, or Periodically.
8. After selecting the frequency 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.
9. In the **Commands** field, type the necessary commands for the job you want run. Each command must be on a new line.
10. Click **Add Job**.

The job is added to the list, and is now available for selection in the Pre-population

Schedule a Pre-Population Job Through 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. Click **System > Tools** and click **Console**.
2. Log into the console.
The prompt `exinda-appliance >` is displayed.
3. To enter configuration mode, type `enable` then `configure terminal`.
The prompt `exinda-appliance (config) #` is displayed.
4. To schedule the pre-population job type the following commands:

```
exinda-appliance (config) # job <job_number> command <command_number>
"acceleration prepopulate <prepopulation_name> start"
```

For example:

```
exinda-appliance (config) # job 100 command 1 "acceleration populate
Documentation start"
exinda-appliance (config) # job 100 schedule daily time 00:30:00
```

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.

- ["Enable MAPI Acceleration on the Exinda Appliances" on page 272](#)
- 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" on page 273](#) and ["Disable encryption on the Exchange server" on page 275](#).
- Verify MAPI traffic is being accelerated.
 - ["View MAPI Acceleration Results" on page 275](#)
 - ["View real-time inbound and outbound conversations" on page 81](#).

Note If you encounter any issues, see ["Troubleshoot problems with MAPI acceleration" on page 279](#)

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. Click **Optimizer**.
2. 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**.
3. 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--
Order:		Policy: ALL - Accelerate	Add To 'WAN outbound'	

4. 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.

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 "[Disable encryption on the Exchange server](#)" on page 275.

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.

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 ["Turn off MAPI encryption in Microsoft Outlook"](#) on page 273

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

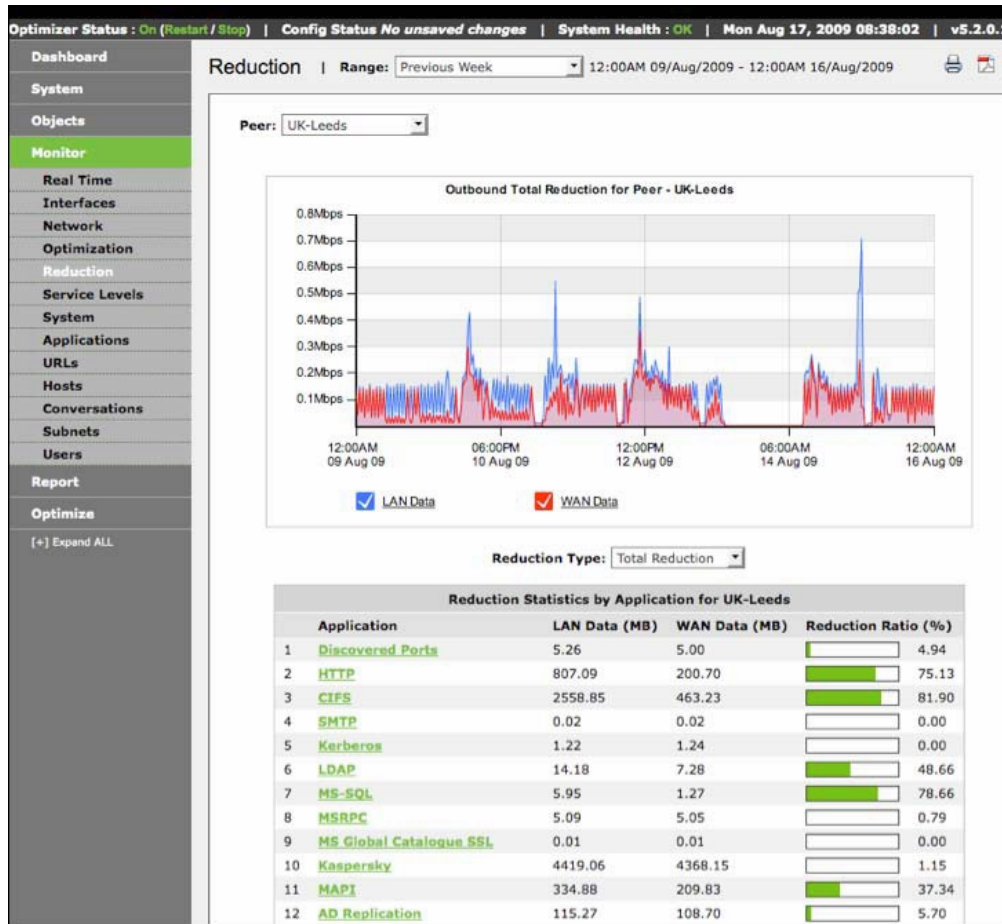
Confirm that the MAPI traffic from Microsoft Outlook and Exchange are being accelerated.

- ["View MAPI Acceleration Results"](#) on page 275
- ["View real-time inbound and outbound conversations"](#) on page 81

View MAPI Acceleration Results

View the reduction in MAPI traffic on the network.

1. Click **Monitor > Optimization** and switch to the **Reduction** tab.
The report displays the reduction in MAPI traffic.









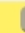



View real-time inbound and outbound conversations










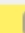
The Real-time Conversations Report shows a breakdown of the Conversations monitored by the Exinda appliance during the last 10 seconds. Conversations are divided into Inbound and Outbound directions.

1. Click **Monitor > Real Time > Conversations**.

By default, the Real-time Conversations Report looks like the example below. Conversations are sorted by throughput. You can also see the packet rate and number of flows for each Conversation. Any extra information about a Conversation (a URL for example) will be shown in square brackets next to the Application.

Inbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			1408.428	284	24	
 	192.168.10.1	192.168.10.128	MAPI	570.834	82	1
 	192.168.10.9	192.168.10.128	MAPI	483.247	54	2
 	192.168.10.7	192.168.10.128	MAPI	275.334	92	2
 	192.168.10.10	192.168.10.128	MAPI	65.153	51	2
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	5.496	1	1
	192.168.10.9	192.168.10.128	LDAP	2.939	1	1
	10.20.4.1	239.255.255.250	udp ports 62612 -> 3702	1.097	0	1
	10.20.4.1	239.255.255.250	udp ports 62610 -> 3702	1.069	0	1
	192.168.10.1	192.168.0.1	NetBIOS	0.623	1	1
	192.168.10.10	192.168.10.128	LDAP	0.556	0	2
	192.168.10.132	255.255.255.255	DHCP	0.541	0	1
	192.168.10.9	192.168.0.1	NetBIOS	0.225	0	1
	10.20.3.118	10.20.255.255	NetBIOS	0.225	0	1
	192.168.10.9	192.168.255.255	NetBIOS	0.225	0	1
	10.20.11.100	224.0.0.252	udp ports 58633 -> 5355	0.212	0	1
	10.20.0.14	10.20.255.255	NetBIOS	0.193	0	1
	192.168.10.9	192.168.10.128	LDAP	0.174	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.106	0	1
	192.168.10.9	192.168.0.1	DNS	0.102	0	1
	10.20.0.181	10.20.255.255	NetBIOS	0.075	0	1

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To view only a specific IP address or subnet, type the address in the **IP/Subnet Filter** field.
The report can be filtered by IPv4 or IPv6 addresses.
- To display the optimization policy the conversation falls into, select **Show Policies**.


Outbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			3594.527	412	14	
 	192.168.10.7	192.168.10.128	MAPI	1826.409	196	2
 	192.168.10.10	192.168.10.128	MAPI	1184.445	125	2
 	192.168.10.1	192.168.10.128	MAPI	564.195	72	1
 	192.168.10.9	192.168.10.128	MAPI	12.200	17	2
	192.168.10.9	192.168.10.128	LDAP	3.316	1	1
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	2.902	1	1
	192.168.10.10	192.168.10.128	LDAP	0.565	0	2
	192.168.10.9	192.168.0.1	DNS	0.197	0	1
	192.168.10.9	192.168.10.128	LDAP	0.188	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.109	0	1

- To display the user name associated with an internal IP, select **Show Users**.


6. To group individual connections within a flow as a single line item or show each connection as a separate line item, select **Group**.

Understanding the Conversation Report

When a conversation has been accelerated by the Exinda appliance, the Conversation are highlighted in yellow and the Application Acceleration technologies being 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 is also been process by WAN Memory.



When a conversation has been processed by Edge Cache it is highlighted in blue.



The following legend describes the meaning of each 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 also appear next to each conversation.

	Asymmetric: The traffic is asymmetric, and is not being accelerated.
	Local: The connection is passing through this appliance in the cluster.
	Remote: The connection is passing through another appliance in the cluster.
	Local/Remote: The connection is passing though both this and other appliances in the cluster.

Troubleshoot problems with MAPI acceleration

If you are experiencing issues with accelerating Microsoft Outlook or Exchange traffic on the Exinda appliance, these troubleshooting topics may help resolve the issue.

- ["Notify administrators of system issues" on page 279](#)
- ["Outlook cannot connect to the Exchange server" on page 283](#)
- ["Outlook slow to send or receive emails" on page 283](#)
- ["Decrease in acceleration of MAPI traffic" on page 283](#)

Notify administrators of system issues

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. SLA and APS email alerts are sent when the set threshold limits are exceeded. Use the form below to disable alerts that you do not wish to trigger or receive emails and SNMP traps for.

Note You must configure valid SMTP and DNS settings prior to receiving email alerts. See ["Add an SMTP server for sending email notifications" on page 281](#) and ["Configure DNS and Domain Names" on page 159](#).

1. Click **System > Setup**, and switch to the **Alerts** tab.
2. To receive alert notifications, select the **Enable** checkbox for the appropriate alert.
For a description of when each alert is triggered, refer to the table below.
3. Select what type of notification to receive: an **Email**, an **SNMP Trap**, or both.
4. For CPU Utilization, Disk Usage, or NIC Collisions alerts, specify the **Trigger Threshold** and **Clear Threshold** levels that cause the notifications to be sent.
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**.

Alert Name	Description
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

Alert Name	Description
	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.
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.
System Startup	Alert raised when the Exinda appliance boots up.
SMB signed connections	Alert raised when SMB signed connections are present.
SLA Latency	Alert raised when the set latency for an SLA object is exceeded.
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.
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).
Connection Limiting	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.
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	Alert raised when encrypted MAPI traffic to a Microsoft Exchange

Alert Name	Description
Connections	server is detected on an Exinda appliance. Encrypted MAPI traffic cannot be accelerated.

Add an SMTP server for sending email notifications

An SMTP server is required for receiving scheduled reports, system alerts and auto-support notifications.

1. Click **System > Network**, and switch to the **Email** tab.
2. In the SMTP Server area, type the SMTP server name.
IPv4 and IPv6 addresses can be used.
3. Type the SMTP server port.
The default port number is 25.
4. Type the email address that system alerts and report notifications will appear to have been sent from.
5. To require authentication against the SMTP server before emails can be sent, select the **SMTP Authentication** checkbox.
After selecting SMTP Authentication, you must provide the username and password for the SMTP server, and select the authentication method.
6. Click **Apply Changes**.
7. To ensure the users can successfully receive notification emails, click **Send Test Email to Add**.

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. Click **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.
2. To view the alert that has triggered the warning, click the alarm name.
Use the information in this alert to help troubleshooting the issue.
3. 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

Alert Name	Description
	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.
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.

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 Networks Support Services.

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 Networks Support Services.

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

1. 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 "[Disable encryption on the Exchange server](#)" on page 275 and "[Turn off MAPI encryption in Microsoft Outlook](#)" on page 273.

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.

Specify application quality based on host

Per Host Quality of Service (QoS) allows you to manage traffic congestion by policing 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.

Tip Per Host QoS can be integrated with Active Directory so bandwidth management can be tailored to users or groups.

1. (Optional) Integrate the Exinda appliance with Active Directory. Refer to the Active Directory guide.
2. "[Set a per-host limit on bandwidth usage](#)" on page 123
3. "[View the number of hosts on a Dynamic Virtual Circuit](#)" on page 287

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. Click **Optimizer**.
2. Click **Create New Virtual Circuit**.
3. Type a name for the virtual circuit.

4. Type the amount of bandwidth to be used by the virtual circuit.
5. To enable Per Host QoS, select the **Dynamic Virtual Circuit** checkbox.
6. Set the amount of bandwidth (in KB per second or percentage of the virtual circuit bandwidth) that each host will receive in the **Per Host Bandwidth** field.

This bandwidth is guaranteed, so it will be available to each host, if required.

To have the amount of bandwidth each host receives calculated by dividing the Virtual Circuit guaranteed bandwidth by the number of active hosts, select **Automatically Share**.

7. Set the maximum amount of bandwidth (in KB per second or percentage of the virtual circuit bandwidth) that each host can burst to in the **Per Host Max Bandwidth** field.

If **No Bursting Allowed** is selected, each host only gets the bandwidth that they have been guaranteed.

8. Set the location of the hosts to allocate bandwidth to.

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.

9. Set the maximum number of hosts that can use the Dynamic Virtual Circuit.

If **Auto** is selected, the maximum number of hosts is calculated by assuming each host gets its guaranteed bandwidth.

If **Automatically Share** is selected, the maximum number of hosts is calculated by assuming each host is entitled to minimum bandwidth, which is 10kbps.

Any host that becomes active after the maximum number of hosts is exceeded do not fall into this Virtual Circuit.

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.

The following examples describe various Dynamic Virtual Circuit configurations.

<p>Name: Example 1</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: Auto</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <p>If there is 1 user, they get the full 1024kbps.</p> <p>If there are 2 users, they each get 512kbps and can burst up to the full 1024kbps (if the other user is not using their guaranteed 512kbps).</p> <p>If there are 10 users, they each get 102kbps and can burst up to the full 1024kbps (if the other users are not using their guaranteed 102kbps).</p>
---	---

<p>Per User Max Bandwidth: 100%</p> <p>Host Location: Internal</p> <p>Max Hosts: Auto</p>	
<p>Name: Example 2</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: 10%</p> <p>Per User Max Bandwidth: No</p> <p>Host Location: Internal</p> <p>Max Hosts: Auto</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <p>If there is 1 user, they get 102kbps and cannot burst.</p> <p>If there are 10 users, they each get 102kbps and cannot burst.</p> <p>If there are 100 users, the first 10 users each get 102kbps and cannot burst. The remaining 90 users will not match this VC.</p>
<p>Name: Example 3</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: 64kbps</p> <p>Per User Max Bandwidth: 50%</p> <p>Host Location: Internal</p> <p>Max Hosts: 16</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <p>If there is 1 user, they get 64kbps and can burst up to 512kbps.</p> <p>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).</p> <p>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.</p>
<p>Name: Example 4</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Application: Citrix</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <p>"Citrix" is an Application that defines Citrix traffic. This VC will match all Internal User's Citrix traffic.</p> <p>If there is 1 user, they get 64kbps for their Citrix traffic and cannot burst.</p> <p>If there are 16 users, they each get 64kbps for their</p>

<p>Dynamic Virtual Circuits Enabled: Yes Per Host Bandwidth: 64kbps Per User Max Bandwidth: No Host Location: Internal Max Hosts: 16</p>	<p>Citrix traffic and cannot burst. 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.</p>
---	--

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. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.
4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.
5. 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
```
6. 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. Click **Monitor > Control > Policies**.
2. 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.

View the data throughput on the interfaces

The Interface Throughput Report provides you with statistics of the total data that has passed through each WAN interface on each bridge. This report allows you to see the inbound and outbound throughput for all traffic on the wire, over time.



1. Click **Monitor > Interfaces > Throughput**.
2. To filter the data in the chart, select the WAN or out-of-path interface to display.
Bridge WAN ports, policy-based routing interfaces, and WCCP interfaces are available. Selecting **All** includes all out-of-path interfaces in the report.
3. To identify how much traffic falls above a specific percentile, select the desired value from the **Select Percentile Marker to Display** list.

The table at the bottom of the report shows the total amount of data transferred into and out of the WAN interface(s), and also the maximum and average throughput values for the selected time period. The values in the table are automatically updated when the interactive flash graphs are manipulated.

Note Given that this report shows all data on the wire, the report may also include traffic that is 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	1200094.45	113.39	286.88
Outbound	553478.38	52.30	136.31

[View the outbound packet rate for all traffic](#)

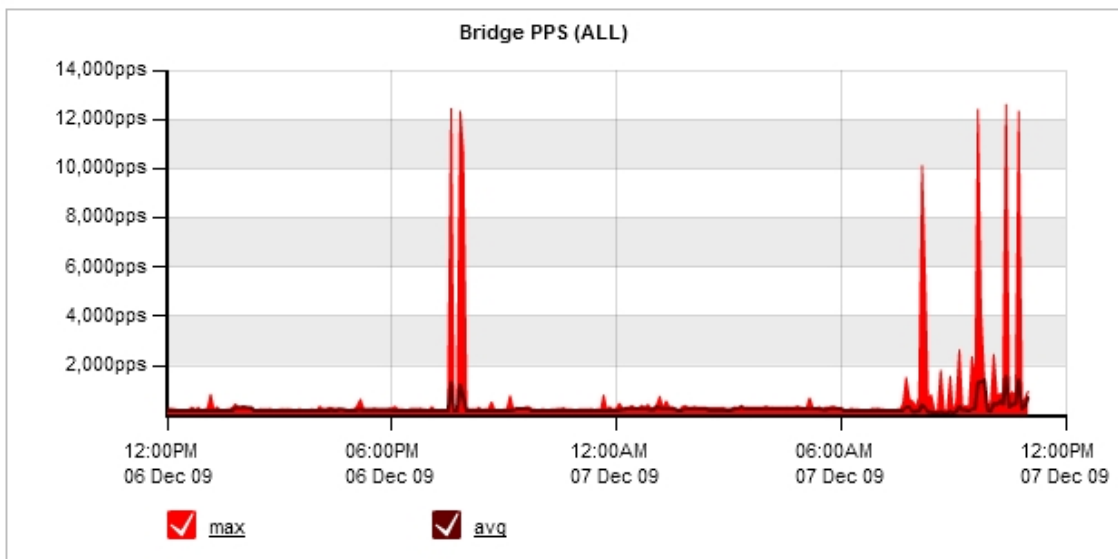
The Interface PPS Report provides you with the packet rate that has passed through each bridge on the Exinda appliance. This report allows you to see the outbound packet rate for all traffic on the wire, over time.

1. Click **Monitor > Interfaces > Packets Per Second**.
2. To filter the data in the chart, select the Bridge or out-of-path interface to display.
Bridge WAN ports, policy-based routing interfaces, and WCCP interfaces are available. Selecting **All** includes all out-of-path interfaces in the report.
3. To identify how much traffic falls above a specific percentile, select the desired value from the **Select Percentile Marker to Display** list.

The table at the bottom of the report shows the maximum and average PPS values through the bridge for the selected time period. The values in the table are automatically updated when the interactive flash graphs are manipulated.

Note Given that this report shows all data on the wire, the report may also include traffic that is not seen on the WAN, such as local LAN broadcasts, etc.

Bridge/Out-of-path Interface Selection: ALL
Select Percentile Marker to Display: None



Bridge PPS Summary (ALL)		
Data Direction	Packets Per Second (Avg)	Packets Per Second (Max)
Outbound	215	12,629

Per Host QoS Usage Examples

The following examples show how Per Host QoS can be used in a variety of situations.

- "Limit Bandwidth Per Host" on page 291
- "Limit Application Bandwidth" on page 292
- "Guarantee Application Bandwidth" on page 295
- "Per Host QoS with Active Directory" on page 298
- "Per Host QoS for Adaptive Response" on page 303

Limit Bandwidth Per Host

Example #1

Limit bandwidth to 100 kbps for each internal host.

Edit Virtual Circuit	
Virtual Circuit Number	10 . 10
Virtual Circuit Name	WAN
Schedule	ALWAYS
Bandwidth Options	
Virtual Circuit Bandwidth	50000 kbps
Oversubscription	<input checked="" type="radio"/> Automatic <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input checked="" type="checkbox"/> Automatically Share 0 kbps
Per Host Max Bandwidth	<input type="checkbox"/> No Bursting Allowed 100 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	ALL
Application	ALL
Direction	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.

Optimize 🖨️ ?

Optimize Policies Wizard

		Operations
Circuit 10 - Default (50000 kbps)		--Actions--
Dynamic Virtual Circuit 10 - WAN (50000 kbps [auto 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)	--Actions--
Order:	Policy: ALL - Accelerate Add To 'WAN'	
Create New Policy...		
Create New Virtual Circuit...		
Create New Circuit...		

Example #2

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.

Optimize 🖨️ ?

Optimize Policies Wizard

		Operations
Circuit 10 - Default (50000 kbps)		--Actions--
Dynamic Virtual Circuit 10 - WAN (50000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'ALL')		--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 Add To 'WAN'	
Create New Policy...		
Create New Virtual Circuit...		
Create New Circuit...		

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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input checked="" type="checkbox"/> Automatically Share 0 kbps
Per Host Max Bandwidth	<input type="checkbox"/> No Bursting Allowed 20 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
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 . 10
Virtual Circuit Name	P2P Overflow
Schedule	ALWAYS
Bandwidth Options	
Virtual Circuit Bandwidth	100 kbps
Oversubscription	<input checked="" type="radio"/> Automatic <input type="radio"/> Manual
Dynamic Virtual Circuit	<input type="checkbox"/>
Filter Options	
VLAN Object	ALL
Network Object	ALL
Application	P2P
Direction	Both

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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input checked="" type="checkbox"/> Automatically Share 0 kbps
Per Host Max Bandwidth	<input type="checkbox"/> No Bursting Allowed 100 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	ALL
Application	ALL
Direction	Both

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.

The screenshot shows the 'Optimize' configuration window with three tabs: 'Optimize', 'Policies', and 'Wizard'. The 'Optimize' tab is active, displaying a list of virtual circuits and their associated policies. Each circuit has a checkbox, a priority value, a policy name, and an 'Add To' button. The 'Operations' column contains dropdown menus for each circuit.

Circuit Name	Priority	Policy	Add To	Operations
Circuit 10 - Default (50000 kbps)	10	ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)	Add To 'P2P'	--Actions--
Dynamic Virtual Circuit 5 - P2P (5000 kbps [auto kbps - 20 kbps per user / auto users max] 'P2P' traffic to / from 'ALL')	10	ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)	Add To 'P2P Overflow'	--Actions--
Virtual Circuit 10 - P2P Overflow (100 kbps 'P2P' traffic to / from 'ALL')	10	ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)	Add To 'WAN'	--Actions--
Dynamic Virtual Circuit 25 - WAN (45000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'ALL')	10	ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)		--Actions--

At the bottom of the window, there are links for 'Create New Policy...', 'Create New Virtual Circuit...', and 'Create New Circuit...'.

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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input type="checkbox"/> Automatically Share 30 kbps
Per Host Max Bandwidth	<input type="checkbox"/> No Bursting Allowed 100 %
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	ALL
Application	Citrix
Direction	Both

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	40000 kbps
Oversubscription	<input checked="" type="radio"/> Automatic <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input checked="" type="checkbox"/> Automatically Share 0 kbps
Per Host Max Bandwidth	<input checked="" type="checkbox"/> No Bursting Allowed 0 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	ALL
Application	ALL
Direction	Both

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.

Optimize 🖨️ ?

Optimize Policies Wizard

		Operations
Circuit 10 - Default (50000 kbps)		--Actions--
Dynamic Virtual Circuit 5 - Citrix (10000 kbps [30 kbps - 100% per user / auto users max] 'Citrix' traffic to / from 'ALL')		--Actions--
<input checked="" type="checkbox"/>	10 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)	--Actions--
Order:	Policy: ALL - Accelerate Add To 'Citrix'	
Create New Policy...		
Dynamic Virtual Circuit 25 - WAN (40000 kbps [auto 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)	--Actions--
Order:	Policy: ALL - Accelerate Add To 'WAN'	
Create New Policy...		
Create New Virtual Circuit...		
Create New Circuit...		

Per Host QoS with Active Directory

Example

Restrict users in the Active Directory 'Students' group to 100 kbps.

Using the Web UI - Advanced Mode, navigate to Objects | Users & Groups. Edit the "Students (DEV)" group.

Welcome to **exinda**, logged in as **admin** (advanced, switch to basic). [Logout](#)

Optimizer Status : On (Restart / Stop) | Config Status No unsaved changes | System Health : OK | Thu Apr

Users & Groups

Network Users | **Network Groups**


Network Groups (Total: 12)

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]

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 guests (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"/>
Guests (DEV)	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
Remote desktop users (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"/>

Active Directory Groups

Create a Network Object from the Active Directory group.

Welcome to exinda , logged in as admin (advanced, switch to basic).  Logout

Optimizer Status : On (Restart / Stop) | Config Status No unsaved changes | System Health : OK | Thu Apr 8, 2010 00:34:32

Dashboard

System

Objects

Network

Users & Groups

VLANs

Protocols

Applications

Schedules

Adaptive Response

Monitor

Report

Optimize

[+] Expand ALL

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

Map to Network Object

Ignore Domain

Map AD Group 'Students' To Network Object 'Students'

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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input checked="" type="checkbox"/> Automatically Share 0 kbps
Per Host Max Bandwidth	<input type="checkbox"/> No Bursting Allowed 100 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	Students
Application	ALL
Direction	Both

Students Dynamic Virtual Circuit

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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input type="checkbox"/>
Filter Options	
VLAN Object	<input type="text" value="ALL"/>
Network Object	<input type="text" value="Students"/>
Application	<input type="text" value="ALL"/>
Direction	<input type="text" value="Both"/>

Students Overflow Virtual Circuit

Edit 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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input type="checkbox"/> Automatically Share 100 kbps
Per Host Max Bandwidth	<input checked="" type="checkbox"/> No Bursting Allowed 0 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	ALL
Application	ALL
Direction	Both

Dynamic Virtual Circuit For Remaining Bandwidth

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**.

Optimize

Optimize Policies Wizard

Circuit 10 - Default (50000 kbps)

Dynamic Virtual Circuit 5 - Students (24000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'Students')

10 ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)

Order: Policy: ALL - Accelerate

[Create New Policy...](#)

Virtual Circuit 10 - Students Overflow (1000 kbps to / from 'Students')

10 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)

Order: Policy: ALL - Accelerate

[Create New Policy...](#)

Dynamic Virtual Circuit 15 - WAN (25000 kbps [100 kbps per user / auto users max] to / from 'ALL')

10 ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)

Order: Policy: ALL - Accelerate

[Create New Policy...](#)

[Create New Virtual Circuit...](#)

[Create New Circuit...](#)

Operations

--Actions--

--Actions--

--Actions--

--Actions--

--Actions--

--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.

Create the Network Object **Students** based on the Active Directory **Students** group as shown in the previous chapter.

Using the Web UI - Advanced Mode, navigate to **Objects > Adaptive Response**.

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.

Welcome to **exinda**, logged in as **admin** (advanced, switch to **basic**). [Logout](#)

Optimizer Status : On (Restart / Stop) | Config Status Unsaved changes (Save) | System Health : OK | Thu Apr 8, 2010 03:59:01 | v5.4.0.13281

Adaptive Response

Adaptive Response Limits are rules which are used to create and populate network objects based on amount of data transferred. They then be used when creating virtual circuits or filters.

Add New AR Limit

Name:

Source Network Object:

Destination Network Object:

Duration:

Direction:

Amount (MB):

Enable:

Name	Source Network	Destination Network	Duration	Direction	Amount	Enabled	Edit	Delete
No AR Limits.								

Create Adaptive Response Object

Create a Dynamic Virtual Circuit, with **Network Object** set to "**Students_Shaped**". Hosts matching this Network Object will fall into this Dynamic Virtual Circuit.

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 <input type="radio"/> Manual
Dynamic Virtual Circuit	<input checked="" type="checkbox"/>
Dynamic Options	
Per Host Bandwidth	<input checked="" type="checkbox"/> Automatically Share 0 kbps
Per Host Max Bandwidth	<input type="checkbox"/> No Bursting Allowed 100 kbps
Host Location	Internal
Max Hosts	<input checked="" type="checkbox"/> Auto 0
Filter Options	
VLAN Object	ALL
Network Object	Students_Shaped
Application	ALL
Direction	Both

Students Dynamic Virtual Circuit

Additional hosts will share bandwidth allocated in the Students Overflow Virtual Circuit.

Edit Virtual Circuit	
Virtual Circuit Number	10 . 10
Virtual Circuit Name	Students Overflow
Schedule	ALWAYS
Bandwidth Options	
Virtual Circuit Bandwidth	1000 kbps
Oversubscription	<input checked="" type="radio"/> Automatic <input type="radio"/> Manual
Dynamic Virtual Circuit	<input type="checkbox"/>
Filter Options	
VLAN Object	ALL
Network Object	Students_Shaped
Application	ALL
Direction	Both

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: Automatic Manual

Dynamic Virtual Circuit:

Filter Options

VLAN Object: ALL

Network Object: ALL

Application: ALL

Direction: Both

Virtual Circuit To Share Remaining Bandwidth

Other users and students who have not used their 100MB daily quota will share 45 Mbps of bandwidth in the WAN Virtual Circuit.

The screenshot shows the 'Optimize' configuration page with three tabs: 'Optimize', 'Policies', and 'Wizard'. The 'Optimize' tab is active, displaying a list of virtual circuits and their associated policies. The circuits are:

- Circuit 10 - Default (50000 kbps)**: Dynamic Virtual Circuit 5 - Students (4000 kbps [auto kbps - 100 kbps per user / auto users max] to / from 'Students_Shaped'). Policy: ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7).
- Virtual Circuit 10 - Students Overflow (1000 kbps to / from 'Students_Shaped')**: Policy: ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5).
- Virtual Circuit 15 - WAN (45000 kbps to / from 'ALL')**: Policy: ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5).

Each circuit entry includes a checkbox, a priority field (set to 10), a policy name, and an 'Add To' button. There are also links for 'Create New Policy...' and 'Create New Virtual Circuit...'.

Controlling Traffic based on Users










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. See "Integrate the Exinda Appliance with Active Directory" on page 165.

1. "Create Network User Objects" on page 11 and "Create Network Group Objects" on page 12.
2. "Optimize Traffic Based on Users and Groups" on page 309

Create Network User Objects

Network Users displays a pre-populated list of Users (and their associated IP addresses) from either the Exinda Active Directory 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, it 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	


To define a user as a Dynamic Network Object

1. In the Exinda WebUI, go to **Objects > Users & Groups > Network Users**.
2. Select the checkbox for the user.
3. Click **Add Network Object**.

The Network Status icon for the user changes to , indicating it is a network object.

To stop identifying a user as a dynamic network object

1. Select the checkbox for the user.
2. Click **Remove Network Object**.

The Network Status icon for the user changes to , indicating it is no longer a network object.

Create Network Group Objects

Network Groups displays a pre-populated list of Groups from either the Exinda Active Directory 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.

To define a group as a Dynamic Network Object


1. In the Exinda WebUI, go to **Objects > Users & Groups > Network Groups**.
2. Locate 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 , indicating it is 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 stop identifying a group as a Dynamic Network Object

1. Locate the group in the list, and click **Delete**.

The Network Status icon for the user changes to , indicating it is no longer a network object.

If the dynamic network object was created from multiple groups, each group is again listed individually in the list.

Optimize Traffic Based on Users and Groups

Create policies that affect the traffic based on the source or destination host.

Note Active Directory must be configured before optimization policies can target specific users and groups. See ["Integrate the Exinda Appliance with Active Directory" on page 165](#).

1. Click **Optimizer > Policies**.
2. Type a name for the policy.
3. Set the required bandwidth and acceleration parameters.
4. 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.
5. Click **Create New Policy**.
6. Once the desired policies are in place on all Exinda appliances, restart the Optimizer. In the appliance

status bar, click **Restart**.

Optimizer Status : On (Restart / Stop)


Add SSL certificates and keys on the Exinda appliances

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's desktop. The SSL is used protocol to secure each session link by automatically establishing connections on-demand using standards-based protocols, encryption techniques, and certificate exchange.

SSL encryption requires a server certificate to authenticate the identity of a server. A certificate is an electronic confirmation that the owner of a public key is who he or she really claims to be, and holds the private key corresponding to the public key in the certificate.

A public key and a particular server can communicate securely by generating a certificate signing request using the server's public key. The identity of the server is verified by a Certificate Authority (CA), and generates a signed certificate. The server offers the generated certificate to clients that recognize the CA's signature, and verifies the identity of the server.

The process of setting up the private connection is automatically initiated by the server communicating directly with the browser. The result is a private, encrypted tunnel used to move information between the server and client desktop. When the session is over, the connection is automatically terminated.



If the connection between the client and the server uses SSL to encrypt the sessions, the benefits that can be gained by application acceleration are limited. For example, Exinda's WAN Memory technology achieves higher reduction on clear text rather than encrypted data.

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.

Note The SSL Acceleration settings on each Exinda appliance be configured exactly the same.

SSL certificates and keys can be added to the Exinda appliance in one of the following ways:

- "Import a certificate onto an appliance" on page 253 that has been [exported from another appliance](#).
- "Generate a self-signed certificate" on page 254

Virtualization

The Virtualization feature allows Virtual Machines to be run on the the Exinda appliance.


Note Virtualization requires an additional, optional license before this feature can be configured and used. Please contact Exinda TAC or your local Exinda representative if you do not have this license and you wish to use this feature.

Most of the Virtualization configuration is performed using the "virt" CLI command. The table below show all configured Virtual Machines and provides the ability to Power them ON, OFF or restart (cycle) them. You can also view the VGA console for each Virtual Machine. The VA console uses a Java-based applet to launch secure SSH-based protocol to encrypt the session. You will need to have Java installed and enabled as well as direct SSH access to the Exinda appliance in order to use this feature. You will also need to authenticate with your username and password.

Virtual Machines			
Name	Comment	Status	Actions
<input type="checkbox"/>	Replify	Replify VA	Running - IP Address: 172.16.1.242 <input type="button" value="Launch VGA Console"/>
<input type="checkbox"/>	WinXP	Windows XP	Running <input type="button" value="Launch VGA Console"/>

Virtualization overview

Exinda Virtualization support allows 3rd party operating systems/products to be installed on selected Exinda hardware appliances.



Once enabled, you can install any standard 32 or 64 bit x86 (Intel based) operating system onto an Exinda Virtualization Partition (EVP).

This How to Guide provides a general overview on how to use Exinda Virtual Partitions and also how to install supported 3rd party products.

Note Virtualization requires an additional, optional license before this feature can be configured and used. Please contact Exinda TAC or your local Exinda representative if you do not have this license and you wish to use this feature.

Virtualization requirements

In order to enable Virtualization support, the following requirements must be satisfied:




- Only selected Exinda 60 series hardware is supported. Currently, the following hardware is available for Virtualization:
 - Exinda 4061*


Note The Exinda 4061 requires a 2GB RAM upgrade in order to enable Virtualization. This RAM is provided free of charge for all 4061s when an Exinda Virtualization license is purchased for that 4061.
 - Exinda 6060
 - Exinda 8060
 - Exinda 10060
- An optional Exinda Virtualization license must be purchased for each Exinda appliance you want to enable Virtualization support.
- ExOS 5.5 (or later) must be installed.

Once an Exinda Virtualization license has been installed on supported hardware, the Virtualization functionality will be enabled.

Products supported on the virtualization partition

Although almost any operating system can be installed onto an Exinda Virtualization Partition, Exinda currently approves and supports the following 3rd party products:

Vendor	Product	Description
	Exinda SDP	Centrally Manage Unified Performance Management Solutions in Distributed Deployments.
	Exinda Mobile Server	Exinda Mobile Servers can be configured as required, to emulate Data Center or Branch Office boxes, offering system administrators the maximum optimization and deployment flexibility benefits the Exinda Mobile Suite has to offer.
	Exinda Mobile Manager	The Exinda Mobile Manager provides centralized configuration of your Exinda

		Mobile Suite deployment, with detailed reporting on all of your Exinda Mobile Client and Exinda Mobile Server.
	Windows Server 2008 R2	Windows Server 2008 R2 builds on the successes and strengths of its Windows Server predecessors while delivering valuable new functionality and powerful improvements to the base operating system.

Installing Virtualization

Most of the Exinda Virtualization configuration is done via the CLI "virt vm" command. These commands can be used to create virtual machines, configure CPU, storage, RAM, network, etc.

```

exinda-16d806 (config) # virt vm MyVM ?
<cr>          Create this virtual machine (if it does not already exist)
arch          Set CPU architecture
boot          Configure boot options
comment       Set a comment describing this virtual machine
console       Configure or connect to the text or graphical console
copy          Make a duplicate copy of this virtual machine
feature       Enable certain virtualization features
install       Install an operating system onto this virtual machine (temporarily attach a CD and boot from it)
interface     Configure virtual interfaces
manufacture   Manufacture this virtual machine with an appliance image
memory        Set memory allowance
power         Turn this virtual machine on or off, plus other related options
rename        Rename this virtual machine
storage       Configure storage for this virtual machine
vcpus         Specify virtual CPUs
    
```

The **System > Virtualization** page on the Web UI, advanced mode, allows you to perform basic operations on the virtual machine, such as power on and off. You can also launch the VGA console.

Configure Virtualization options.

Virtualization

Virtualization Enable

The following table lists the configured Virtual Machines.

Virtual Machines			
Name	Comment	Status	Actions
<input type="checkbox"/> Exinda-Mobile-SRV	Exinda Mobile Server	Running	<input type="button" value="Launch VGA Console"/>

Note The 'Launch VGA Console' command requires Java to be installed and enabled. This will create 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.

Exinda SDP



Exinda's Service Delivery Point (SDP) is a revolutionary platform for centrally managing Exinda appliances distributed throughout the corporate network. As a virtual appliance, SDP further simplifies the task of installing, configuring, monitoring and reporting WAN optimization.

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.

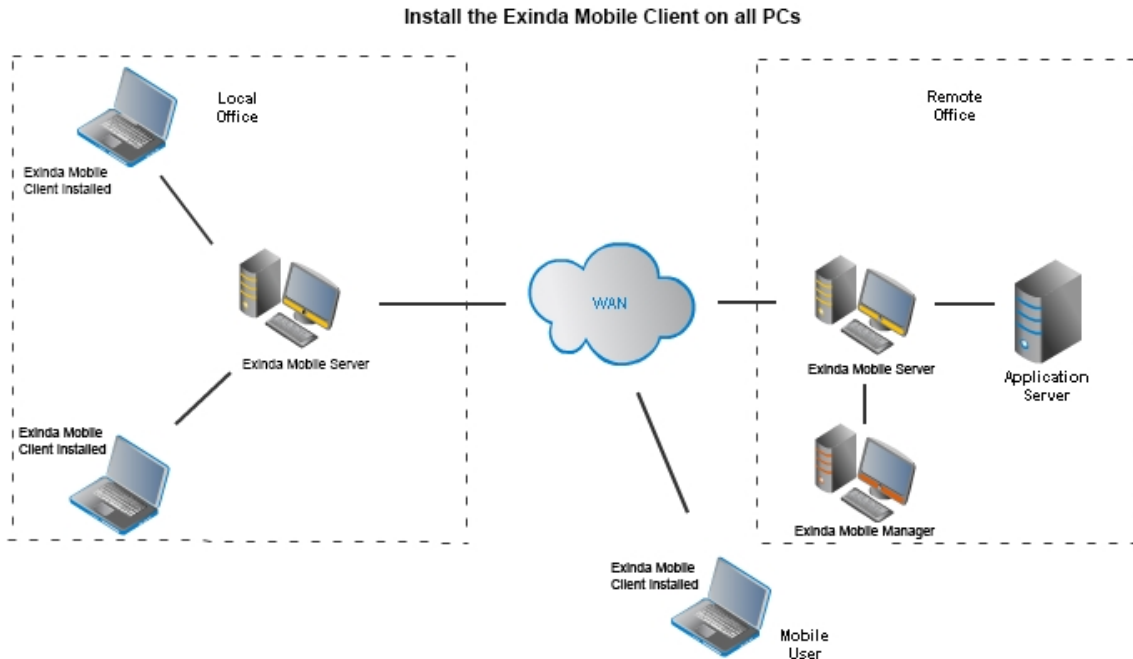
Note

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.

Exinda Mobile Suite



Exinda Mobile is a fully featured, WAN Optimization Software Client for mobile workers and branch offices.



The Exinda Mobile Suite vastly improves remote user experience through compression, cross-protocol data reduction, protocol manipulation and TCP optimization. Users will experience vastly improved response times and download speeds over a wide range of applications - including email, CRM, ERP, and collaboration applications.

Note

For more information how to configure Exinda Mobile, consult the [Exinda Mobile User Guide](#).

Install the Exinda Mobile Server on an Exinda appliance

Exinda Virtualization support allows 3rd party operating systems and products to be installed on selected Exinda hardware appliances. For additional information, see the Exinda *Virtualization* how-to guide.

Exinda Mobile Server (EMS) is the traffic-handling component of Exinda Mobile, and is responsible for intercepting and accelerating network traffic for remote users.

Note The text of the underlined CLI commands may wrap on the page, but should be entered as a single line.

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

- To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config) #` prompt is displayed.

- Assign a physical interface to the virtual infrastructure so that the virtual machine can have network connectivity.

The following command add the physical interface specified to a bridge that can later be attached to virtual machines. 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". Use this bridge later on when configuring the virtual machine's network interfaces.

```
virt interface eth2
```

- Download the 500MB image file from http://updates.exinda.com/exinda_mobile/KVM/ExindaMobileServer.tar.gz, and unzip the file.

- If you have shell access, unzip the file to `/var/opt/tms/virt/pools/default/`
- If you do not have shell access, unzip the file to a local server from which you can secure copy, FTP, or access it through HTTP with the following command:

```
virt volume fetch <file-location>
```

where `<file-location>` can be one of the following:

```
http://<hostname>/<path>/ExindaMobileServer.img
```

```
https://<hostname>/<path>/ExindaMobileServer.img
```

```
scp://<username>[:<password>]@<hostname>/<path>/ExindaMobileServer.img
```

```
ftp://<username>[:<password>]@<hostname>/<path>/ExindaMobileServer.img
```

```
tftp://<username>[:<password>]@<hostname>/<path>/ExindaMobileServer.img
```

```
sftp://<username>[:<password>]@<hostname>/<path>/ExindaMobileServer.img
```

- Create an empty 40GB virtual disk image for additional storage.

```
virt volume create disk file ExindaMobileServerDisk2.img size-max 40000
```

- After the previous commands have completed, create a VM called "ExindaMobileServer" with the following commands:

```
virt vm ExindaMobileServer
```

```
virt vm ExindaMobileServer arch x86_64
```

```
virt vm ExindaMobileServer boot auto-power on
```

```
virt vm ExindaMobileServer boot device order hd
```

```
virt vm ExindaMobileServer comment "Exinda Mobile Server"
```

```
virt vm ExindaMobileServer console graphics vnc
```

```
virt vm ExindaMobileServer console text tty
```

```
virt vm ExindaMobileServer feature acpi enable
```

```
virt vm ExindaMobileServer feature apic enable
```

```

no virt vm ExindaMobileServer feature pae enable
# Specify the bridge created above when setting the virtual interface.
virt vm ExindaMobileServer interface 1 bridge brvm2
virt vm ExindaMobileServer interface 1 model e1000
no virt vm ExindaMobileServer interface 2
virt vm ExindaMobileServer memory 1024
virt vm ExindaMobileServer storage device bus ide drive-number 1 source file
ExindaMobileServer.img
virt vm ExindaMobileServer storage device bus ide drive-number 2 source file
ExindaMobileServerDisk2.img
virt vm ExindaMobileServer vcpus count 2

```

9. Click **System > Virtualization**.

Select the virtual machine, and click **Power On**.

Launch the VGA console to view the ExindaMobileServer VM boot up, then login to the CLI by pressing **ALT-F2** after boot-up. The default CLI login credentials are:

Username: root

Password: default

10. From the CLI, change the IP address and network connectivity settings using the `configure-network` command, and expand the Exinda Mobile Server file system to make use of the additional 40GB storage using the `add-new-disks` command.
11. After IP address and network connectivity settings have been saved, you can then access the web-based UI by pointing your browser to `http://<ExindaMobileServer-IP>`. The default Exinda Mobile Server login credentials are:

Username: admin

Password: default

Note The Exinda Mobile Server's management 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 Exinda Mobile Server UI.

Install the Exinda Mobile Manager on an Exinda appliance

Exinda Virtualization support allows 3rd party operating systems and products to be installed on selected Exinda hardware appliances. For additional information, see the Exinda *Virtualization* how-to guide.

The Exinda Mobile Manager (EMM) is the system controller responsible for licensing Exinda Mobile Servers and Exinda Mobile Clients, configuring components, and reporting on overall system behavior.

Note The text of the underlined CLI commands may wrap on the page, but should be entered as a single line.

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

5. Assign a physical interface to the virtual infrastructure so that the virtual machine can have network connectivity.

The following command add the physical interface specified to a bridge that can later be attached to virtual machines. 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". Use this bridge later on when configuring the virtual machine's network interfaces.

```
virt interface eth2
```

6. Download the base virtual disk image file with one of the following methods:

- a. Download the 2GB image file from:

```
virt volume fetch url http://updates.exinda.com/exinda_
mobile/KVM/ExindaMobileManager.img
```

- b. Download the 500MB image file from `http://updates.exinda.com/exinda_
mobile/KVM/ExindaMobileManager.tar.gz`, and unzip the file.

- If you have shell access, unzip the file to `/var/opt/tms/virt/pools/default/`
- If you do not have shell access, unzip the file to a local server from which you can secure copy, FTP, or access it through HTTP with the following command:

```
virt volume fetch <file-location>
```

where `<file-location>` can be one of the following:

```
http://<hostname>/<path>/ExindaMobileManager.img
https://<hostname>/<path>/ExindaMobileManager.img
scp://<username>[:<password>]
@<hostname>/<path>/ExindaMobileManager.img
ftp://<username>[:<password>]
@<hostname>/<path>/ExindaMobileManager.img
tftp://<username>[:<password>]
@<hostname>/<path>/ExindaMobileManager.img
sftp://<username>[:<password>]
@<hostname>/<path>/ExindaMobileManager.img
```

7. Create an empty 3GB virtual disk image for additional storage.

```
virt volume create disk file ExindaMobileManagerDisk2.img size-max 3000
```

8. After the previous commands have completed, create a VM called "ExindaMobileManager" with the following commands:

```

virt vm ExindaMobileManager
virt vm ExindaMobileManager arch x86_64
virt vm ExindaMobileManager boot auto-power on
virt vm ExindaMobileManager boot device order hd
virt vm ExindaMobileManager comment "Exinda Mobile Manager"
virt vm ExindaMobileManager console graphics vnc
virt vm ExindaMobileManager console text tty
virt vm ExindaMobileManager feature acpi enable
virt vm ExindaMobileManager feature apic enable
no virt vm ExindaMobileManager feature pae enable
# Specify the bridge created above when setting the virtual interface.
virt vm ExindaMobileManager interface 1 bridge brvm2
virt vm ExindaMobileManager interface 1 model e1000
no virt vm ExindaMobileManager interface 2
virt vm ExindaMobileManager memory 1024
virt vm ExindaMobileManager storage device bus ide drive-number 1 source file
ExindaMobileManager.img
virt vm ExindaMobileManager storage device bus ide drive-number 2 source file
ExindaMobileManagerDisk2.img
virt vm ExindaMobileManager vcpus count 2

```

9. Click **System > Virtualization**.

Select the virtual machine, and click **Power On**.

Launch the VGA console to view the ExindaMobileManager VM boot up, then login to the CLI by pressing **ALT-F2** after boot-up. The default CLI login credentials are:

Username: root

Password: default

10. From the CLI, change the IP address and network connectivity settings using the `configure-network` command, and expand the Exinda Mobile Manager file system to make use of the additional 3GB storage using the `add-new-disks` command.
11. After IP address and network connectivity settings have been saved, you can then access the web-based UI by pointing your browser to `http://<ExindaMobileManager-IP>`. The default Exinda Mobile Manager login credentials are:

Username: admin

Password: default

Note The Exinda Mobile Manager's management 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 Exinda Mobile Manager UI.

Microsoft Windows Server



Windows Server 2008 R2 builds on the successes and strengths of its Windows Server predecessors while delivering valuable new functionality and powerful improvements to the base operating system.

Microsoft Windows Server 2008 R2

To install Microsoft Windows Server 2008 R2 (Enterprise Edition), 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 Windows Server 2008 R2 Enterprise Edition (180 day evaluation) virtual disk image. The file is approximately 2.7GB, and once downloaded, will be uncompressed to about 6.7GB.

```
virt volume fetch url http://updates.exinda.com/vm/microsoft/Windows-Server-2008-R2-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 "Windows-Server-2008-R2".

```
virt vm Windows-Server-2008-R2
virt vm Windows-Server-2008-R2 arch x86_64
virt vm Windows-Server-2008-R2 boot auto-power on
virt vm Windows-Server-2008-R2 boot device order hd
virt vm Windows-Server-2008-R2 comment "Windows Server 2008 R2"
virt vm Windows-Server-2008-R2 console graphics vnc
virt vm Windows-Server-2008-R2 console text tty
virt vm Windows-Server-2008-R2 feature acpi enable
virt vm Windows-Server-2008-R2 feature apic enable
no virt vm Windows-Server-2008-R2 feature pae enable
# Specify the bridge created above when setting the virtual interface.
virt vm Windows-Server-2008-R2 interface 1 bridge brvm2
virt vm Windows-Server-2008-R2 interface 1 model e1000
no virt vm Windows-Server-2008-R2 interface 2
virt vm Windows-Server-2008-R2 memory 2048
```

```
virt vm Windows-Server-2008-R2 storage device drive-number 1 source file
Windows-Server-2008-R2-x86_64-0.img
virt vm Windows-Server-2008-R2 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. You should then launch the VGA console and continue with the configuration of the Windows Server. The default login credentials are:

Username: Administrator

Password: Pass@word1

By default, the Windows Server's license has not been activated. This is a special 180 day evaluation version of Windows Server 2008 R2 Enterprise Edition. There is a 10 day grace period before the Windows Server will require activation, so you should activate the Windows Server as soon as possible after installation in order to start the 180 day evaluation period (leave the Product Key blank when activating).

Note The Windows Server'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 Windows Server.

The virtual disk image used for this Windows Server is an automatically expanding image. When first installed, it will be about 6.7GB, however, it is capable of automatically growing up to 130GB as space is used. Ensure you allow enough free space on the Exinda appliance for the image to grow.

Authentication

The Authentication section of the Exinda appliance System Setup allows you to configure Local User Accounts, as well as remote authentication mechanisms such as LDAP (including Active Directory), Radius or TACACS+. The various configuration pages include:

- [Active Users](#): View a list of currently logged in users.
- [Local User Accounts](#): Configure local usernames and passwords.
- [AAA](#): Specify how users are authenticated and authorized on the Exinda appliance.
- [LDAP](#): Configure remote LDAP authentication servers.
- [Radius](#): Configure remote Radius authentication servers.
- [TACACS+](#): Configure remote TACACS+ authentication servers.

Display a List of Active Users

Active Users lists the users currently logged into either the Web UI or the CLI.

1. Click **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

Local User Accounts

Local User Accounts allows you to add/remove local user accounts as well as change local user's passwords.

1. Click **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"/>

2. 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**.
3. 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"/>

4. 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	admin ▼
New Password	<input type="text"/>
Confirm Password	<input type="text"/>
<input type="button" value="Change Password"/>	

AAA

AAA configures how remote users should authenticate to the Exinda appliance and what privileges they should receive.

1. To configure AAA, navigate to **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	Local ▼
Second Method	Local ▼
Third Method	Local ▼
Fourth Method	Local ▼
<input type="button" value="Apply Changes"/>	

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	remote-first ▼
Map Default User	admin ▼

Apply Changes

	remote-first	Apply user privileges supplied by the remote authentication mechanism first. If that fails, use the 'Map Default User' setting below.
Map Order	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.

- Click **Apply Changes**.

LDAP Authentication

LDAP authentication allows you to configure the Exinda appliance to authenticate user login attempts with a remote LDAP (including Active Directory) server.

- Ensure LDAP is selected as an Authentication Method on the [AAA](#) page.
- Click **System > Authentication** and switch to the **LDAP** tab.
- Define the global LDAP authentication options.
Click **Apply Changes**.
- Specify the hostname or IP address of the remote LDAP server.
IPv4 or IPv6 addresses can be specified. Multiple LDAP servers may be defined.
- Click **Add New LDAP Server**.
- To remove an LDAP servers from the Exinda appliance, select the checkbox for the server and click **Remove Server**.
- To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Radius Authentication

Radius authentication allows you to configure the Exinda appliance to authenticate user login attempts with a remote Radius server.

1. Ensure RADIUS is selected as an Authentication Method on the [AAA](#) page.
2. Click **System > Authentication** and switch to the **Radius** tab.
3. Define the global RADIUS settings.
4. Click **Apply Changes**.
5. Specify the hostname or IP address of the remote Radius server.
IPv4 addresses can be specified. Multiple Radius servers may be defined.
6. Click **Add New RADIUS Server**.
7. To remove Radius servers from the Exinda appliance, select the checkbox for the server and click **Remove Server**.
8. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

TACACS+ authentication

TACACS+ authentication allows you to configure the Exinda appliance to authenticate user login attempts with a remote TACACS+ server.

1. Ensure TACACS+ is selected as an Authentication Method on the [AAA](#) page.
2. Click **System > Authentication** and switch to the **TACACS+** tab.
3. Define global TACACS+ authentication options.
4. Click **Apply Changes**.
5. Specify the hostname or IP address of the remote TACACS+ server.
IPv4 addresses can be specified. Multiple TACACS+ servers may be defined.
6. Click **Add New TACACS+ Server**.
7. To remove TACACS+ servers from the Exinda appliance, select the checkbox for the server and click **Remove Server**.
8. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Logging

The System Logging section of the Exinda appliance System Setup allows you to view and configure the System Log files.

- ["View System Log Files" on page 327](#): View and search the System Log files.
- ["Live Log" on page 327](#): View real-time entries to the System Log file.
- ["Tail Log" on page 327](#): View the most recent entries to the System Log file.

- "System Logging Configuration" on page 327: Configure various System Logging settings, including remote syslog servers.

View 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.

1. Click **System > Logging** and switch to the **View** tab.
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 filters are examples of common filters, and 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**.

Live Log

The Live Log page allows you to view new entries to the System Log in real-time.

1. Click **System > Logging** and switch to the **Live Log** tab.
A dot/period (.) character is displayed after a few seconds of inactivity to indicate the Live Log is still active.

Tail Log

The Tail Log page allows you to view the most recent entries in the system log file.

1. Click **System > Logging** and switch to the **Tail Log** tab.
2. Configure how many lines to view and in which order to display the log entries.

View Last: Lines View Log Order:

3. To refresh this page and ensure any new log entries since the list time this page was refreshed are displayed, click **Go**.

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" on page 328
- "Add a remote syslog server" on page 328
- "Remove a remote syslog server" on page 328

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. Click **System > Logging** and switch to the **Setup** tab.
2. 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.
3. 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.
4. Select when the logs are rotated.
To force System Log rotation immediately, click **Force Rotation Now**.
5. Specify how many log files should be kept before they are permanently removed from the Exinda appliance.
6. Click **Apply Changes**.
7. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

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. Click **System > Logging** and switch to the **Setup** tab.
2. In the Add New Remote Sink area, enter the Hostname or IPv4 address of the remote syslog server.
IPv6 addresses are not supported for remote sinks.
3. 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.
4. Click **Add New Remote Sink**.
5. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Remove a remote syslog server

To stop forwarding system log entries to a remote syslog server, remove the server from the Exinda

appliance.

1. Click **System > Logging** and switch to the **Setup** tab.
2. Select the server from the Remote Log Sinks list, and select **Remove Selected**.
3. Click **Add New Remote Sink**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

System Diagnostics

System Diagnostics provide troubleshooting assistance in cases of system and network issues. The various diagnostics pages include:

- ["Notify administrators of system issues" on page 279](#): View status of system alerts
- ["Diagnostics Files" on page 331](#): Generate and download diagnostics files
- ["TCP Dump" on page 332](#): Run and download a TCP dump
- ["View the status of the community" on page 333](#): View community diagnostics
- ["Verify acceleration configuration" on page 334](#): View TCP acceleration, WAN memory, and SMB acceleration diagnostics
- ["Monitor" on page 337](#): View monitoring diagnostics
- ["Optimizer Diagnostics" on page 338](#): View optimizer diagnostics
- ["NIC Diagnostics" on page 339](#): View NIC diagnostics
- ["RAID Diagnostics" on page 339](#): View RAID adapter, logical and physical drive information

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. Click **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.
2. To view the alert that has triggered the warning, click the alarm name.
Use the information in this alert to help troubleshooting the issue.
3. 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

Alert Name	Description
	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.
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.

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"/>	 sysdump-ex240-20091120-152626.tgz	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.

TCP Dump

The following form can be used to generate a TCP Dump on the Exinda appliance. 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 TCP Dump	
Interface	ALL <input type="button" value="v"/>
Timeout	60 Seconds <input type="button" value="v"/>
Filter	<input type="text"/>
Status	Stopped

Note when ALL is selected for the Interface, only those interfaces which are link up will be included.

Interface	Select an interface to run the TCP dump on. Select ALL to capture packets on all (link up) interfaces.
Timeout	Select the time for which the TCP Dump will run.
Filter	Set a filter if required. More information on tcpdump filters is available at www.tcpdump.org
Status	Shows the status of a running TCP Dump

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"/>	capture-ex240-20091123-192334.tar.gz	Mon Nov 23 19:23:36 EST 2009	23850263 bytes
<input type="checkbox"/>	capture-ex240-20091123-194306.tar.gz	Mon Nov 23 19:44:46 EST 2009	619548490 bytes
<input type="checkbox"/>	capture-ex240-20091123-194926.tar.gz	Mon Nov 23 19:49:26 EST 2009	7023 bytes
<input type="checkbox"/>	capture-ex240-20091123-195327.tar.gz	Mon Nov 23 19:53:27 EST 2009	2917 bytes
<input type="checkbox"/>	capture-ex240-20091123-200247.tar.gz	Mon Nov 23 20:02:47 EST 2009	13135 bytes
<input type="checkbox"/>	capture-ex240-20091123-200833.tar.gz	Mon Nov 23 20:08:34 EST 2009	2412339 bytes
<input type="checkbox"/>	capture-ex240-20091124-152842.tar.gz	Tue Nov 24 15:29:03 EST 2009	217979641 bytes
<input type="checkbox"/>	capture-ex240-20091124-171922.tar.gz	Tue Nov 24 17:19:30 EST 2009	116281545 bytes
<input type="checkbox"/>	capture-ex240-20091127-135953.tar.gz	Fri Nov 27 13:59:53 EST 2009	240871 bytes
<input type="checkbox"/>	capture-ex240-20091202-135242.tar.gz	Wed Dec 02 13:52:42 EST 2009	276 bytes
<input type="checkbox"/>	capture-ex240-20091202-135345.tar.gz	Wed Dec 02 13:53:45 EST 2009	280 bytes

Note Valid SMTP and DNS settings are required for TCP Dumps to be sent to Exinda TAC.

View the status of the community

Display the state of the community and details of the individual hosts that have joined.

1. Click **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:      jb-exinda
  Version:       5.5.0.12035
```

Verify acceleration configuration

Acceleration diagnostics aid in troubleshooting TCP Acceleration, SMB Acceleration and WAN Memory issues by displaying the current configuration for those areas.

In this area of the Exinda Web UI you can:

- ["View the SMB configuration and connections" on page 334](#)
- ["View the TCP acceleration configuration and connections" on page 335](#)
- ["View the WAN memory configuration and reduction statistics" on page 336](#)

View the SMB configuration and connections

The SMB Acceleration diagnostics display the current configuration settings as well as the number of new and concurrent accelerated connections. If SMB signed connections are present, the total number of signed connections is also displayed.

Note To configure CIFS acceleration settings, navigate to **System > Optimization > SMB** on the Web UI, advanced mode.

1. Click **System > Diagnostics** and switch to the **Acceleration** tab.
2. From the Module drop-down, select **SMB Acceleration**.

The configuration settings for SMB and SMB2 are displayed. The connections statistics are broken down 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 first time an attempt to validate the domain credentials fails, the connection is identified as being signed, but is not accelerated. This attempt and 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 were accelerated.
- **Unhandled** — After a signed connection has failed validating the domain credentials the first time, and the connection is marked as **Bypassed**, all subsequent attempts to validate credentials of a signed connection against the IP address of the server are marked as **Unhandled**.

Note The statistics reported on this page are reset each time the SMB Acceleration service is restarted.

3. "[View System Log Files](#)" on page 327.

View the TCP acceleration configuration and connections

The TCP Acceleration diagnostics display the current configuration settings as well as the number of new and concurrent accelerated connections.

- Note**
- To configure TCP acceleration settings, navigate to **System > Optimization > TCP** on the Web UI, advanced mode.
 - To view the licensed limit of accelerated connections, navigate to **System > Setup > License** on the Web UI, advanced mode.

1. Click **System > Diagnostics** and switch to the **Acceleration** tab.
2. From the Module drop-down, select **TCP Acceleration**.

The configuration settings for TCP acceleration are displayed.

Configuration

Congestion Control Algorithm:	cubic
Transport Mode:	transparent
Window Scale:	5 (2M)W
Appliance Discovery Enabled:	yes
Dual Bridge Bypass Enabled:	yes
TCP-Keep-Alive Enabled:	yes
TCP-Keep-Alive Timeout:	3600
Concurrent Accelerated Connections:	3/4000
New connections per second:	0

Peak Concurrent Accelerated Connections: 127

Reduction Statistics

TX reduction:	87.53%
TX bytes in:	38.56M
TX bytes out:	4926.02k
RX reduction:	83.39%
RX bytes in:	7063.57k
RX bytes out:	41.53M

3. ["View System Log Files" on page 327.](#)

View the WAN memory configuration and reduction statistics

The WAN memory Acceleration diagnostics display the current configuration settings as well as reduction statistics for the individual hosts.

Note To configure WAN memory acceleration settings, navigate to **System > Optimization > WAN memory** on the Web UI, advanced mode.

1. Click **System > Diagnostics** and switch to the **Acceleration** tab.
2. From the Module drop-down, select **WAN Memory**.

The configuration settings for WAN memory acceleration are displayed.

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:	104.52G
Disk used:	384k (0.000%)
Persistence active:	yes
Persistence loading:	no
Persistence clear pending:	no

TX reduction:	88.36%
TX bytes in:	50.7M
TX bytes out:	6041.19k
TX small matcher bytes in:	0
TX small matcher bytes out:	0

```
RX reduction:          84.50%
RX bytes in:           8587.15k
RX bytes out:          54.12M
```

```
Peer 0010f3074208 state
```

```
Status:                ONLINE
Last status change:    2011/05/25 12:33:10 (23h 42m 27s)
IP address:            192.168.0.208
Connections:           0
Disk cache status:     Active
Version info:          0k/24 (local:23-25 remote:23-24)
```

```
TX reduction:          34.35%
TX bytes in:           1000.17k
TX bytes out:          656.6k
```

```
RX reduction:          32.22%
RX bytes in:           167.53k
RX bytes out:          247.16k
```

```
Peer 0010f3117e48 state
```

```
...
```

3. ["View System Log Files" on page 327.](#)

Monitor

The monitor diagnostics display the current monitor settings and the status of monitor and collector processes.

Note To configure Monitor settings, navigate to **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

```

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, 6121M 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 1185M 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, 46G 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 18G 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

```

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 **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
```

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]

```

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 Networks Support Services.

1. Click **System > Diagnostics > Log a Case**.
2. Complete the fields in the form.
3. Ensure a brief summary of the problem or question is included in the **Subject** field.
4. Provide a detailed description of the question or the problem you are experiencing in the **Description** field.
5. Select whether to attach a diagnostics file or monitoring report to the case.
6. 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.

Maintenance

The Maintenance section of the Exinda appliance System Setup allows you to perform various system maintenance tasks. These include:



- "[Manage System Configuration](#)" on page 341: Allows you to save, activate, switch, revert and delete system configuration files.
- "[Import System Configuration](#)" on page 342: Allows you to import previously saved or backed-up system configuration files.
- "[Cluster and High Availability](#)" on page 343: View the status of Exinda clustering.
- "[Install an update to the Exinda appliance software](#)" on page 344: Upgrade the ExOS software on the Exinda appliance.
- "[Factory Defaults](#)" on page 345: Restore the Exinda appliance to factory default settings.
- "[Reboot the Exinda appliance](#)" on page 346: Reboot or Shutdown the Exinda appliance.

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 **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"/> initial.bak		
<input type="checkbox"/> initial	<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.

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 **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"/> Upload local binary file:	<input type="text"/> <input type="button" value="Browse..."/> (To be saved as separate file with its original name)
<input type="radio"/> Upload local text file: (CLI commands)	<input type="text"/> <input type="button" value="Browse..."/> (To be executed immediately in the running configuration)

Upload local

Use this option to upload a saved binary configuration file. This file would

binary file	have been downloaded as a binary file from the System Maintenance Manage Config page. Once this file is uploaded, it will appear in the list of available configuration files on the System Maintenance Manage Config 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 System Maintenance Manage Config 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)

Cluster and High Availability

Clustering allows multiple Exinda appliances to operate as if they were a single appliance. This allows for seamless deployment into High Availability and Load Balanced environments. A typical deployment topology is illustrated below.



In this example, there are two physical links. 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 and acceleration policies, as if they were a single appliance.

Refer to the following topics for example topologies:

- ["Redundancy through multiple Exinda appliances" on page 73](#)
- ["Load balancing and fail-over with multiple Exinda appliances" on page 75](#)
- ["High availability mode" on page 77](#)

Install an update to the Exinda appliance software

Exinda software is called ExOS and is updated regularly with new product features as well as system and performance improvements. The Firmware Update screen allows you to update the software on your Exinda appliance.

Exinda appliances have two partitions for installing ExOS updates. The current, running ExOS version will be installed on one partition, which means you can install a newer ExOS update on the other, unused partition.

Note Valid Software Subscription (SS) is required to install new ExOS updates. You can view your SS expiry date at the top of this page.

1. Click **System > Maintenance** and switch to the **Firmware Update** tab.
2. Locate the new image to install.
 - If you know the URL for the image file, select **Install from URL** and type the URL.
This URL is usually published in the Release Notes as part of the ExOS release. If you click **Check for Latest Update**, and a newer ExOS software update is available, this field is populated automatically.
 - If the image file has previously been downloaded 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 file.
3. (Optional) Schedule the download and installation of the ExOS update for a later date or time.
 - a. Select the **Schedule Installation** checkbox and specify the **Date** and **Time**.
 - b. By default, the download of the image will happen straight away and only the installation will be scheduled. To schedule the download of the ExOS image to happen at the scheduled time, select **Schedule Image Download**.
 - c. By default, the Exinda appliance will not reboot following a scheduled installation. To restart the appliance after the scheduled installation, select **Reboot After Installation**.
4. Before installing or scheduling a new ExOS update, you must accept the End User License Agreement (EULA).

5. Click **Install**.

The image is installed on the appliance. This process may take a few minutes to complete.

Note If the network connection fails while retrieving the latest ExOS 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.

6. To finalize the ExOS install, you must reboot the appliance. See "[Reboot the Exinda appliance](#)" on page 346.

Return to the previously installed version of ExOS

If you have updated the version of ExOS that is 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 which means you can revert to the older ExOS version.

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

1. Click **System > Maintenance** and switch to the **Firmware Update** tab.
2. In the Current Installed Images area, click **Switch Boot Partition**.
3. To finalize the ExOS install, you must reboot the appliance. See "[Reboot the Exinda appliance](#)" on page 346.

After the Exinda appliance reboots, the partition with the older build is running on the appliance.

Factory Defaults

The Factory Defaults screen allows you to restore the Exinda appliance's configuration back to factory default settings. This includes removing any system logs, WAN Memory cache and monitoring statistics.

Note To restore Factory Defaults, navigate to **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.

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 Reboot or Shutdown without saving the changes first.

1. Click **System > Maintenance** and switch to the **Reboot / Shutdown** tab.
2. (Optional) Schedule the Exinda appliance to reboot at a specific date or time.
 - a. Select the **Schedule Reboot** checkbox.
 - b. Type the date and time that the appliance should reboot.
3. 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.
4. 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. Click **System > Maintenance** and switch to the **Reboot / Shutdown** tab.
2. In the System Watchdog area, select **Enable**.
3. Click **Apply Changes**.
4. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Shutdown the Exinda appliance

If the Exinda appliance needs to be powered off, shut it down from within the Exinda Web UI.

Caution Any unsaved configuration changes will be lost if the Exinda appliance is Reboot or Shutdown without saving the changes first.

1. Click **System > Maintenance** and switch to the **Reboot / Shutdown** tab.
2. Click **Shutdown**.

The Exinda appliance will not restart, and must be physically powered on again.

Tools

There are a set of tools installed on the Exinda appliance to assist with configuration and troubleshooting. These tools include:

- ["Ping" on page 347](#): A tool to ping a network host from the Exinda appliance.
- ["Traceroute" on page 347](#): A tool to perform a traceroute to a network host from the Exinda appliance.
- ["DNS Lookup" on page 348](#): A tool to query the configured DNS servers from the Exinda appliance.
- ["Access the Command Line Interface" on page 348](#): A tool to connect to the Exinda appliance's CLI from the Web UI.
- ["IPMI" on page 349](#): A tool to issue remote power commands to an IPMI enabled appliance.

Ping

Use the Ping Tool to test network connectivity from the Exinda appliance to other hosts on the WAN or Internet.

1. Click **System > Tools > Ping**.

```
IPv4 Host:  Ping
IPv6 Host:  Ping

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
```

2. In the **IPv4 host** or **IPv6 host** field, specify an IP address or fully qualified domain name to attempt to ping.
3. Click **Ping**.

It may take a few seconds for the ping operation to complete and display the results.

Traceroute

Use the Traceroute Tool to determine the network hops from the Exinda appliance to other hosts on the WAN or Internet.

1. Click **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

```

2. In the **Host** field, specify an IPv4 or IPv6 Address, or fully qualified domain name to attempt to traceroute.
3. Click **Traceroute**.

It may take a few seconds for the operation to complete and display the results.

DNS Lookup

Use the DNS Lookup Tool to have the Exinda appliance query it's configured DNS servers to resolve the specified domain name.

1. Click **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

```

2. In the Domain field, specify a fully qualified domain name to look up.
3. Click **Lookup**.

It may take a few seconds for the operation to complete and display the results.


Access 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



login:

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

IPMI

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](#). Select the desired action from the drop down selection, enter the IPMI authentication details for the remote appliance and click on the Do Power Action button.

Power Control Options	
Command	<input type="text" value="Get Status"/>

Remote IPMI Login Details	
IPv4 Address	<input type="text"/>
Username	<input type="text" value="admin"/>
Password	<input type="password" value="*****"/>

IPv4 Address	The IPMI IPv4 address of the remote appliance
Username	The IPMI username for the remote appliance (default: admin)

Password	The IPMI password for the remote appliance (default: exinda)
----------	--

Chapter 7: Object Definitions

Objects are used by the Exinda appliance to define various items of interest such as IP addresses, Users/Groups, Applications, VLANs, Schedules, etc. These Objects are then often used in the Monitoring reports or in the Optimizer policies. The following Objects can be defined:

- "Network Objects" on page 7: Used to define IP addresses and subnets.
- "Users and Groups" on page 11: Used to define users and groups.
- "VLAN Objects" on page 13: Used to define 802.1Q VLANs.
- "Protocol Objects" on page 14: Used to define network layer protocols.
- "Configuring application objects" on page 14: Used to define applications.
- "Schedule Objects" on page 23: Used to define times during days of the week.
- "Adaptive Response" on page 26: Used to define and manage user data transfer quota.

Network Objects

Network objects are used to represent hosts, subnets and groups for monitoring and optimization. Thus, a network object can either be a single addressable network host, or a subnet of network hosts or a combination of both. Once defined, a network object may be used throughout the Exinda appliance for monitoring, optimization, and configuration purposes.

There are 2 types of Network Objects:

- **Static**: These are user-defined Network Objects. Users must manually specify subnets and/or IP addresses to define the Network Object.
- **Dynamic**: These Network Objects are automatically created and maintained by the Exinda appliance.

Create a static network object

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, optimization, and configuration purposes.

Traffic between hosts set to 'ignore' will pass through the unit without monitoring, QoS, or compression.

The **ALL**, **private net** and **local** Network Objects are automatically created by the appliance.

- **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.

Additional network objects can be added to the Exinda appliance as needed.

1. Click **Objects > Network** and switch to the **Network Objects** tab.
2. Specify a name for the Network Object.
3. Select whether the subnet is on the LAN side of the appliance (internal) or the WAN side (external).

There are three options for the location field: Internal, External, and Inherit.

- **Internal** — All subnets and hosts defined by the Network Object that exist on the LAN side of the appliance. Subnets defined as Internal when the network object was created on the Exinda appliance.

When the **Ignore Internal-to-Internal** option is set on the Monitoring page, all traffic between Network Objects marked as Internal is ignored and passed through the Exinda appliance. See "[Monitoring Configuration](#)" on page 216.

- **External** — All subnets and hosts defined by the Network Object that exist on the WAN side of the appliance. Subnets defined as External when the network object was created on the Exinda appliance.
- **Inherit** — The subnet and hosts location is determined by closest match to other Network Objects. If no Network Objects match then the location defaults to external.

Note

When creating Network Objects that have location set to "Inherit", use the CLI command `show network-object` to show the location.

- If all subnets in the Network Object are contained within another Network Object that is internal, the location will be internal.
- If all subnets match a Network Object that is external, the location will be external.
- If some subnets match a Network Object that is internal, and some match a Network Object that is external, the location will be shown as mixed.

Packets are matched to a Network Object, and the closest subnet (see Example below) within that Network Object determines the location.

4. Select whether the subnet is included in the Subnet reports. See "[Subnets Report](#)" on page 87.
5. 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.

6. Click **Add new Network Object**.
7. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

Example 1

Create a Network Object that defines two internal proxy servers, 192.168.1.10 and 192.168.1.11:

```
Name: Web Proxies
Location: Internal
Subnet Report: Yes

Subnets: 192.168.1.10 /32
Subnets: 192.168.1.11 /32
```

Example 2

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
Subnet Report: Yes

Subnets: 10.0.100.0 /24
```

Example 3

Create a Network Object that defines the internal IPv6 server at 2001:db8::1234:5678

```
Name: FileServer6
Location: Internal
Subnet Report: Yes

Subnets: 2001:db8::1234:5678 /128
```

Example 4

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: User-1 Subnets: 10.0.1.200/32 Location: Inherit
```

Subnets are matched by decreasing netmask length. The host 10.0.1.200 will be internal, as it most closely matches the Office-A Network Object which is internal. Since the User-1 Network Object contains a single subnet that can be matched to Office-A, its location is shown as internal.

```
(config) # show network-object User-1
Network Object: User-1
Location:      internal (inherited)
Subnet Report: no
Subsystem:    static
```

```
Subnets:
 10.0.1.200/32
```

Restrict access to management services by IP address

When there are only specific IP addresses that should have access to management services, remove the ability for all other IP addresses from accessing them. Limiting access can help eliminate denial of service attacks and attempted password guesses from random IP addresses by restricting access to known addresses from which management will be performed.

1. [Create a network object](#) that includes only the IP addresses that are allowed access to the management services.
2. Click **Tools > Console**.
3. Type the appliance username and password at the prompts.
4. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

5. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

6. To restrict access to the management services, at the prompt type the following commands for the network object created in step 1:

```
web http restrict <network-object>
web https restrict <network-object>
ssh server restrict <network-object>
snmp-server restrict <network-object>
telnet-server restrict <network-object>
```

7. Save the running configuration to the active configuration file.

```
write memory
```

Dynamic 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. This page allows you to 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.

Currently there are two types of Dynamic Network Objects.

1. 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 further information, see the [Adaptive Response](#) page.

2. 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 [System > Network > Active Directory](#) page and the [Objects > Users and Groups](#) page.

Users and Groups

Users and Groups Objects are used to define pre-populated users and groups such that they can be used for monitoring and optimization.

Currently, 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](#)" on page 11.
- To define groups as Dynamic Network Objects, see "[Create Network Group Objects](#)" on page 12.

Create Network User Objects

Network Users displays a pre-populated list of Users (and their associated IP addresses) from either the Exinda Active Directory 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, it 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	


To define a user as a Dynamic Network Object

1. In the Exinda WebUI, go to **Objects > Users & Groups > Network Users**.
2. Select the checkbox for the user.
3. Click **Add Network Object**.

The Network Status icon for the user changes to , indicating it is a network object.

To stop identifying a user as a dynamic network object

1. Select the checkbox for the user.
2. Click **Remove Network Object**.

The Network Status icon for the user changes to , indicating it is no longer a network object.

Create Network Group Objects

Network Groups displays a pre-populated list of Groups from either the Exinda Active Directory 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.

To define a group as a Dynamic Network Object

1. In the Exinda WebUI, go to **Objects > Users & Groups > Network Groups**.
2. Locate 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 , indicating it is 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 stop identifying a group as a Dynamic Network Object

1. Locate the group in the list, and click **Delete**.

The Network Status icon for the user changes to , indicating it is no longer a network object.

If the dynamic network object was created from multiple groups, each group is again listed individually in the list.

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.

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 by using the form at the top of the page.

Add New VLAN

Name:

Type:

Details:

VLAN ID (0-4094)	VLAN Priority (0-7)
<input style="width: 40px;" type="text"/> - <input style="width: 40px;" type="text"/>	<input style="width: 40px;" type="text"/> - <input style="width: 40px;" type="text"/>

Add New VLAN

Name	Specify a meaningful name for the VLAN Object.
Type	Specify the type of VLAN to define. Currently only 802.1Q VLANs are available.
VLAN ID	Specify the range of VLAN IDs to define. To define all VLAN IDs, leave this field blank or enter 0 - 4094. A single VLAN ID can be defined by entering the same value in both fields.
VLAN Priority	Specify the VLAN Priority range to define. To define all VLAN Priorities, leave this field blank or enter 0 - 7. A single VLAN Priority can be defined by entering the same value in both fields.

Example

If VoIP traffic has a VLAN ID of 10, you'll need to 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)
```

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 be defined using the form at the top of the page.

Add New Protocol

Name:

Number:

Name	Give the protocol a meaningful name.
Number	Specify the IPv4 protocol number.

Note Protocol numbers are unique and can only be defined once.

All the defined Protocol Objects are available to view in the table on this page. Each Protocol Object can be Edited or Deleted by clicking the appropriate button in the table. Some Protocol Objects are protected and cannot be edited or deleted.

Example

SCTP (Stream Control Transport Protocol) is undefined by default and needs to be defined in the Exinda appliance.

```
Name: SCTP
Number: 132
```

Configuring application objects

Application Objects are used to represent applications for monitoring and optimization. They are used to

define applications that run over the network and are made up of TCP/UDP port numbers/ranges and layer 7 signatures. A Network Object can also be specified which can be used to tie an Application to a specific IP address/subnet. Application Objects can also be grouped together to form Application Groups.

Add a new application

Applications are used to define applications that run over the network, and are made up of TCP/UDP port numbers and ranges and layer 7 signatures. For example the HTTP application will match any traffic over TCP port 80 OR any traffic that matches the inbuilt HTTP Layer 7 signature.

A Network Object can also be specified to tie an Application to a specific IP address/subnet. Network Objects can only be combined with TCP/UDP port numbers/ranges, not with Layer 7 signatures, and use AND in the query to create a match. An *external* Network Object is matched as external, and an *internal* Network Object is matched as internal. If a TCP/UDP port number/range is specified, it is always matched against the host matching the network object.

There are hundreds of predefined Applications built into the Exinda appliance. Add any applications you want to monitor that are not in the list.

1. Click **Objects > Applications**.
2. On the **Applications** tab, in the Add New Application area type a name for the new application.
3. Select the **Network Object** or **L7 Signature** for the application.

Network Objects and Layer 7 Signatures are mutually exclusive. Only one can be selected.

citrix	application	Allows you to define an Application Object based on a published Citrix application name.
	priority ^	Allows you to define an Application Object based on a published Citrix priority. Citrix priorities are 0=High, 1=Medium, 2=Low, 3=Background. The Citrix priority detection will only work if Citrix is running without session-reliability, over TCP port 1494.
	user	Allows you to define an Application Object based on the user running the Citrix published application.
flash	host	Allows you to define an Application Object based on the 'host' field in the HTTP header (where flash is running over http).
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"> ■ A string literal is enclosed in quotes ("). ■ Internal quotes can be escaped with the backslash (\") character. ■ A backslash can be included in the string by escaping it with another backslash (\\). ■ Keywords are bare (<code>common_name</code>) with no quotes. ■ Keywords are bare (<code>host</code>) with no quotes. ■ Grouping is supporting using parenthesis ■ Operators supported are <code>or</code> and <code>and</code> <code>and</code> has higher precedence than <code>or</code> ■ The comparison operators that are available are: <table border="1" data-bbox="735 951 1354 1734"> <thead> <tr> <th data-bbox="735 951 964 999">Description</th> <th data-bbox="964 951 1146 999">Syntax</th> <th data-bbox="1146 951 1354 999">Example</th> </tr> </thead> <tbody> <tr> <td data-bbox="735 999 964 1104">equals</td> <td data-bbox="964 999 1146 1104"><code>= host</code></td> <td data-bbox="1146 999 1354 1104"><code>= "example.com"</code></td> </tr> <tr> <td data-bbox="735 1104 964 1209">does not equal</td> <td data-bbox="964 1104 1146 1209"><code>!= host</code></td> <td data-bbox="1146 1104 1354 1209"><code>!= "example.com"</code></td> </tr> <tr> <td data-bbox="735 1209 964 1314">contains substring</td> <td data-bbox="964 1209 1146 1314"><code>=% host</code></td> <td data-bbox="1146 1209 1354 1314"><code>=% "example.com"</code></td> </tr> <tr> <td data-bbox="735 1314 964 1419">does not contain substring</td> <td data-bbox="964 1314 1146 1419"><code>!% host</code></td> <td data-bbox="1146 1314 1354 1419"><code>!% "example.com"</code></td> </tr> <tr> <td data-bbox="735 1419 964 1608">Right side is a regular expression and it matches the full left side</td> <td data-bbox="964 1419 1146 1608"><code>=~ host</code></td> <td data-bbox="1146 1419 1354 1608"><code>=~ "example.*"</code></td> </tr> <tr> <td data-bbox="735 1608 964 1734">Right side is a regular expression and it</td> <td data-bbox="964 1608 1146 1734"><code>!~ host</code></td> <td data-bbox="1146 1608 1354 1734"><code>!~ "example.*"</code></td> </tr> </tbody> </table>	Description	Syntax	Example	equals	<code>= host</code>	<code>= "example.com"</code>	does not equal	<code>!= host</code>	<code>!= "example.com"</code>	contains substring	<code>=% host</code>	<code>=% "example.com"</code>	does not contain substring	<code>!% host</code>	<code>!% "example.com"</code>	Right side is a regular expression and it matches the full left side	<code>=~ host</code>	<code>=~ "example.*"</code>	Right side is a regular expression and it	<code>!~ host</code>	<code>!~ "example.*"</code>
Description	Syntax	Example																					
equals	<code>= host</code>	<code>= "example.com"</code>																					
does not equal	<code>!= host</code>	<code>!= "example.com"</code>																					
contains substring	<code>=% host</code>	<code>=% "example.com"</code>																					
does not contain substring	<code>!% host</code>	<code>!% "example.com"</code>																					
Right side is a regular expression and it matches the full left side	<code>=~ host</code>	<code>=~ "example.*"</code>																					
Right side is a regular expression and it	<code>!~ host</code>	<code>!~ "example.*"</code>																					

		<table border="1"> <thead> <tr> <th>Description</th> <th>Syntax</th> <th>Example</th> </tr> </thead> <tbody> <tr> <td>does not match the full left side</td> <td></td> <td></td> </tr> </tbody> </table> <ul style="list-style-type: none"> Regular expressions use the perl syntax The keywords for HTTP are: host, file, user_agent, content_type, method, content_len and encoding <p>Examples:</p> <ul style="list-style-type: none"> (url =% "index" or file =% "login") and host =% "example.org" and content_type.case = "MyContentType" (host =% "facebook.com" and file !=% "cgi-bin/abcd") or host =% "facebook2.com" 	Description	Syntax	Example	does not match the full left side				
Description	Syntax	Example								
does not match the full left side										
ssl	common_name	Allows you to define an Application Object based on the 'common name' field in the SSL certificate.								
	content_type	Allows you to define an Application Object based on the 'content-type' field in the HTTP header.								
	advanced	<p>Define custom criteria with the following syntax:</p> <ul style="list-style-type: none"> A string literal is enclosed in quotes ("). Internal quotes can be escaped with the backslash (\") character. A backslash can be included in the string by escaping it with another backslash (\\). Keywords are bare (common_name) with no quotes. Grouping is supporting using parenthesis Operators supported are or and and and has higher precedence than or The keywords for SSL are common_name (cn) and organization_name (o) The comparison operators that are available are: <table border="1"> <thead> <tr> <th>Description</th> <th>Syntax</th> <th>Example</th> </tr> </thead> <tbody> <tr> <td>equals</td> <td>= common_name</td> <td>= "John"</td> </tr> <tr> <td>does not equal</td> <td>!= common_name</td> <td>!= "John"</td> </tr> </tbody> </table>	Description	Syntax	Example	equals	= common_name	= "John"	does not equal	!= common_name
Description	Syntax	Example								
equals	= common_name	= "John"								
does not equal	!= common_name	!= "John"								

		Description	Syntax	Example
		contains substring	<code>=% common_name</code>	<code>=% "John"</code>
		does not contain substring	<code>!% common_name</code>	<code>!% "John"</code>
		Right side is a regular expression and it matches the full left side	<code>=~ common_name</code>	<code>=~ "John*"</code>
		Right side is a regular expression and it does not match the full left side	<code>!~ common_name</code>	<code>!~ "John*"</code>
<ul style="list-style-type: none"> Regular expressions use the perl syntax 				
	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).		

4. 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.

For 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**.

Add New Application

Name:

Network Object:

L7 Signature:

Ports/Protocols: eg. 80,8080,3127-3128

[Show a List of Common Port Numbers](#)

- 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.

Note

- Ports, Port Ranges and L7 Signatures are OR'd together, that is, traffic only has to match one of the conditions for it to be considered a match. Take the HTTP example above. Traffic needs to be on TCP port 80 OR match the 'http' L7 signature for the appliance to classify it as HTTP. If a Network Object is specified, it is AND'd with any specified ports or port ranges.
- TCP and UDP port pairs can only be defined once. So if you define an Application Object with a port range TCP 500 -> 510, you cannot then define another Application Object on TCP port 505. You can however, define UDP port 505 as TCP and UDP are treated separately. You can also define duplicate ports/port ranges if a Network Object is also specified.

- Click **Add New Application**.

Application sub-types


Sub-type classification takes reporting and Layer 7 visibility to a whole new level of granularity. Just like reporting on specific web applications, where most vendors can only report on port 80 traffic, Exinda allows a deeper look into Layer 7 applications.

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
- Exinda s Layer 7 with Sub-type classification report on web applications as: Yahoo video, Yahoo voice, or Yahoo webchat. Similarly it breaks down other applications into their piece parts.

The following two graphics show the simplified configuration screen used by the Exinda to identify an applications sub-type and, following, what a report on a sub-type classification may look like (notice Skype video and Skype voice broken out as separate reportable entries).

The screenshot shows the 'Add New Application' configuration interface. The 'Name' field contains 'Yahoo Voice'. The 'Network Object' field is empty. The 'L7 Signatures' field is set to 'yahoo --->' and has a dropdown menu open with the following options: 'file-transfer', 'unknown', 'video', 'voice' (which is highlighted in blue), and 'webchat'. The 'Ports/Protocols' field is set to 'Protocol'. Below the 'L7 Signatures' field, there is a link that says 'Show a List of Common'. At the bottom left of the form, there is a button labeled 'Add New Application'.



Add an application group

Application Groups are used to group together Individual Application Objects into a logical group. There are several predefined Application Groups, such as Mail, P2P, Voice, etc. Using this page, you can edit existing Application Groups or create new ones.

1. Click **Objects > Applications**.
2. On the **Application Groups** tab, 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, then select the Edit button next to the newly created Application Group. You will be presented with four additional drop-downs to add more Application Objects.

4. Click **Add New Application Group**.

Once defined, Application Objects are automatically used in various monitoring reports and can also be used in the Optimizer policies.

Add an application to an application group

Application Groups allow you to group one or more defined Applications for easy inclusion into Optimizer policies. Add any applications you have added to the Exinda appliance `

Caution For monitoring visibility, groups must not contain applications which are already a member of another group being monitored.

1. Click **Objects > Applications**.
2. On the **Application Groups** tab, locate the group to add applications to and click **Edit**.
3. In the Edit Application Group area, select the applications that belong in the group.
4. Click **Apply Changes**.

Enable anonymous proxy classification

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 can't easily be traced back to their original IP.

The Anonymous Proxy Application is a special Application Object that is used to detect Anonymous Proxy websites and services. 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 be created. 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 **System > Setup > Monitoring** page.

1. Click **Objects > Applications** and switch to the **Anonymous Proxy** tab.
2. Click **Enable**.

As the appliances communicates with the Exinda web servers, the status of the anonymous proxy data is updated.

The 'Renumerate' button can be used to force the Anonymous Proxy Service to fetch the Anonymous Proxy definitions immediately.

Note Valid Software Subscription (SS) is required in order to fetch the Anonymous Proxy list from the Exinda website. The Exinda appliance will also require access to the Exinda web servers to fetch new Anonymous Proxy definitions either via a direct Internet connection or via a HTTP Proxy.

Enable Anonymous Proxy Detection

The Anonymous Proxy service is disabled by default. In order to enable this feature, navigate to **Objects > Applications > Anonymous Proxy** using the Web UI, Advanced mode.

Service: **Running**

Settings	
URL	http://www.exinda.com/ap/apdata.tar.gz
Last Check	2009/12/03 09:32:08 (57m 45s ago)
Last Update	2009/12/03 09:32:11 (57m 42s ago)
Status	Ok

The **renumerate** button refreshes the Anonymous Proxy list immediately

Figure 1: The form to configure and enable the Anonymous Proxy service.

This page allows you to start the Anonymous Proxy service and also check when the definitions were last updated. The 'Renumerate' button allows you to force the Anonymous Proxy service to fetch the latest definitions immediately.

Given that Anonymous Proxies are constantly changing, the Anonymous Proxy service will automatically retrieve the latest Anonymous Proxy definitions from the Exinda servers on a daily basis. If the Anonymous proxy service is stopped or disabled the last retrieved definitions will be used for detection of Anonymous proxy.

Note In order to receive daily Anonymous Proxy definition updates, the Exinda appliance must be able to contact the www.exinda.com web servers and the appliance must also have valid software subscription.

The Anonymous proxy ASAM is another component of the Anonymous Proxy detection. This works in combination with the Anonymous Proxy service and it is enabled by default.

To disable this ASAM, navigate to **System > Setup > Monitoring** using the Web UI - Advanced mode. If the service is stopped and Anonymous proxy detection is no longer required, disabling the ASAM will clear the existing definitions.

ASAM	
HTTP	<input checked="" type="checkbox"/> Enable
Citrix	<input checked="" type="checkbox"/> Enable
Anonymous Proxy	<input checked="" type="checkbox"/> Enable

Figure 2: The form to enable/disable the Anonymous Proxy ASAM.

Control Anonymous Proxy Traffic

Once the Exinda appliance identifies traffic as an Anonymous Proxy, it is classified as the "Anonymous Proxy" application. This means that any Anonymous Proxy traffic will show up in the real-time monitoring screen and other monitoring reports as "Anonymous Proxy".

Inbound Applications					Outbound Applications				
Application Name	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)	Application Name	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)
Total	177.348	52	129		Total	78.713	57	131	
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>
Anonymous Proxy	2.184	2	7	<div style="width: 7%;"></div>	Anonymous Proxy	2.745	2	7	<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>

Figure 3: The Anonymous Proxy application is shown on the real-time monitoring screen.

It is also possible to create Optimizer Policies using the Anonymous Proxy application, like you would any other application. The Optimizer Policy configuration form below shows how to create an Optimizer Policy that will block Anonymous Proxies.

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="< - >"/>	<input type="text"/>	<input type="text"/>	Anonymous Proxy
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	

Figure 4: Blocking Anonymous Proxies using the Optimizer.

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 Proxies, the discard policy must be above any policy that references the Recreational application group.

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, whilst still allowing automated backup services and email to function.

By default, there are 3 Schedule Objects defined.

Name	From Day	To Day	From Time	To Time	Edit	Delete
After Hours					<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"/>
ALWAYS						
	Sunday	Saturday	0:00	24:00		
Work Hours					<input type="button" value="Edit"/>	<input type="button" value="Delete"/>
	Monday	Friday	08:00	18:00		<input type="button" value="Delete"/>

ALWAYS	This Schedule Object defines an 'Always On' schedule. This schedule is active 24 hours a day, 7 days a week. This Schedule is not editable and cannot be deleted.
After Hours	This Schedule Object defines a typical after hours schedule. It is active all day on Saturday and Sunday and from 6pm to 8am on Mondays to Fridays.
Work Hours	This Schedule Object defines a typical working hours schedule. It is active from 8am to 6pm on Mondays to Fridays.

Additional Schedule Objects can easily be added by using the form at the top of the page.

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"/>
<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>
<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>
<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>	<input style="width: 80%;" type="text"/>

Name	Give the Schedule Object a meaningful name.
Times	Specify one or more times this Schedule Object is to be active for.
From Day	Specify a day of the week this Schedule Object should begin.
To Day	Specify a day of the week this Schedule Object should end.

From Time	Specify a time this Schedule Object should begin.
To Time	Specify a time this Schedule Object should end.

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.
- By default, there are 4 drop-downs 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 4 additional drop-downs to add more schedule times.

Adaptive Response Rules

Adaptive Response allows administrators to specify rules based on data transfer which dynamically populate Network Objects. These Dynamic Network Objects may then be used when configuring Optimizer Policies.

This functionality allows the system administrator to create policies which automatically restrict a user's bandwidth once a set transfer limit has been exceeded within a specified period of time. Users are identified by IP address.

The following steps are required to implement such policies:

1. Create a static Network Object that defines the subnet(s) that will be monitored. See ["Create a Source Network Object" on page 26](#).
- OR
- Map an Active Directory group to a Dynamic Network Object.
2. ["Create an Adaptive Response limit rule" on page 27](#)
 3. Include the Dynamic Network Object in the Optimizer Policies. See ["Use Adaptive Response with Active Directory" on page 29](#)

To demonstrate how to configure Adaptive Response using the Web UI, the following example will be used as a guide.

Example

An educational institution has a group of students who have IP addresses in the subnet 192.168.0.0/16. Each student shall be allowed 10 GB data transfer (uploads and downloads) per month.

Adaptive Response

Adaptive Response Objects are used to define and manage user data transfer quotas. You can use Adaptive Response Objects to create a limit on how much data users transfer per day, week or month. Use the form at the top of the page to create a new Adaptive Response Object.

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"/>
Amount (MB):	<input type="text"/>
Enable:	<input type="text" value="No"/>

Add New Limit

Name	The name of the Adaptive Response object
Source Network Object	Specify a source Network Object to use as a list of users for whom to apply the quota. This can be a Static Network Object (such as a subnet) or a Dynamic Network Object (such as an Active Directory group).
Destination Network Object	Specify a name for the Dynamic Network Object that will be created, which will hold the list of users that have exceeded their quota.
Duration	Specify the duration to use when accounting the quota. Daily, weekly or monthly.
Direction	Specify which direction should be used when accounting the quota. Inbound, outbound or both.
Amount	Specify the quota amount (in MB) for this rule.
Enable	Specify if this rule should be enabled or not.

Any users from the 'Source Network Object' who have exceeded their quota within the 'Duration' period, will be placed into the 'Destination Network Object', which can then be used in the Optimizer policies. To view a list of the users who have exceeded their quota, see the [Objects > Network > Dynamic](#) page.

Create a Source Network Object

Create a Network Object that defines the Student subnet as 192.168.0.0/16.

1. Click **Object > Network > Network Objects**.
2. In the Add New Network Object area, type a name for the object.
3. Select whether the subnet is on the LAN side of the appliance (internal) or the WAN side (external). Packets are matched to a Network Object, and the closest subnet within that Network Object determines the location.

There are 3 options for the location field: Inherit, Internal, and External.

- Internal means all subnets/hosts defined by this Network Object exist on the LAN side of the appliance.
 - External means all subnets/hosts defined by this Network Object exist on the WAN side of the appliance.
 - Inherit means that a subnet/hosts location is determined by closest match to other Network Objects. If no Network Objects match then the location defaults to external.
4. To include traffic matching this network object in the Subnets Report, select the **Subnet Report** checkbox.
 5. Type the network IP addresses and netmask length of the subnet in the fields.
IPv4 and IPv6 addresses are accepted.
 6. Click **Add New Network Object**.

Note When creating or editing a network object, you will be presented with 4 input lines. To add more than 4 objects, you need to save and then re-edit to be presented with an extra 4 lines.

Create an Adaptive Response limit rule

Adaptive Response Limits are rules which are used to create and populate network objects based on amount of data transferred. These dynamic network objects may then be used when creating virtual circuits or filters. For example, create a rule that ensures that any user in the Students Network Object gets placed in the Students-Over-Quota Dynamic Network Object once they have transferred (uploaded and downloaded) more than 10 GB in a calendar month.

At the end of the calendar month, the Students-Over-Quota Network Object is reset.

Add New AR Limit	
Name:	<input type="text" value="Students-AR"/>
Source Network Object:	<input type="text" value="Students"/>
Destination Network Object:	<input type="text" value="Students-Over-Quota"/>
Duration:	<input type="text" value="Monthly"/>
Direction:	<input type="text" value="Both"/>
Amount (MB):	<input type="text" value="10000"/>
Enable:	<input type="text" value="Yes"/>

1. Click **Object > Adaptive Response**.
2. Type a name for the new limit.
3. Specify a **Source Network Object** to use as a list of users for whom to apply the quota.
This can be a Static Network Object (such as a subnet) or a Dynamic Network Object (such as an Active Directory group).
4. Specify a name for the **Dynamic Network Object**, which holds the list of users that have exceeded their quota.
5. Specify the duration to use when accounting the quota: **daily**, **weekly**, or **monthly**.
6. Specify which direction should be used when accounting the quota: **inbound**, **outbound**, or **both**.
7. Specify the quota amount (in MB) for this rule.
8. To enable the rule, select **Yes**.
When a rule is disabled all IPs will be removed from the Destination Network Object.
9. Click **Add New Limit**.

Use the Adaptive Response Rule in the Optimizer

Add the new Dynamic Network Object to the Optimizer Policies using the Advanced Web UI and navigated to the Optimizer page.

Circuit 10 - Students Network (10240 kbps)

Virtual Circuit 10 - Students Over Quota (512 kbps to / from 'Students-Over-Quota')

ALL - Guarantee Low 5%-100% (Optimize 5% - 100%, Priority 7)

Order: **Policy:**

[Create New Policy...](#)

Virtual Circuit 20 - Other Students (10240 kbps to / from 'ALL')

ALL - Guarantee Med 8%-100% (Optimize 8% - 100%, Priority 5)

Order: **Policy:**

[Create New Policy...](#)

[Create New Virtual Circuit...](#)

[Create New Circuit...](#)

Figure 5: The Students-Over-quota Dynamic Network Object used in an Optimizer Virtual Circuit.

In this example, the Students that have exceeded their monthly limit get placed in a 512 kbps Virtual Circuit whereas all other students (the ones who have not exceeded their monthly limit) are placed in a 10Mbps Virtual Circuit.

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. Click **Objects > Users & Groups**.
2. Beside the "Students (DEV)" group click **Edit**.

Welcome to **exinda**, logged in as **admin** (advanced, switch to basic). Logout

Optimizer Status : On (Restart / Stop) | Config Status No unsaved changes | System Health : OK | Thu Apr

Users & Groups

Network Users | **Network Groups**

Network Groups (Total: 12)

[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]

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"/>
Guests (DEV)	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
Remote desktop users (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"/>

3. Select the **Map to Network Object** and **Ignore Domain** checkboxes.
4. Click **Apply**.

A Network Object named 'Students' 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.

Create 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
```

Add 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
```

Disable 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
```

Exclude 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 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
```

Service Levels

The Service Levels section allows you to configure parameters and alerts for Application Performance Score, Application Performance Metrics and Service Level Agreement Objects.

- [Service Level Agreements](#)
- "Create an Application Performance Score object" on page 113
- "Create an Application Performance Metric object" on page 106

Service Level Agreements

Service Level Agreement (SLA) Objects can be configured on this page.

Note To configure SLA Sites, navigate to **Objects > Service Levels > Service Level Agreements** on the Web UI, advanced mode.

The table at the top of this page lists the currently defined SLA Objects, from here, you can edit/delete them. To configure new SLA Objects, click on the 'Add New SLA Object...' link.

Add New SLA Site

Name:

Destination IP:

Latency Threshold (ms):

Ping Size:

Duration:
(Duration for which the threshold is exceeded)

Enable:

Add New SLA Site

Cancel

Name	Specify a meaningful name for the SLA Site.
Destination IP	Specify an IP Address to use as the host to ping.
Latency Threshold	Specify a Latency Threshold. An alert will be triggered when the latency for the SLA Site exceeds the Latency Threshold for longer then the Duration. The default is 500 milliseconds.
Ping Size	Specify the ping packet size to use. The default is 64 bytes.
Duration	Specify a Duration. An alert will be triggered when the latency for the SLA Site exceeds the Latency Threshold for longer then the Duration. The default is one hour.
Enable	Make this SLA Site active by clicking 'Enable'.

Note Email alerts are sent when the specified threshold is exceeded for the set duration.

Valid SMTP and email settings are required for email alerts. To configure, see ["Add an SMTP server for sending email notifications"](#) on page 281.

Chapter 8: Monitoring and Reporting

After installing and configuring your Exinda appliance you can start monitoring your network. You will have full visibility into the applications that users are accessing and the amount of inbound and outbound throughput that they reach. It's recommended that you monitor your network for an adequate period before customizing Optimizer policies.

Set the Time Period Reflected in the Report

The data displayed in the reports can be focused on specific periods of time. Date ranges are available on all reports except the Real Time reports.

1. Select a report from the Monitor list.
2. 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

3. 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.

Range: 12:00AM 25/Oct/2010 - 12:00AM 26/Oct/2010

After the date range is select, the graphs and charts are immediately updated.

Data Granularity

The Exinda appliance stores data for the following amount of time:

- 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

Interactive Reports

The time graphs allow you to zoom in to a custom time range. Drag your mouse over the top of the graph and select the desired time range. To return to the initial time range click the magnifier glass icon.




Any tabular data displayed below these interactive graphs will automatically be updated with the data for the selected time range.

Note The interactive feature is only applicable to flash graphs. To change the graph display option navigate to **System > Setup > Monitoring** on the Web UI, advanced mode.

Printable Reports

All 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-hand corner 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.

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 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 to the reports displayed on the Web UI.

Real Time Monitoring

The Real Time Monitoring Report displays a breakdown of the traffic that passed through your monitored links during the last 10-60 seconds. There are several tabs for viewing the different types of network traffic in real time. These include inbound and outbound applications, hosts, conversations, and reduction by application. Real time monitoring also allows you to filter the view by IP address or subnet, policy, and active directory user.

There are six Real Time Reports available:

- [Inbound and Outbound Applications](#)
- [Internal and External Hosts/Users](#)
- [Inbound and Outbound Conversations](#)
- [Reduction](#)
- [Application Response](#)
- [Host Health](#)

It is useful during times of network congestion to know exactly what is going through the link just as it happens. Use these reports to see which internal and external hosts and/or applications are currently using the link the most and at what speeds they are doing so.

Applications

The Real Time Applications Report shows a breakdown of the Applications monitored by the Exinda appliance during the last 10 seconds. Applications are divided into Inbound and Outbound directions. The menu at the top of the report allows you to adjust the report auto-refresh rate (every 30 seconds by default).

Applications are sorted by throughput. You can also see the packet rate and number of flows for each Application. The Distribution shows the percentage of throughput an Application consumed relative to all the other Applications.

Inbound Applications				
Application Name	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)
Total	202.846	119	183	
HTTP	146.626	38	73	
HTTPS	30.860	30	18	
SMTP	18.102	42	15	
ICMP	3.216	5	36	
Skype	1.846	3	26	
Twitter	1.645	1	1	
Unclassified	0.204	0	9	
IKE	0.163	0	1	
ExindaCom	0.104	0	3	
IMAP-SSL	0.080	0	1	

Outbound Applications				
Application Name	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)
Total	824.998	146	185	
HTTPS	486.747	52	18	
SMTP	217.695	32	15	
HTTP	109.099	47	73	
ICMP	4.987	8	36	
Skype	3.326	4	26	
Twitter	1.301	1	1	
Unclassified	0.885	1	9	
IKE	0.375	0	1	
ExindaCom	0.198	0	3	
Print	0.150	0	1	
Other	0.236	0	2	

Real-time Traffic by Hosts

The Real Time Hosts/Users Report shows a breakdown of the Hosts/Users monitored by the Exinda appliance during the last 10 seconds. Hosts/Users are divided into Internal and External Hosts/Users.

Auto-Refresh Rate: | Show Users

1. Click **Monitor > Real Time > Hosts/Users**.

Hosts/Users are sorted by throughput, and display the packet rate and number of flows for each Host/User. The Distribution column shows the percentage of throughput a Host/User consumed relative to all the other Hosts/Users.

Inbound Hosts/Users				
IP Address (User)	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)
Total	138.037	46	117	
172.16.0.246 (Ksiakou)	105.324	10	5	
172.16.0.134 (Pforto)	13.909	3	4	
172.16.1.70 (Selfservice)	6.639	18	3	
172.16.1.240	3.771	6	34	
172.16.0.211	3.554	3	12	
172.16.0.244 (Cniko)	1.295	2	15	
172.16.0.127 (Sshannon)	1.060	2	20	
172.16.1.74	0.684	0	1	
172.16.0.239 (Jbothe)	0.593	1	5	
172.16.0.63 (Lenehan)	0.493	0	1	
Other	0.715	2	9	

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To display the user name associated with an internal IP, select **Show Users**.









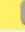

Note Active Directory must be configured on the Exinda appliances before user names can be displayed in reports. See "[Integrate the Exinda Appliance with Active Directory](#)" on page 165.

View real-time inbound and outbound conversations







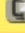



The Real-time Conversations Report shows a breakdown of the Conversations monitored by the Exinda appliance during the last 10 seconds. Conversations are divided into Inbound and Outbound directions.

- Click **Monitor > Real Time > Conversations**.

By default, the Real-time Conversations Report looks like the example below. Conversations are sorted by throughput. You can also see the packet rate and number of flows for each Conversation. Any extra information about a Conversation (a URL for example) will be shown in square brackets next to the Application.

Inbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			1408.428	284	24	
 	192.168.10.1	192.168.10.128	MAPI	570.834	82	1
 	192.168.10.9	192.168.10.128	MAPI	483.247	54	2
 	192.168.10.7	192.168.10.128	MAPI	275.334	92	2
 	192.168.10.10	192.168.10.128	MAPI	65.153	51	2
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	5.496	1	1
	192.168.10.9	192.168.10.128	LDAP	2.939	1	1
	10.20.4.1	239.255.255.250	udp ports 62612 -> 3702	1.097	0	1
	10.20.4.1	239.255.255.250	udp ports 62610 -> 3702	1.069	0	1
	192.168.10.1	192.168.0.1	NetBIOS	0.623	1	1
	192.168.10.10	192.168.10.128	LDAP	0.556	0	2
	192.168.10.132	255.255.255.255	DHCP	0.541	0	1
	192.168.10.9	192.168.0.1	NetBIOS	0.225	0	1
	10.20.3.118	10.20.255.255	NetBIOS	0.225	0	1
	192.168.10.9	192.168.255.255	NetBIOS	0.225	0	1
	10.20.11.100	224.0.0.252	udp ports 58633 -> 5355	0.212	0	1
	10.20.0.14	10.20.255.255	NetBIOS	0.193	0	1
	192.168.10.9	192.168.10.128	LDAP	0.174	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.106	0	1
	192.168.10.9	192.168.0.1	DNS	0.102	0	1
	10.20.0.181	10.20.255.255	NetBIOS	0.075	0	1

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To view only a specific IP address or subnet, type the address in the **IP/Subnet Filter** field.
The report can be filtered by IPv4 or IPv6 addresses.
- To display the optimization policy the conversation falls into, select **Show Policies**.






Outbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			3594.527	412	14	
 	192.168.10.7	192.168.10.128	MAPI	1826.409	196	2
 	192.168.10.10	192.168.10.128	MAPI	1184.445	125	2
 	192.168.10.1	192.168.10.128	MAPI	564.195	72	1
 	192.168.10.9	192.168.10.128	MAPI	12.200	17	2
	192.168.10.9	192.168.10.128	LDAP	3.316	1	1
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	2.902	1	1
	192.168.10.10	192.168.10.128	LDAP	0.565	0	2
	192.168.10.9	192.168.0.1	DNS	0.197	0	1
	192.168.10.9	192.168.10.128	LDAP	0.188	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.109	0	1

- To display the user name associated with an internal IP, select **Show Users**.

- To group individual connections within a flow as a single line item or show each connection as a separate line item, select **Group**.

Reduction

The Real Time Reduction Report shows a breakdown of the Reduction Ratio by Application achieved by the Exinda appliance during the last 5 minutes. The menu at the top of the report allows you to adjust the report auto-refresh rate (every 30 seconds by default) as well as an option to view the inbound, outbound or bi-directional traffic reduction ratio (bi-directional by default).

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

Reduction Ratio is a ratio that compares After Exinda (AE) to Before Exinda (BE).

$$\text{Reduction Ratio} = (\text{Data Transfer Size BE} - \text{Data Transfer Size AE}) / \text{Data Transfer Size BE}.$$

Application Response

The Real Time Application Response Report shows the Round-trip Time (RTT), Normalized Network Delay, Normalized Server Delay, Normalized Total Delay, Network delay, Server delay, Transaction Delay, Transaction Count, and Flow Count for each application monitored by the Exinda appliance during the last 10 seconds. The menu at the top of the report allows you to adjust the report auto-refresh rate (every 30 seconds by default).

The Applications are sorted by Round-trip Time.

Application Name	RTT (ms)	Normalized Network (ms/kb)	Normalized Server (ms/kb)	Application Response			Transaction Delay (ms)	Transaction Count	Flows
				Normalized Delay Total (ms/kb)	Network (ms)	Server (ms)			
HTTPS	192.49	1.07	7.88	8.94	1.88	13.90	15.78	1	4

Note These statistics are only available if the Performance Metrics ASAM Module is enabled on the [System | Setup | Monitoring](#) page.

Host Health

The Real Time Host Health Report shows the Retransmitted Bytes, Aborted Connections, Refused Connections, Ignored Connections and Flow Count for each Internal and External Host monitored by the Exinda appliance during the last 10 seconds. The menu at the top of the report allows you to adjust the report auto-refresh rate (every 30 seconds by default).

The Hosts are sorted by Retransmitted Bytes.

Health					
Internal IP	Retransmitted (bytes)	Aborted	Refused	Ignored	Flows
192.168.0.59	0	0	0	0	1
192.168.0.87	0	0	0	0	1
192.168.0.1	0	0	0	0	1
192.168.0.35	0	0	0	0	1
192.168.0.209	0	0	0	0	1
192.168.60.59	0	0	0	0	1
192.168.10.206	0	0	0	0	1
172.16.0.222	0	0	0	0	1

Aborted Connections	Connections that were unexpectedly aborted by either the client or server sending a TCP reset.
Refused Connections	Connections that were refused by the server (TCP SYN sent, received ICMP refused or TCP reset in response).
Ignored Connections	Connections that were ignored by the server (TCP SYN sent, received nothing in response).

Note These statistics are only available if the Performance Metrics ASAM Module is enabled on the [System | Setup | Monitoring](#) page.

Interface Reports

The Interface Reports provides you with statistics on the total amount of traffic that has passed through the monitored network interfaces. This page allows you to see the inbound and outbound throughput for all traffic on the wire as seen by the Exinda appliance.

The following Interface Reports are available on all models:

- [Interface Throughput Reports](#)
- [Interface Packets Per Second \(PPS\) Reports](#)

View the data throughput on the interfaces

The Interface Throughput Report provides you with statistics of the total data that has passed through each WAN interface on each bridge. This report allows you to see the inbound and outbound throughput for all traffic on the wire, over time.



1. Click **Monitor > Interfaces > Throughput**.
2. To filter the data in the chart, select the WAN or out-of-path interface to display.

Bridge WAN ports, policy-based routing interfaces, and WCCP interfaces are available. Selecting **All** includes all out-of-path interfaces in the report.

3. To identify how much traffic falls above a specific percentile, select the desired value from the **Select Percentile Marker to Display** list.

The table at the bottom of the report shows the total amount of data transferred into and out of the WAN interface(s), and also the maximum and average throughput values for the selected time period. The values in the table are automatically updated when the interactive flash graphs are manipulated.

Note Given that this report shows all data on the wire, the report may also include traffic that is not seen on the WAN, such as local LAN broadcasts, etc.

WAN/Out-of-path Interface Selection: Select Percentile Marker to Display: 

WAN Interface Throughput Summary (ALL)			
Data Direction	Total Data (MB)	Throughput Avg (Mbps)	Throughput Max (Mbps)
Inbound	1200094.45	113.39	286.88
Outbound	553478.38	52.30	136.31

View the outbound packet rate for all traffic


The Interface PPS Report provides you with the packet rate that has passed through each bridge on the Exinda appliance. This report allows you to see the outbound packet rate for all traffic on the wire, over time.

1. Click **Monitor > Interfaces > Packets Per Second**.
2. To filter the data in the chart, select the Bridge or out-of-path interface to display.
Bridge WAN ports, policy-based routing interfaces, and WCCP interfaces are available. Selecting **All** includes all out-of-path interfaces in the report.
3. To identify how much traffic falls above a specific percentile, select the desired value from the **Select Percentile Marker to Display** list.

The table at the bottom of the report shows the maximum and average PPS values through the bridge for the selected time period. The values in the table are automatically updated when the interactive flash graphs are manipulated.

Note Given that this report shows all data on the wire, the report may also include traffic that is not seen on the WAN, such as local LAN broadcasts, etc.

Bridge/Out-of-path Interface Selection: ALL
Select Percentile Marker to Display: None




Bridge PPS Summary (ALL)		
Data Direction	Packets Per Second (Avg)	Packets Per Second (Max)
Outbound	215	12,629

Network

The Network Reports display a time series of the top 10 inbound and outbound Applications, Application groups, Hosts, Conversations, URLs or Users on the network. The Applications series bundles applications not in the top 10 in an "Other" category.


1. Click **Monitor > Network**.
2. Select the graph you wish to view from the drop-down list.
3. Select the time range for the graph.

The LAN and WAN graphs are updated to display the selected parameters.



4. Remove specific types of traffic from the graph by deselecting their checkbox in the legend below the graph.
5. 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.



Control Reports

There are three Control Reports available:

- ["Policies Report" on page 47](#): Shows Circuit, Virtual Circuit and Policy throughput over time.
- ["Discard Report" on page 50](#): Shows discard (drop) statistics over time.
- ["Prioritization Report" on page 51](#): Shows policy prioritization statistics over time.

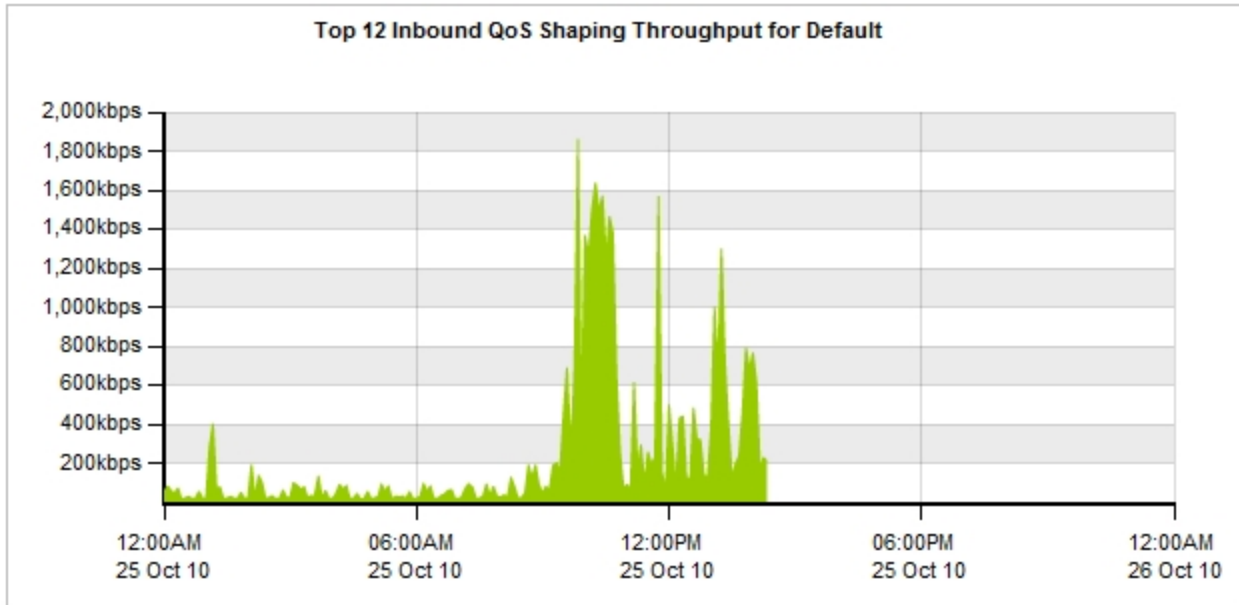
Policies Report

The Policies 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.

Note To view the Policies Report, navigate to **Monitor > Control > Policies** on the Web UI, advanced mode.


The menu at the top of the report allows you to filter the report by Circuit, Virtual Circuit or Policy.

When viewing Circuits, the report shows all Virtual Circuits within the selected Circuit, in both the graph and the table.




Optimizer Inbound QoS Summary (Default)				
VC Name	Maximum BW	Avg Rate / Max Rate (kbps)	Current Rate (kbps) / Utilization (%)	
<input checked="" type="checkbox"/> Internet inbound	10000kbps	237.00 / 1,861.00	<div style="width: 2%; border: 1px solid black; display: inline-block;"></div>	484.00 / 4.84

If the Virtual Circuit selected is a Dynamic Virtual Circuit, then the following graph will appear above the throughput graph.



The 'Active Hosts' line shows the number of Active Hosts for this Dynamic Virtual Circuit. The 'Maximum Exceeded' line shows the number of Hosts that have exceeded the Hosts limit for this Dynamic Virtual Circuit.

If a Virtual Circuit is selected, the report shows all Policies within the selected Virtual Circuit in both the graph and the table.



Optimizer Inbound QoS Summary (Internet inbound)				
Policy Name [+] Show Details	Avg Rate / Max Rate (kbps)	Current Rate (kbps) / Utilization (%)		
<input checked="" type="checkbox"/> 10 - P2P - Choke 1%-3%	0.00 / 0.00	<input type="text"/>	0.00 / 0.00	
<input checked="" type="checkbox"/> 20 - Recreational - Limit Low 2%-10%	40.00 / 915.00	<input type="text"/>	0.00 / 0.00	
<input checked="" type="checkbox"/> 30 - Software Updates - Limit Med 3%-50%	7.00 / 1,258.00	<input type="text"/>	0.00 / 0.00	
<input checked="" type="checkbox"/> 40 - Voice - Guarantee Critical 15%-100%	0.00 / 22.00	<input type="text"/>	0.00 / 0.00	
<input checked="" type="checkbox"/> 50 - Thin Client - Guarantee High 10%-100%	0.00 / 1.00	<input type="text"/>	0.00 / 0.00	
<input checked="" type="checkbox"/> 60 - Files - Guarantee Med 8%-100%	0.00 / 0.00	<input type="text"/>	0.00 / 0.00	
<input checked="" type="checkbox"/> 70 - Web - Guarantee High 10%-100%	118.00 / 1,547.00	<input type="text"/>	57.00 / 0.57	
<input checked="" type="checkbox"/> 80 - Mail - Guarantee Med 8%-100%	44.00 / 472.00	<input type="text"/>	93.00 / 0.93	
<input checked="" type="checkbox"/> 200 - ALL - Guarantee Low 5%-100%	23.00 / 223.00	<input type="text"/>	23.00 / 0.23	

The table underneath the graph shows some additional information for the selected time period. 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 10 seconds.

Discard Report


The Discard Report shows how many packets were discarded as a result of Optimizer Discard Policies.

Note To view the Optimizer Discard Report, navigate to **Monitor > Control > Discard** on the Web UI, advanced mode.

Using the menu at the top of the Report, you can filter the Report by Virtual Circuit or Policy. Only Virtual Circuits and Polices that contain discard actions are displayed.

Virtual Circuit: Policy:

The graph shows the number of packets discarded over time. The table below show the total number of discarded packets over the selected time period.



Optimizer Outbound Discarded Summary	
Policy Name	Total Packets Discarded
5 - Discard	6

It is also possible to "zoom in" to a particular time period using the interactive flash-based graphs and disable certain plots by de-selecting their checkbox in the legend.


Prioritization Report

The Prioritization Report tells you how often critical applications were prioritized (also referred to re-ordering or re-queuing).

Note To view the Prioritization Report, navigate to **Monitor > Control > Prioritization Ratio** on the Web UI, advanced mode.

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.

$$\text{Prioritization Ratio} = 100 \times \frac{\text{Number of Packets Re-ordered}}{\text{Number of Total Packets}}$$



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.

Optimization Reports

There are two Optimization Reports available:

- [Reduction](#): Shows Optimizer Reduction saving and statistics over time.
- [Edge Cache](#): Shows Edge Cache saving and statistics over time.

Reduction Report

The Optimization Reduction Report shows the WAN Memory reduction throughput and percentage, the reduction statistics and peers (remote sites) and Applications with reduced traffic. You can choose to view a particular time period using the time range selection bar at the top of the page.

Note To view the Optimization Reduction Report, navigate to **Monitor > Optimization > Reduction** on the Web UI, advanced mode.

There are two options when it comes to configuring the Reduction Report.


- **Direction**: Choose which direction you wish to view reduction statistics. Choices are Bi-directional, Inbound or Outbound.

- Reduction Type: Choose how you want reduction statistics to be displayed. Choices are Percentage and Throughput.


The drop-down menus at the top of this report allow you to configure these options.

Direction: **Reduction Type:**

When reduction statistics are displayed as a percentage, the following graph will be visible:





























When reduction statistics are displayed as throughput, the following graph will be visible:



You can choose to view a particular time period using the time range selection bar at the top of the page. It is also possible to "zoom in" to a particular time period using the interactive flash-based graphs and disable certain plots by de-selecting their checkbox in the legend.

The table below the graphs show reduction statistics broken down by peer (remote Exinda appliance) as well as 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
	Total	88,831.21	59,957.9		32.50

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


Reduction Ratio is a ratio that compares After Exinda (AE) to Before Exinda (BE).

$$\text{Reduction Ratio} = \frac{\text{Data Transfer Size BE} - \text{Data Transfer Size AE}}{\text{Data Transfer Size BE}}$$


View Edge Cache Statistics

The Edge Cache report shows the amount of data reduced by the Edge Cache.

1. Click **Monitor > Optimization > Edge Cache**.
2. Select the type of report to view.
 - **Edge Cache Throughput**—the amount of WAN and LAN traffic that was sent from the Edge Cache instead of from the HTTP application.




- **Edge Cache Reduction Ratio (%)**—the percentage of requested data that was sent from the Edge Cache instead of from the HTTP application.



3. To display the number of requests per second and the number of hits per second, review the

Edge Cache Statistics graph. A hit occurs when a request is made for an object stored in the Edge Cache.



- Review the table to see 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

Service Level Reports

The Exinda appliance provides various ways of measuring Application Performance and Network Service Levels:

- [APS](#): Application Performance Score passively measures application response times and scores application performance against preset criteria.
- [SLA](#): Service Level monitoring actively measures network response times.
- [TCP Efficiency](#): Passively measures TCP retransmissions and calculates TCP Efficiency.
- ["View TCP health" on page 100](#): Passively measures TCP connection states and calculates TCP Health.

View the Application Performance Score results

The Application Performance Score combines selected Application Performance Metrics to form an overall score that is used to monitor the performance of a networked application.

Note To monitor a new application, see ["Create an Application Performance Score object" on page 113](#).

The Exinda appliance calculates the APS by comparing the results of each metric against the threshold for the metric and is classified into one of three categories:

- Good — The baseline for the application is good, and the application is performing within the expected levels (below the threshold), and users will be happy with the performance of the application.
- Tolerated — The performance of the application is less than expected, but is still performing within a range that you should be able to tolerate (between the threshold and four times the threshold). The performance isn't great, but users will be OK with it.
- Frustrated — The application is performing really poorly (more than four times the threshold), and users will be frustrated with the performance.

The APS score is a number between 0 and 10 that measures the network performance of an application:

$$\text{aps} = 10 * (\text{satisfied samples} + (\text{tolerated samples} / 2)) / \text{total samples}$$

Example: Calculating an APS

A threshold is configured for Network Delay as $T = 300$ msec for HTTP.

In one 10s period, 11 flows are sampled for HTTP with the following results:

- 5 flow samples are > 300 ms but < 1200 ms
- 6 flow samples are < 300 ms

The APS score is calculated as follows:

$$\text{aps} = 10 * (6 + 5/2) / 11 = 7.7$$

The Application Performance Score report displays the performance of each APS object over time in a chart, and the table below lists each APS Object and includes the individual metrics used to calculate the score.

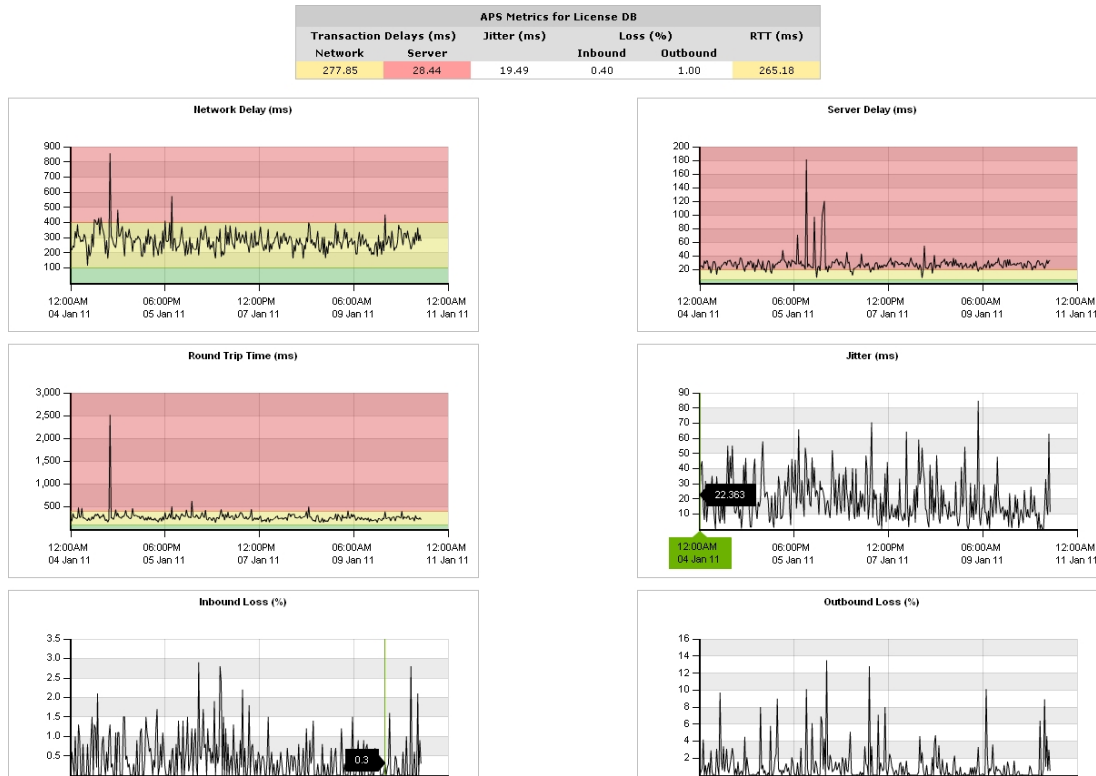
1. Click **Monitor > Service Levels** and switch to the **Application Performance Score (APS)** tab.



APS Scores							
Name	Score	Transaction Delays (ms)		Jitter (ms)	Loss (%)		RTT (ms)
		Network	Server		Inbound	Outbound	
<input checked="" type="checkbox"/> HTTP	9.52	66.32	125.11	40.18	0.50	0.60	70.25
<input checked="" type="checkbox"/> License DB	4.13	277.85	28.44	19.49	0.40	1.00	265.18
<input checked="" type="checkbox"/> SMTP	9.98	41.86	2.57	1.40	2.10	0.00	250.79

The colors indicate the category for each metric: Good is green, Tolerable is yellow, and Frustrated is red. When no color is used it indicates a metric that does not contribute to the APS score because no threshold has been configured for that metric.

2. To change the time period that the report covers, select the **Range** from the list.
3. To view the time series graphs for individual metrics, click the APS name.



Network Response (SLA) Reports



SLA monitoring is a valuable tool for ensuring the performance of your IP provider against predefined criteria.

The SLA monitoring works by sending 1 ICMP ping every 10 seconds (each of 64-bits length by default) to the remote site to collect statistics.

To add an SLA Site, click the **Add/Edit SLA Site...** link.

SLA Statistics for DNS					
Site Name	IP Address	Availability	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)
DNS	203.2.192.124	100.00 %	44.34	57.65	113.15


- 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 average and the maximum network delay from the local Exinda appliance to a remote host and back again.



TCP Efficiency Report

The TCP Efficiency Report shows the total efficiency of all TCP connections over time.

The Report can be viewed by Applications, Internal Hosts or External Hosts and the view can be changed by selecting the desired category from the drop-down at the top of the page.



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	
SSH	0.007	10.254	2.627	98.108	97.57
ICMP	0.000	0.001	0.000	0.002	98.51
RDP	0.000	0.075	0.025	1.656	98.52
Replify	0.000	0.337	0.100	7.769	98.76
SMTP	1.315	107.018	0.024	14.494	98.90
Quicktime	0.231	22.299	0.000	0.385	98.98
HTTP	1.745	388.939	7.281	540.749	99.03
LinkedIn	0.015	2.740	0.007	0.687	99.38
Facebook	0.031	6.423	0.006	0.990	99.50
Jabber	0.000	0.038	0.000	0.033	99.56
Twitter	0.002	0.582	0.001	0.361	99.73
HTTP-ALT	0.000	0.074	0.001	0.161	99.73
Skype	0.001	0.533	0.000	0.031	99.75
SalesForce	0.006	3.245	0.000	0.433	99.84

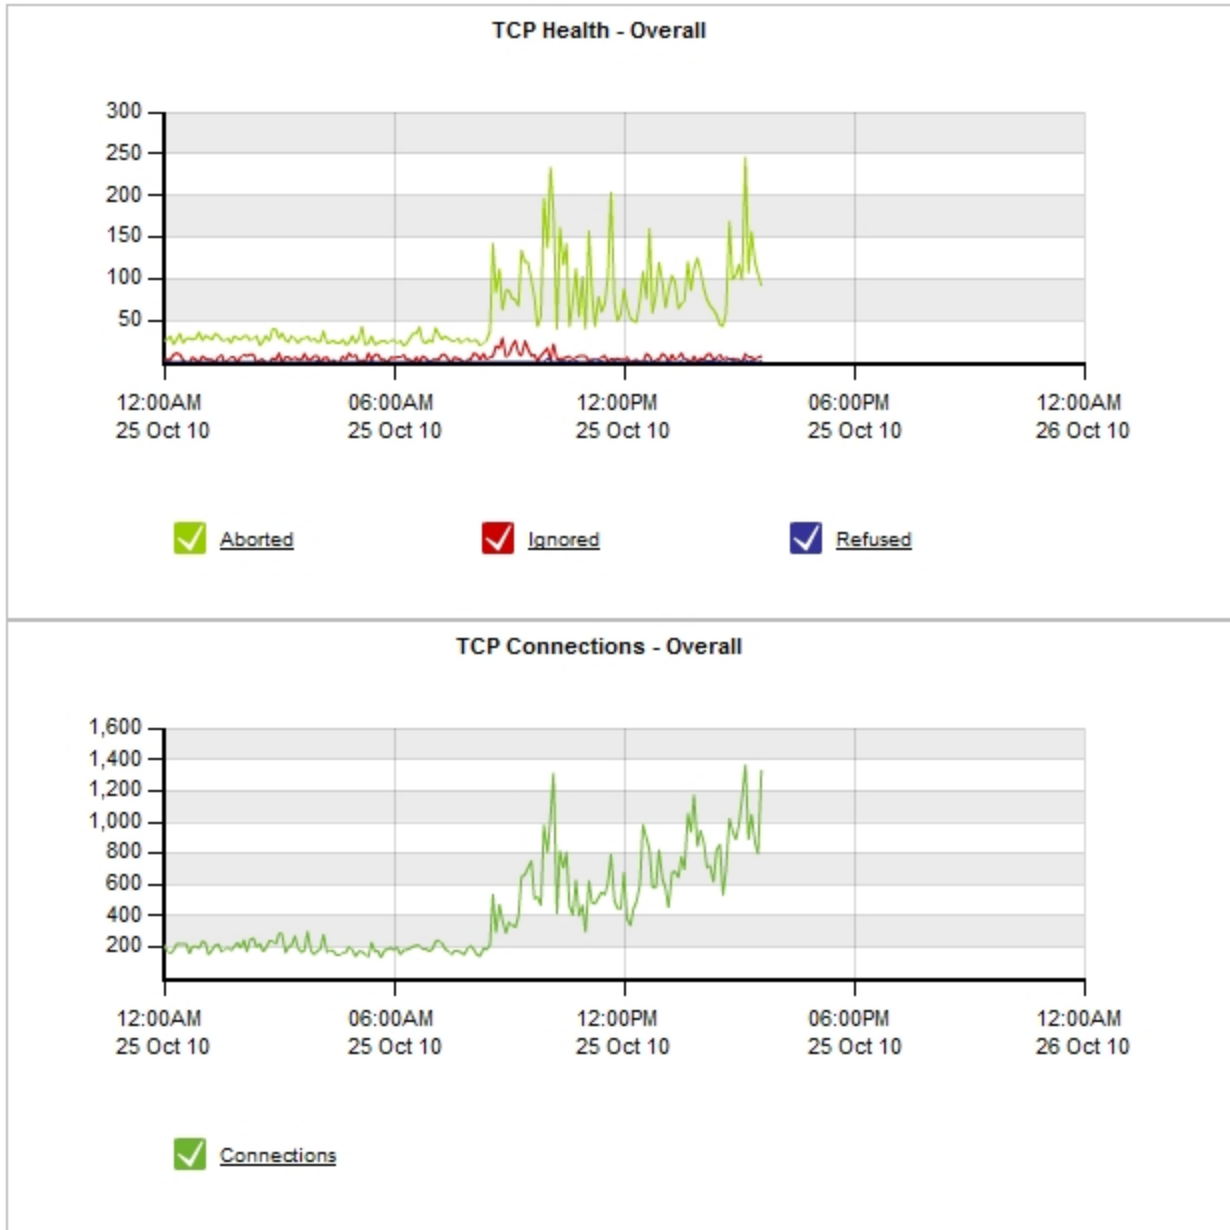
View TCP health

When monitoring the overall performance of the TCP traffic on your network, watch for connections that are

aborted, refused, or ignored over time.

- **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.

The TCP Health Report displays aborted, ignored, and refused TCP connections for the selected time period, as well as displaying the total TCP connections.



1. Click **Monitor > Service Levels** and switch to the **TCP Health**.
2. From the Category list select **Applications**.
The Report can be viewed by Applications, Internal Hosts, or External Hosts.
3. To change the time period that the report covers, select the **Range** from the list.

4. Click the name of the application to view the TCP Health details and a graph for that item.

Top 50 Applications				
	Connections	Aborted	Ignored	Refused
HTTP	34263	5172	1	6
HTTPS	21925	4941	8	1
ExindaCom	4440	0	822	0
Flash	820	233	0	0
HTTP-ALT	695	0	112	0
POP-SSL	147	110	0	0
LinkedIn	371	93	0	0
MAPI	674	91	0	0
Facebook	688	89	0	0
CIFS	64	3	61	0
SMTP	517	32	0	0
Replify	769	0	0	31
Windows Updates	58	26	0	0
SalesForce	198	21	0	0


System Reports

The System Reports show various statistics about system performance and health. There are 6 System Reports available:

- [Connections Report](#): Shows the number of concurrent and new connections.
- [Accelerated Connections Report](#): Shows the number of concurrent and new accelerated connections.
- [CPU Usage Report](#): Shows the CPU usage over time.
- [CPU Temperature Report](#): Shows the CPU temperature over time.
- [RAM Usage Report](#): Shows the RAM usage over time.
- "Disk IO" on page 68: Shows the read and write disk usage over time.
- [Swap Usage Report](#): Shows the Swap usage over time.

Connections Report

This report shows the number of concurrent connections as well as the connection establishment rate through the Exinda appliance over time.




Systems that report unusually high spikes in the number of connections or the connection rate may be experiencing some form of DoS attack or network problem.

Report: Accelerated Connections

This report shows the number of concurrent accelerated connections as well as the accelerated connection establishment rate through the Exinda appliance over time.

1. Click **Monitor > System** and switch to the **Accelerated Connections** tab.


The reports are displayed.



The second graph shows the number of connections through each of the specific application acceleration modules, such as SMB or SMB2 Acceleration, SSL Acceleration, and NCP Acceleration.

CPU Usage Report


This reports shows the Exinda appliance's CPU usage over time.



Systems running at high CPU usage for long periods of time may be overloaded or may be experiencing a problem. Contact Exinda TAC if the CPU usage constantly reports close to 100%.

CPU Temperature Report


This report show the CPU temperature over time in degrees Celsius.



Systems running at very high temperatures may be experiencing a problem and system performance may be affected. Contact Exinda TAC if the CPU temperature constantly reports hot.

RAM Usage Report


This report shows how much system RAM the Exinda appliance is consuming over time.



Systems that are running low on RAM may experience performance problems. Contact Exinda TAC if RAM usage increases close to the maximum.


Disk IO

This report shows the disk usage for each service in kB/s (read and write), as well as the total usage for each system disk.



Swap Usage Report

This report shows the swap or page file usage of the Exinda appliance over time.



Excessive amounts of swapping may impact system performance so this report provides a way to determine how much swapping the Exinda appliance is doing.

Applications Report

The Application Reports displays a breakdown of the traffic that has passed through the monitored links by Application Object or Application Group. There are also special Application Reports for URLs and VoIP. Use the Application Reports to determine what applications are using the link the most and at what speeds.

There are 4 Application Reports:

- [Application Groups Report](#): Shows the top Application Groups.
- [Individual Applications Report](#): Shows the top Application Objects.
- [URLs Report](#): Shows the top URLs.
- [VoIP Report](#): Shows VoIP flows and MOS information.

View Unclassified Applications

The Applications Report may also contain links to Discovered Ports. These are links to inbound and outbound applications which have not been classified.

1. Click **Monitor > Applications**.
2. If unclassified applications are sending traffic through the Exinda appliance, a link to Discovered Ports is displayed. To display the unclassified applications, click **Displayed Ports**.

The Discovered Ports report is displayed with source and destination ports for each unclassified applications.


Each discovered port can be broken down to the internal/external host by clicking on the link. This helps to drill down and understand what the application really is and what it should be classified as.

Top 10 Inbound Discovered Ports					
	Application	Packets	Data (MB)	Throughput Avg (kbps)	Throughput Max (kbps)
1	tcp ports 49471 & 6323	366	0.028	0.21	0.30
2	tcp ports 1794 & 41686	257	0.018	0.20	0.30
3	tcp ports 4443 & 6030	186	0.015	0.92	1.74
4	tcp ports 44371 & 10752	152	0.012	0.20	0.30
5	tcp ports 34688 & 33192	129	0.010	0.22	0.33
6	tcp ports 2800 & 10752	158	0.009	0.16	0.32
7	tcp ports 4692 & 48537	161	0.009	0.13	0.21
8	tcp ports 55099 & 19029	107	0.009	0.21	0.28
9	tcp ports 1283 & 62371	59	0.005	0.24	0.33
10	tcp ports 50236 & 48537	18	0.002	0.21	0.37

Note When deciding how to classify a discovered port, look for a common destination port. If more than two entries appear with the same destination port, the chances are that by adding that port to an Application Object, the application will be classified correctly.

Application Groups Report

The Top 10 Application Groups are shown in a pie chart. You can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie.



Each table shows the top Application Groups together with the number of packets, number of flows data transferred and throughput statistics. Click on the 'Show Details' link to expose RTT, Normalized Delays, Transaction Delays, and Efficiency statistics for each Application Group.

Click on the Application Group name to drill-down into the Individual Applications Report, filtered by that particular Application Group.

Top 50 Inbound Application Groups					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
[+] Show Details					
Web	9442485	6807.044	14.25	8801.78	15165
Streaming	3331989	4439.998	152.12	930.83	1388
Mail	3951889	3930.893	54.41	3384.16	578
Recreational	2227319	2971.643	107.79	1377.49	2501
Social Networking...	1266370	1690.538	126.17	1377.49	1296
File Services	1338624	1096.368	494.85	7970.57	32
Software Updates	431603	604.262	788.23	4346.87	181
VPN	736575	596.622	31.05	1326.80	30
Voice	1069909	140.829	11.51	409.92	85
Interactive	859128	81.518	4.22	448.42	119


Note To customize the Individual Application Objects that make up an Application Group, see the [Objects | Applications | Application Groups](#) page.

View All Network Activity for a Specific User

The Applications reporting displays the inbound and outbound traffic that has passed through the Exinda broken down by application for a specific user. The applications graphs can be used to determine which applications are currently using the link the most and at what speeds.

The table view shows the amount of packets and Megabytes transferred as well as the average and maximum throughput. The Top Application Objects are shown in a pie chart. You can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie.

1. Click **Monitor > Users**.
2. In the Select Users to View list, select **Internal**.
3. "[Set the Time Period Reflected in the Report](#)" on page 76.
After the date range is select, the graphs and charts are immediately updated.
4. In the details table, click the user name.
The **Traffic Analysis - Users - Applications** report is displayed and lists all the applications that have generated inbound and outbound traffic for that user.
5. On the Traffic Analysis - Users report there are multiple reports available. Select whether to display the user's **Applications**, **Conversations**, **URLs**, or **Hosts** report.
 - Applications**—All applications sending inbound and outbound traffic on the user's computer.
 - Conversations**—All inbound and outbound data sent and received from the user's computer.
 - URLs**—All inbound and outbound URL requests.
 - Hosts**—All internal and external hosts communicating with the user's computer.
6. Hover over the pie slices to view the amount of data transferred in megabytes and percentage.



The table shows the top Application Objects together with the number of packets, number of flows, data transferred and throughput statistics.

Top 50 Inbound Applications					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
[+] Show Details					
HTTP	6817611	5334.422	18.45	8801.78	12960
Shoutcast	1873255	2523.348	258.10	307.57	9
SMTP	1927337	2297.050	297.60	2584.53	291
YouTube	1088274	1512.809	367.74	1377.49	300
HTTPS	2619277	1472.302	7.89	1407.26	2197
IMAP-SSL	1769795	1459.845	26.97	534.73	186
Flash	914997	1199.138	76.21	828.16	1354
FTP	728912	1034.463	5432.26	7970.57	12
MPEG	542279	716.296	203.89	930.83	14
OpenVPN	372946	471.839	96.49	148.52	19

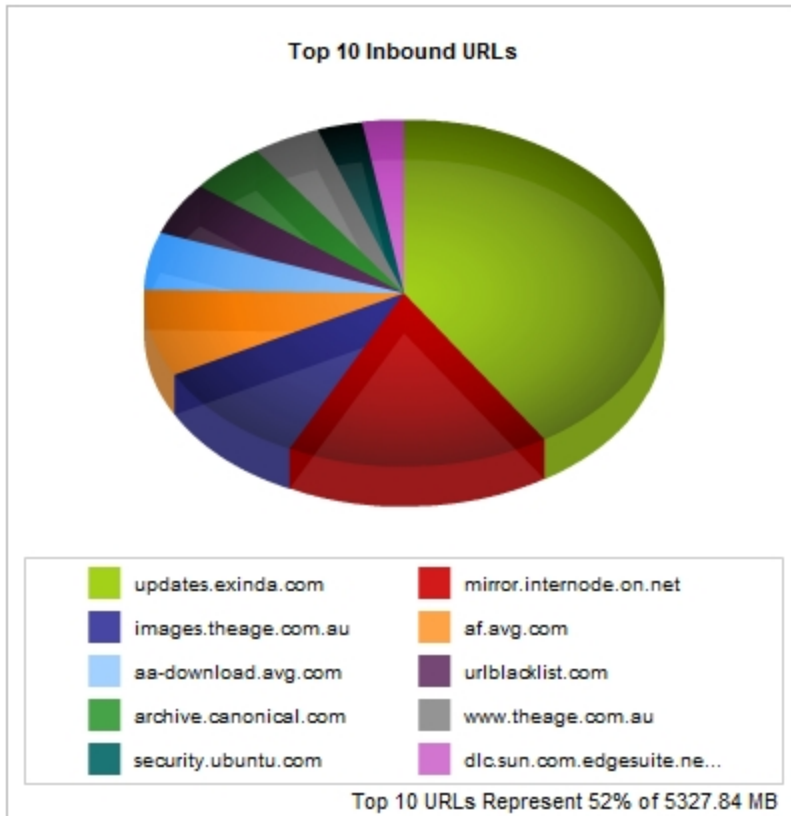
Note If a previously defined Application Object has been deleted, it will appear in these reports as 'Deleted Application'.

- To display additional details in the table for each user such as round trip time (RTT), transaction delay and efficiency statistics, click **Show Details**.

URLs Report

The URLs report displays a breakdown of the inbound and outbound traffic that has passed through the Exinda to and from various URLs that users are accessing. URLs are requested when users visit HTTP sites. They are stored in the form of a domain/host name on the Exinda appliance.

The Top 10 URLs are shown in a pie chart for inbound and outbound traffic. You can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie.



Each table shows the top URLs together with the amount of data transferred, number of packets, number of flows and throughput statistics. Click on the 'Show Details' link to expose RTT, Normalized Delays, Transaction Delays, and Efficiency statistics for each URL.


Click on a specific URL to drill-down into the [Hosts Report](#), filtered by that particular URL. You have the option to view Internal or External hosts in the drill-down.

Top 50 Inbound URLs					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
updates.exinda.com	815189	1124.869	4981.04	8175.02	7
mirror.internode.on.net	317642	449.981	2595.95	7900.50	72
images.theage.com.au	448747	273.494	83.35	745.06	309
af.avq.com	158276	222.598	2762.92	8295.91	40
aa-download.avq.com	118371	147.128	1883.24	1617.50	2
urlblacklist.com	98498	136.172	672.00	782.32	7
archive.canonical.com	92519	131.373	417.13	2834.83	24
www.theage.com.au	91250	115.655	55.37	101.83	315
security.ubuntu.com	54822	77.100	168.43	2422.13	50
dlc.sun.com.edgesuite.net	50849	71.411	823.94	713.44	1

VoIP Report

The VoIP Report shows call quality over time. The VoIP Report automatically includes any RTP-based VoIP call, including SIP, H.323 and Cisco Skinny.

The graph shows 3 series, the number of "Good", "Tolerable" and "Bad" calls over time. The table below lists the worst quality inbound and outbound VoIP calls over the selected time period.



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

Green (Good)	Number of calls with a MoS greater than 4.
Yellow (Tolerable)	Number of calls with a MoS between 2 and 4.
Red (Bad)	Number of calls with a MoS less than 2.

Report on Network Activity by User

Use the information in reports to determine how the policies on your Exinda can improve the quality of service and the experience of your network users.

The following reports identify user activity on the network:


- ["View All Network Activity for a Specific User" on page 77](#)
- ["Top Users Generating Traffic" on page 75](#)
- ["Top Internal and External Users on the Network" on page 79](#)
- ["Real-time Traffic by Hosts" on page 80](#)
- ["View real-time inbound and outbound conversations" on page 81](#)

Top Users Generating Traffic

The top users generating inbound and outbound traffic are displayed in a graph. The table view shows the amount of packets and Megabytes transferred as well as the average and maximum throughput per user.

1. Click **Monitor > Users**.
2. Select whether the charts display **Internal** or **External** users.
3. ["Set the Time Period Reflected in the Report" on page 76](#).
After the date range is select, the graphs and charts are immediately updated.
4. Hover over the pie slices to view the amount of data transferred in megabytes and percentage.

Top 8 Internal Users Receiving Inbound Traffic



Each table shows the top Users together with the amount of data transferred, number of packets, number of flows and throughput statistics.

- To display additional details for each user such as round trip time (RTT), transaction delay and efficiency statistics, click **Show Details**.
- To display the applications that are generating the inbound or outbound traffic for user, click the user-name. You have the option to view Applications, URLs, Hosts and Conversations in the drill-down.

Top 50 Internal Users Receiving Inbound Traffic					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
[+] Show Details					
MELB\Pforto	256599	328.319	264.20	4034.39	224
themis	224420	142.290	28.87	1656.61	1194
MELB\Kdeikos	76578	101.077	79.85	887.10	433
MELB\Csiakos	90826	89.588	21.56	3961.84	564
rbyrne	55507	75.806	167.84	946.30	84
MELB\Asavant	110339	71.434	7.47	612.92	1124
scott	60506	54.297	12.48	695.97	689
MELB\Matt	33342	27.351	9.67	1854.87	455
salah-pc	34944	26.560	11.65	855.17	498
MELB\Avish	41831	14.614	1.58	172.73	2243

Set the Time Period Reflected in the Report

The data displayed in the reports can be focused on specific periods of time. Date ranges are available on all

reports except the Real Time reports.

1. Select a report from the Monitor list.
2. 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

3. 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.

Range: 12:00AM 25/Oct/2010 - 12:00AM 26/Oct/2010

After the date range is select, the graphs and charts are immediately updated.

Data Granularity

The Exinda appliance stores data for the following amount of time:

- 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

View All Network Activity for a Specific User

The Applications reporting displays the inbound and outbound traffic that has passed through the Exinda broken down by application for a specific user. The applications graphs can be used to determine which applications are currently using the link the most and at what speeds.

The table view shows the amount of packets and Megabytes transferred as well as the average and maximum throughput. The Top Application Objects are shown in a pie chart. You can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie.

1. Click **Monitor > Users**.
2. In the Select Users to View list, select **Internal**.
3. "[Set the Time Period Reflected in the Report](#)" on page 76.

After the date range is select, the graphs and charts are immediately updated.

4. In the details table, click the user name.

The **Traffic Analysis - Users - Applications** report is displayed and lists all the applications that have generated inbound and outbound traffic for that user.

5. On the Traffic Analysis - Users report there are multiple reports available. Select whether to display the user's **Applications**, **Conversations**, **URLs**, or **Hosts** report.


Applications—All applications sending inbound and outbound traffic on the user's computer.

Conversations—All inbound and outbound data sent and received from the user's computer.

URLs—All inbound and outbound URL requests.

Hosts—All internal and external hosts communicating with the user's computer.

6. Hover over the pie slices to view the amount of data transferred in megabytes and percentage.



The table shows the top Application Objects together with the number of packets, number of flows, data transferred and throughput statistics.

Top 50 Inbound Applications					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
[+] Show Details			Average	Max	
HTTP	6817611	5334.422	18.45	8801.78	12960
Shoutcast	1873255	2523.348	258.10	307.57	9
SMTP	1927337	2297.050	297.60	2584.53	291
YouTube	1088274	1512.809	367.74	1377.49	300
HTTPS	2619277	1472.302	7.89	1407.26	2197
IMAP-SSL	1769795	1459.845	26.97	534.73	186
Flash	914997	1199.138	76.21	828.16	1354
FTP	728912	1034.463	5432.26	7970.57	12
MPEG	542279	716.296	203.89	930.83	14
OpenVPN	372946	471.839	96.49	148.52	19

Note If a previously defined Application Object has been deleted, it will appear in these reports as 'Deleted Application'.

- To display additional details in the table for each user such as round trip time (RTT), transaction delay and efficiency statistics, click **Show Details**.


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.

- Click **Monitor > Network**.
- In the Select Graph to Display list, select **Users - Internal** or **Users - External**.
- "Set the Time Period Reflected in the Report" on page 76.
After the date range is select, the graphs and charts are immediately updated.
- Remove specific types of traffic from the graph by deselecting their checkbox in the legend below the graph.
- 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



Real-time Traffic by Hosts

The Real Time Hosts/Users Report shows a breakdown of the Hosts/Users monitored by the Exinda appliance during the last 10 seconds. Hosts/Users are divided into Internal and External Hosts/Users.

Auto-Refresh Rate: | Show Users

1. Click **Monitor > Real Time > Hosts/Users**.

Hosts/Users are sorted by throughput, and display the packet rate and number of flows for each Host/User. The Distribution column shows the percentage of throughput a Host/User consumed relative to all the other Hosts/Users.

Inbound Hosts/Users				
IP Address (User)	Transfer Rate (kbps)	Packet Rate (pps)	Flows	Distribution (%)
Total	138.037	46	117	
172.16.0.246 (Ksiakou)	105.324	10	5	
172.16.0.134 (Pforto)	13.909	3	4	
172.16.1.70 (Selfservice)	6.639	18	3	
172.16.1.240	3.771	6	34	
172.16.0.211	3.554	3	12	
172.16.0.244 (Cniko)	1.295	2	15	
172.16.0.127 (Sshannon)	1.060	2	20	
172.16.1.74	0.684	0	1	
172.16.0.239 (Jbothe)	0.593	1	5	
172.16.0.63 (Lenehan)	0.493	0	1	
Other	0.715	2	9	

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To display the user name associated with an internal IP, select **Show Users**.









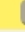

Note Active Directory must be configured on the Exinda appliances before user names can be displayed in reports. See "[Integrate the Exinda Appliance with Active Directory](#)" on page 165.

View real-time inbound and outbound conversations







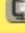



The Real-time Conversations Report shows a breakdown of the Conversations monitored by the Exinda appliance during the last 10 seconds. Conversations are divided into Inbound and Outbound directions.

- Click **Monitor > Real Time > Conversations**.

By default, the Real-time Conversations Report looks like the example below. Conversations are sorted by throughput. You can also see the packet rate and number of flows for each Conversation. Any extra information about a Conversation (a URL for example) will be shown in square brackets next to the Application.

Inbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			1408.428	284	24	
 	192.168.10.1	192.168.10.128	MAPI	570.834	82	1
 	192.168.10.9	192.168.10.128	MAPI	483.247	54	2
 	192.168.10.7	192.168.10.128	MAPI	275.334	92	2
 	192.168.10.10	192.168.10.128	MAPI	65.153	51	2
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	5.496	1	1
	192.168.10.9	192.168.10.128	LDAP	2.939	1	1
	10.20.4.1	239.255.255.250	udp ports 62612 -> 3702	1.097	0	1
	10.20.4.1	239.255.255.250	udp ports 62610 -> 3702	1.069	0	1
	192.168.10.1	192.168.0.1	NetBIOS	0.623	1	1
	192.168.10.10	192.168.10.128	LDAP	0.556	0	2
	192.168.10.132	255.255.255.255	DHCP	0.541	0	1
	192.168.10.9	192.168.0.1	NetBIOS	0.225	0	1
	10.20.3.118	10.20.255.255	NetBIOS	0.225	0	1
	192.168.10.9	192.168.255.255	NetBIOS	0.225	0	1
	10.20.11.100	224.0.0.252	udp ports 58633 -> 5355	0.212	0	1
	10.20.0.14	10.20.255.255	NetBIOS	0.193	0	1
	192.168.10.9	192.168.10.128	LDAP	0.174	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.106	0	1
	192.168.10.9	192.168.0.1	DNS	0.102	0	1
	10.20.0.181	10.20.255.255	NetBIOS	0.075	0	1

- To set how often the data updates in the table, select the frequency from the **Auto-Refresh Rate** list.
- To view only a specific IP address or subnet, type the address in the **IP/Subnet Filter** field.
The report can be filtered by IPv4 or IPv6 addresses.
- To display the optimization policy the conversation falls into, select **Show Policies**.


Outbound Conversations						
External IP	Internal IP	Application	Transfer Rate (kbps)	Packet Rate (pps)	Flows	
Total			3594.527	412	14	
 	192.168.10.7	192.168.10.128	MAPI	1826.409	196	2
 	192.168.10.10	192.168.10.128	MAPI	1184.445	125	2
 	192.168.10.1	192.168.10.128	MAPI	564.195	72	1
 	192.168.10.9	192.168.10.128	MAPI	12.200	17	2
	192.168.10.9	192.168.10.128	LDAP	3.316	1	1
	192.168.10.7	192.168.10.128	HTTPS[perf-win2k8]	2.902	1	1
	192.168.10.10	192.168.10.128	LDAP	0.565	0	2
	192.168.10.9	192.168.0.1	DNS	0.197	0	1
	192.168.10.9	192.168.10.128	LDAP	0.188	0	1
	192.168.10.1	192.168.10.128	MSRPC	0.109	0	1

- To display the user name associated with an internal IP, select **Show Users**.


6. To group individual connections within a flow as a single line item or show each connection as a separate line item, select **Group**.

Understanding the Conversation Report






When a conversation has been accelerated by the Exinda appliance, the Conversation are highlighted in yellow and the Application Acceleration technologies being 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 is also been process by WAN Memory.







When a conversation has been processed by Edge Cache it is highlighted in blue.



The following legend describes the meaning of each 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 also appear next to each conversation.


	Asymmetric: The traffic is asymmetric, and is not being accelerated.
	Local: The connection is passing through this appliance in the cluster.
	Remote: The connection is passing through another appliance in the cluster.
	Local/Remote: The connection is passing though both this and other appliances in the cluster.

Hosts Report

The Hosts report displays the inbound and outbound traffic that has passed through the Exinda appliance broken down by hosts. The hosts graphs is used to determine which hosts are currently using the link the most and at what speeds. The table view shows the amount of packets and Megabytes transferred as well as the average and maximum throughput per host. Hosts are IP address endpoint's in IP transactions and are usually client PCs or servers.

The Top 10 Hosts are shown in a pie chart for inbound and outbound traffic. You can hover over the pie slices to view the amount of data transferred as well as view the percentage of the pie.

You can choose to view either Internal or External Hosts in this Report by using the drop-down at the top of the page. You can choose to view a particular time period using the time range selection bar at the top of the page.



Each table shows the top Hosts together with the amount of data transferred, number of packets, number of flows and throughput statistics. Click on the 'Show Details' link to expose RTT, Normalized Delays, Transaction Delays, and Efficiency statistics for each Host.


Click on the 'Host Name' to drill-down into the [Applications Report](#), filtered by that particular Host. You have the option to view Applications, URLs, Hosts and Conversations in the drill-down.

Top 50 Internal Hosts Receiving Inbound Traffic						
Hostname	Address	Packets	Data (MB)	Throughput (kbps)		Flows
				Average	Max	
[+] Show Details						
PAUL-PC	172.16.0.67	256679	328.333	263.18	4034.39	224
172.16.0.96	172.16.0.96	170136	176.929	63.93	4916.38	342
Themis	172.16.1.85	224508	142.331	28.80	1656.61	1198
172.14.1.10	172.14.1.10	76604	101.079	79.24	887.10	433
CS-MACBOOK-PRO	172.16.0.182	90826	89.588	21.56	3961.84	564
STEPHEN-UBUNTU	172.16.0.114	170381	86.534	24.67	4830.92	456
RBYRNE	172.16.0.213	55507	75.806	167.84	946.30	84
172.16.0.164	172.16.0.164	56105	66.111	48.40	6803.19	406
Test	172.16.0.236	74463	61.797	15.66	612.92	934
172.16.0.127	172.16.0.127	61851	55.486	12.57	695.97	708

Conversations Report

The Conversations page displays a breakdown of the top conversations that have passed through the Exinda appliance. A conversation is defined as data that is transacted between two host machines using the same application within a specified time period. Conversations may also be referred to as sessions. The table view at the bottom of the report shows the amount of packets and Megabytes transferred as well as the average and maximum throughput during the conversation.

The pie chart shows the top 10 inbound and outbound Conversations that have occurred on your monitored interfaces. You can choose to view a particular time period using the time range selection bar at the top.



Each table shows the top Conversations together with the amount of data transferred, number of packets, number of flows and throughput statistics. Click on the 'Show Details' link to expose RTT, Normalized Delays, Transaction Delays, and Efficiency statistics for each Conversation.

Click on the 'Host Name' to drill-down into the [Hosts Report](#), filtered by that particular Host. You have the option to view Applications, URLs, Hosts and Conversations in the drill-down.

Click on the Application Object name to drill-down into the [Applications Report](#), filtered by that particular Application Object. You have the option to view Internal or External Hosts in the drill-down.

Top 50 Inbound Conversations						
External Host	Internal Host	Application	Data (MB)	Throughput (kbps)		Flows
				Average	Max	
207.5.72.97	172.16.0.67	HTTPS	264.435	461.89	828.50	5
207.5.72.97	172.16.0.96	IMAP-SSL	131.470	185.05	454.53	6
202.158.214.106	172.16.0.164	HTTP	49.833	5831.85	6803.19	2
150.101.98.86	172.16.0.213	Flash	48.308	842.00	946.30	1
204.2.179.40	172.16.0.67	Adobe Updates	44.323	4034.39	4034.39	1
150.101.98.219	172.16.0.114	YouTube	43.571	830.07	857.89	1
82.93.98.79	172.16.1.85	SMTP	40.777	175.81	257.79	15
12.155.200.151	172.16.1.85	SMTP	33.122	343.46	1000.48	15
172.16.20.132	172.16.0.63	OpenVPN	22.522	104.24	148.52	15
78.31.107.132	172.16.1.85	HTTP	19.458	797.01	817.33	1

Subnets Report

A network object, referred to as a subnet for monitoring purposes, can include multiple network subnets and/or multiple IP addresses. With the subnets report you can monitor traffic by these defined subnets, and see network usage for each of your branch offices, your business departments, a class of network devices such as printers, or any other network object that you define.

Inbound / Outbound Traffic - The inbound and outbound traffic for these subnets are reported separately, where the inbound and outbound traffic is relative to the subnet, not relative to the Exinda.

Internal / External Subnets - When subnets are defined, they can be specified as being either internal (on the LAN side of the Exinda) or external (on the WAN side of the Exinda). When viewing the subnets monitoring charts, you can filter by internal subnets or external subnets using the drop-down at the top of the page.

Drill-in Details - You can also drill into the subnets to see the applications, hosts, conversations, and users for the specific subnet or host by selecting the appropriate link within the table below the chart. The applications can be viewed as either throughput over time in a stacked area chart or as data volume in a pie chart. When viewed as throughput, you can show just the top applications or the top applications with an 'Other' group to show the total application data for the specified subnet. All other drill down data types will be shown as pie charts only.

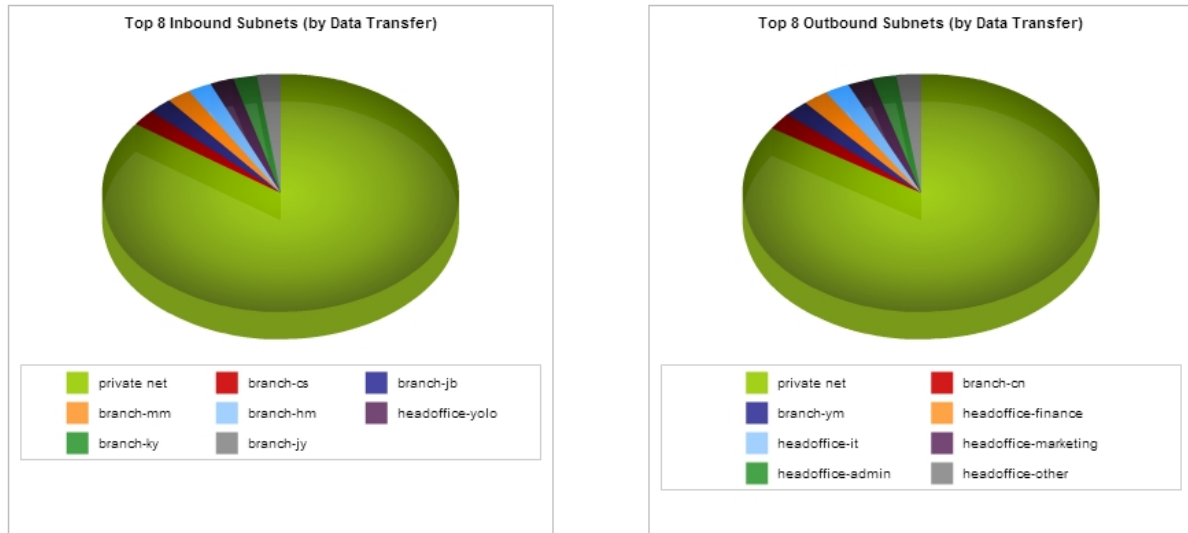
Note

- To configure a subnet for monitoring, you must create a Network Object. See ["Create a static network object"](#) on page 7.
- The network object must have the Subnet Report checkbox selected to be included in the report.
- If Network Objects/Subnet statistics collection is disabled, the Subnets report will not include application data for the time period the collection was disabled. See [Identify the statistics to collect](#).

To view the subnet summary charts

1. Go to **Monitor > Subnets**.
2. Select whether the report displays traffic for Internal or External subnets.

The traffic distribution of all subnets is displayed in a chart for either inbound and outbound traffic.



The tables below the charts show the amount of inbound or outbound data transferred and throughput statistics for each subnet. You can drill down deeper into the report to see what is happening on the network, and view the applications, hosts, conversations, URLs, and users for the specific subnet.

To modify the number of subnets that appear in the subnet summary pie chart

The maximum number of wedges that appear in the pie chart can be modified to between 1 - 10 by using the **System > Setup > Monitoring** tab. See [Monitoring Configuration](#).

To sort the subnet tables according to subnet data volume or alphabetically

The subnet tables below the subnet summary charts can be sorted by data volume or alphabetically by using the **System > Setup > Monitoring** tab. See [Monitoring Configuration](#).

To drill into the subnet to see other data within the subnet (such as applications, hosts, conversations, or users)

- **Applications within the subnet** - View the applications for the specified subnet by clicking on the View Applications link in the table. The applications can be viewed as throughput over time in a time series chart or as data volume in a pie chart.
 - **To set the chart display type:**

Go to **System > Setup > Monitoring** tab and select either **Time series chart** (stacked area chart) or **Pie graph** from the **Display for applications details per subnet** field.
 - **To show the top applications and the other remaining application traffic in the time series chart:**

Check the **Show Other** checkbox at the top of the page and press the **Apply Changes** button. The top applications and the other remaining application category will be shown in the chart and in the table below the chart.

- **To show the applications within the other remaining application traffic category:**

Uncheck the **Show Other** checkbox at the top of the page and press the **Apply Changes** button. The applications will be shown in the table below the chart. Only the top application will appear in the time series chart.

Note: The number of top applications shown in the time series chart and the number of top applications shown in the table are set independently on the **System > Setup > Monitoring** page.

- **Top Hosts within the subnet** - View the top hosts for the specified subnet by clicking on the View Hosts link in the table. The data will be displayed as data volume in pie charts.
- **Top Users within the subnet** - View the top users for the specified subnet by clicking on the View Users link in the table. The data will be displayed as data volume in pie chart.
- **Top Conversations within the subnet** - View the top conversations for the specified subnet by clicking on the View Conversations link in the table. The data will be displayed as data volume in pie charts.
- **Top URLs within the subnet** - View the top URLs for the specified subnet by clicking on the View URLs link in the table. The data will be displayed as data volume in pie charts.

PDF Reporting

PDF reports can be generated and downloaded on demand or they can be generated and emailed at scheduled intervals. The content of the PDF reports can be configured in two ways – either by exploring the data in the monitor screens and requesting a report or by going to the Report page to configure the details of the PDF report. Therefore, the following PDF report generation scenarios are supported:

- Explore the data in the monitor screens and generate an on-demand PDF report of what is shown on the screen.
- Explore the data in the monitor screens and schedule a custom PDF report to be generated using the configuration and filters shown on the screen. (Available in v6.4.3)
- Configure a scheduled PDF report using the Report page.
- Given a scheduled PDF report on the Report page, 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 1am.

On-demand reports from the monitor pages can include any time range available to the monitoring screens, including custom time ranges.


Reports scheduled from the report page, can contain one or more charts in the PDF by selecting any number of charts.

Scheduled PDF reports can be branded by uploading your logo to be displayed on the title page of the reports.

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.
For example, set the date range; Select the Internal or External selector for hosts, user charts, and subnets; Drill 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.
For example, set the date range; Select the Internal or External selector for hosts, user charts, and subnets; Drill into the data by selecting the links in the data tables under the charts; and so on.
2. Click on the schedule PDF icon in the upper-right of the screen. 
3. Optionally protect PDF documents by specifying a password.

PDF Security Option

PDF Password Protected


Enter Password:

Re-enter Password:

4. 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.

Report Details

Report Name:

Report Frequency: Daily 

Email Addresses:

5. The system will add this scheduled report to the Reporting page (**Report > PDF Report**)

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 **Report > 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.

Report Details

Report Name:

Report Frequency: Daily ▼

Email 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 **Report > 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
Daily Report (Scheduled Daily)	Interface Throughput: ALL Bridge PPS: ALL Reduction: Detailed Statistics Subnet Detailed (ALL): Applications URLs Summary Reports: Network Summary VCircuit Detailed (ALL): Optimizer Policy Throughput Appliance: Concurrent Connections CPU Usage RAM Usage	exinda_report@exinda.com		<input type="button" value="Edit"/>	<input type="button" value="Delete"/>

3. To view the report, click the Adobe PDF icon.
4. To email the report to the 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 **Report > Custom Logo**

Custom Logo

Upload New Custom Logo:

- Upload your custom logo.
- The system will insert the logo on the cover page of any scheduled PDF report.

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
cureentweek (Current Week)	Summary Reports: flows		 	<input type="button" value="Edit"/>	<input type="button" value="Delete"/>
currentday (Today)	Summary Reports: 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:	<input type="text"/>
Report Frequency:	<input type="text" value="Schedule Range"/> ▼
Email Addresses:	<input type="text"/>
[Email Addresses is Optional for On-Demand Report]	

[Add New Report](#)

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.

Monitor application performance on the network

Every company has applications that are considered business-critical that need to be performing at their best at all times. Analysing the performance of networked applications is a common task faced by network administrators. Often the root cause of an application's poor performance is not understood, and a common response is to undertake an expensive, often unnecessary upgrade of network capacity.

The Exinda appliance can monitor several properties of an application's TCP flows and collect 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 can also monitor a single metric value within TCP flows for a specified application, known as Application Performance Metrics (APM).

IT departments use the Application Performance Score (APS) to determine what is performing well, and what is performing poorly. 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.

Reports on the Application Performance Score 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 of the network that is impacting the performance of the application, for instance server delay, network delay, or jitter. This makes it easier to fix any network issues and get the application back to optimum performance levels.

- ["How are the metrics calculated?" on page 95](#)
- ["View the Application Performance Score results" on page 102](#)
- ["Capture Application Performance Metrics" on page 105](#)
- ["Calculate an Application Performance Score" on page 105](#)

How are the metrics calculated?

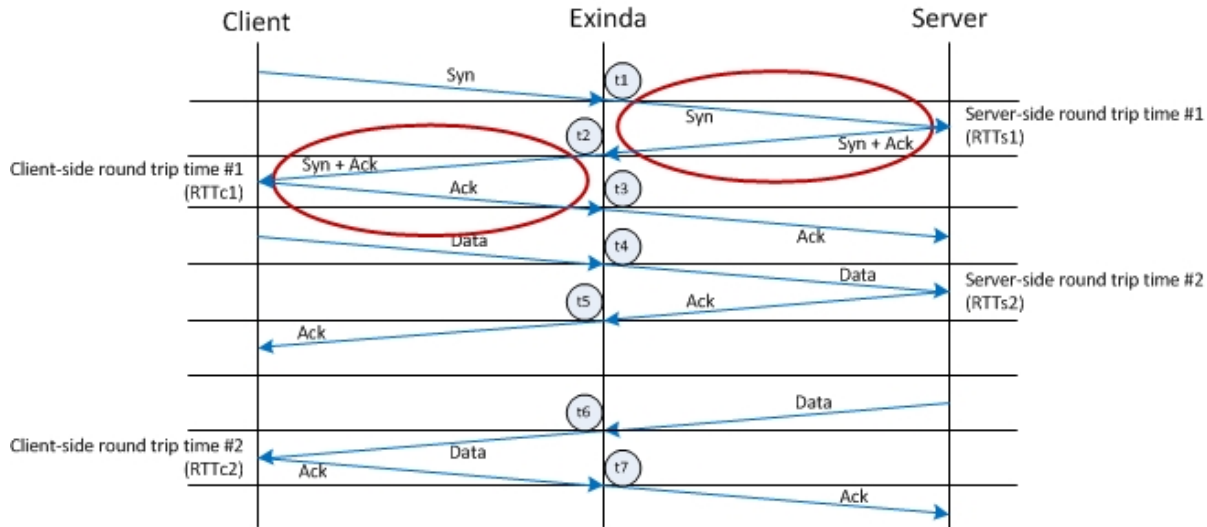
Application performance metrics measure a single aspect of the TCP flows generated by an application, while Application Performance Scores measure multiple areas of the flows for a particular application. The metrics that can be measured are:

- ["Round trip time" on page 95](#)
- ["Read and write transactions" on page 96](#)
- ["Packet loss" on page 99](#)
- ["View TCP health" on page 100](#)

Round trip time

Round trip time (RTT) is the measure of how long it takes for a very small packet to travel from a client computer, Exinda appliance, or other networking component, to cross a network and return. As each packet is intercepted, it is time stamped with a highly accurate nanosecond resolution clock source. With the Exinda appliance positioned between the client and server, two individual calculations are made to determine a total round-trip time: the round trip time from the Exinda appliance to the client and back (Client RTT), and the round trip time from the Exinda appliance to the server and back (Server RTT).

The APM estimates the round-trip time when the connection is established as part of a TCP connection, and continually updates this estimate as data is sent across the connection. The following diagram illustrates how the round trip time is calculated:



Server RTT:

- $RTTs1 = t2 - t1$
- $RTTs2 = t5 - t4$

Client RTT:

- $RTTc1 = t3 - t2$
- $RTTc2 = t7 - t6$

Average Server RTT = $(RTTs1 + RTTs2) / 2$

Average Client RTT = $(RTTc1 + RTTc2) / 2$

Average Total RTT = $avRTTs + avRTTc$

Read and write transactions

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:

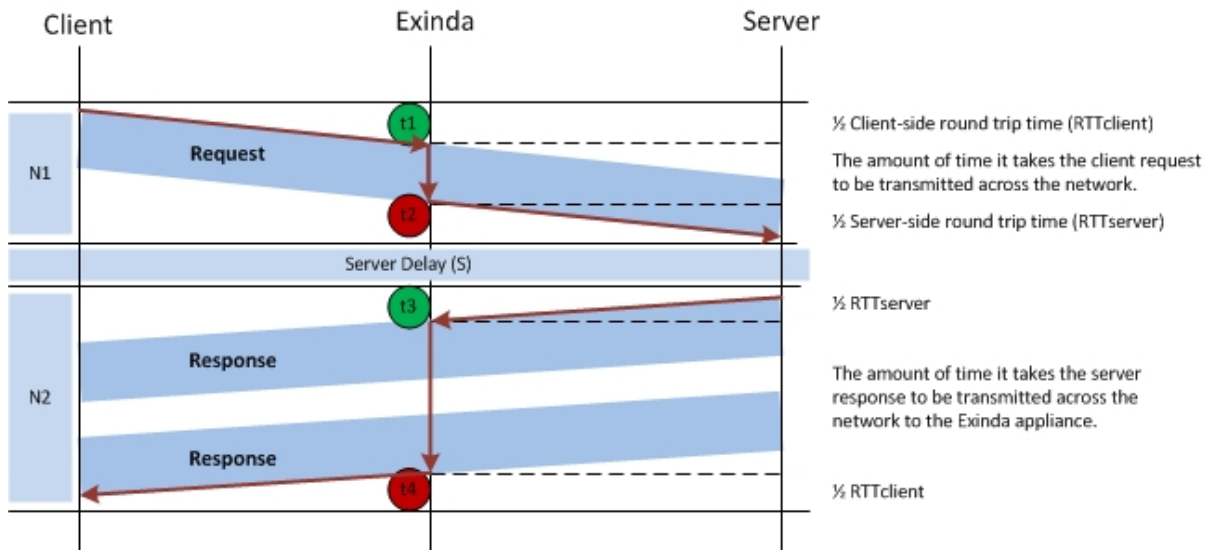
- **Network Delay** — the overall time taken for data to cross from a client to a server, or from the server to a client.
- **Server Delay** — the time taken for a server to respond to a request.
- **Network Jitter** — measures the variability of the network delay time. This is expressed as a multiple of one standard deviation.

To understand the details of the calculations, see ["Calculating read transaction metrics"](#) on page 96 and ["Calculating write transaction metrics"](#) on page 98.

Calculating read transaction metrics

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.



Network Delay for Request (N1)

1. The client sends a request to the server.
2. When the request passes through the Exinda, the time stamp is noted as the beginning of the request (t_1).
3. 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.
4. The server receives the complete client request.

Server Delay (S)

5. 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 Response (N2)

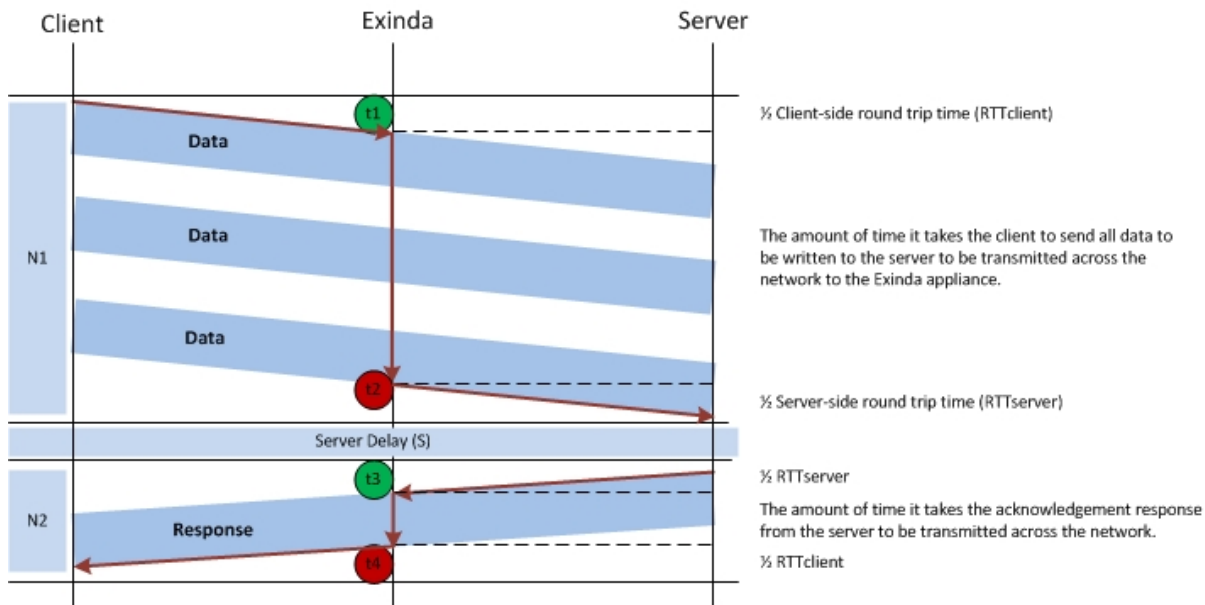
6. The server's response to the client request is sent, and may be sent in a number of packets.
7. When the first response passes through the Exinda, the time stamp is noted (t_3).
8. When the end of the last response passes through the Exinda, the time stamp is noted (t_4).
 $t_4 - t_3 =$ The amount of time it takes the data requested by the client to pass through the Exinda appliance.
9. The client receives the data requested from the server.

The total transaction time for a Read transaction is calculated as $\text{Transaction time} = N1 + S + N2$ where:

- $N1 = \frac{1}{2} RTT_{client} + (t2 - t1) + \frac{1}{2} RTT_{server}$
- $S = (t3 - t2) - RTT_{server}$
- $N2 = \frac{1}{2} RTT_{server} + (t4 - t3) + \frac{1}{2} RTT_{client}$

Calculating write transaction metrics

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 acknowledgment of receiving the data back to the client. The diagram below shows the flow of information between the client, Exinda appliance, and the server, and identifies the points in the transaction where time stamps are acquired.



Network Delay for Request (N1)

1. The client sends data to be written on the server, and may be sent in a number of packets.
2. When the first data packet starts passing through the Exinda, the time stamp is noted as the beginning of the packet (t1).
3. When the end of the last data packet passes through the Exinda, the time stamp is noted (t2).
 $t2 - t1 =$ The amount of time it takes the client to send data through the Exinda appliance.

4. The server receives all the data from the client.

Server Delay (S)

5. 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 Response (N2)

6. The server's acknowledgement response to the client that the data has been received is sent.
7. When the response passes through the Exinda, the time stamp is noted (t3).
8. When the end of the response passes through the Exinda, the time stamp is noted (t4).
 $t4 - t3 =$ The amount of time it takes the server response to pass through the Exinda appliance.
9. The client receives the response from the server.

The total transaction time for a Write transaction is calculated as $\text{Transaction time} = N1 + S + N2$ where:

- $N1 = \frac{1}{2} \text{RTT}_{\text{client}} + (t2 - t1) + \frac{1}{2} \text{RTT}_{\text{server}}$
- $S = (t3 - t2) - \text{RTT}_{\text{server}}$
- $N2 = \frac{1}{2} \text{RTT}_{\text{server}} + (t4 - t3) + \frac{1}{2} \text{RTT}_{\text{client}}$

Normalization of the Network Delay

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:

$$\text{Normalized Network Delay} = \text{Total Network delay} * 1024 / \text{transaction bytes}$$

Note The number of bytes used to normalize the calculation of the network delay during a transaction can be configured through the CLI. See [CLI: APM](#).

Packet loss

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.

$$\text{Efficiency} = 100\% * (\text{transferred} - \text{retransmitted}) / \text{transferred}$$

and

$$\text{Network Loss} = 100 - \text{Efficiency}$$


Note Network Loss, not Efficiency, is used when calculating APS.

View TCP health

When monitoring the overall performance of the TCP traffic on your network, watch for connections that are aborted, refused, or ignored over time.

- **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.

The TCP Health Report displays aborted, ignored, and refused TCP connections for the selected time period, as well as displaying the total TCP connections.



1. Click **Monitor > Service Levels** and switch to the **TCP Health**.
2. From the Category list select **Applications**.
The Report can be viewed by Applications, Internal Hosts, or External Hosts.
3. To change the time period that the report covers, select the **Range** from the list.

4. Click the name of the application to view the TCP Health details and a graph for that item.

Top 50 Applications				
	Connections	Aborted	Ignored	Refused
HTTP	34263	5172	1	6
HTTPS	21925	4941	8	1
ExindaCom	4440	0	822	0
Flash	820	233	0	0
HTTP-ALT	695	0	112	0
POP-SSL	147	110	0	0
LinkedIn	371	93	0	0
MAPI	674	91	0	0
Facebook	688	89	0	0
CIFS	64	3	61	0
SMTP	517	32	0	0
Replify	769	0	0	31
Windows Updates	58	26	0	0
SalesForce	198	21	0	0

View the Application Performance Score results

The Application Performance Score combines selected Application Performance Metrics to form an overall score that is used to monitor the performance of a networked application.

Note To monitor a new application, see "[Create an Application Performance Score object](#)" on page 113.

The Exinda appliance calculates the APS by comparing the results of each metric against the threshold for the metric and is classified into one of three categories:

- **Good** — The baseline for the application is good, and the application is performing within the expected levels (below the threshold), and users will be happy with the performance of the application.
- **Tolerated** — The performance of the application is less than expected, but is still performing within a range that you should be able to tolerate (between the threshold and four times the threshold). The performance isn't great, but users will be OK with it.
- **Frustrated** — The application is performing really poorly (more than four times the threshold), and users will be frustrated with the performance.

The APS score is a number between 0 and 10 that measures the network performance of an application:

$$\text{aps} = 10 * (\text{satisfied samples} + (\text{tolerated samples} / 2)) / \text{total samples}$$

Example: Calculating an APS

A threshold is configured for Network Delay as $T = 300$ msec for HTTP.

In one 10s period, 11 flows are sampled for HTTP with the following results:

- 5 flow samples are > 300 ms but < 1200 ms
- 6 flow samples are < 300 ms

The APS score is calculated as follows:

$$\text{aps} = 10 * (6 + 5/2) / 11 = 7.7$$

The Application Performance Score report displays the performance of each APS object over time in a chart, and the table below lists each APS Object and includes the individual metrics used to calculate the score.

1. Click **Monitor > Service Levels** and switch to the **Application Performance Score (APS)** tab.









APS Scores							
Name	Score	Transaction Delays (ms)		Jitter (ms)	Loss (%)		RTT (ms)
		Network	Server		Inbound	Outbound	
<input checked="" type="checkbox"/> HTTP	9.52	66.32	125.11	40.18	0.50	0.60	70.25
<input checked="" type="checkbox"/> License DB	4.13	277.85	28.44	19.49	0.40	1.00	265.18
<input checked="" type="checkbox"/> SMTP	9.98	41.86	2.57	1.40	2.10	0.00	250.79

The colors indicate the category for each metric: Good is green, Tolerable is yellow, and Frustrated is red. When no color is used it indicates a metric that does not contribute to the APS score because no threshold has been configured for that metric.

2. To change the time period that the report covers, select the **Range** from the list.

3. To view the time series graphs for individual metrics, click the APS name.

APS Metrics for License DB					
Transaction Delays (ms)		Jitter (ms)	Loss (%)		RTT (ms)
Network	Server		Inbound	Outbound	
277.85	28.44	19.49	0.40	1.00	265.18



Calculate an Application Performance Score

The Application Performance Score combines selected application performance metrics to form an overall score that is used to monitor the performance of a networked application. To calculate the APS the results of each metric is compared against the threshold for the metric and is classified into one of three categories:

- **Good** — The baseline for the application is good, and the application is performing within the expected levels (below the threshold), and users will be happy with the performance of the application.
- **Tolerated** — The performance of the application is less than expected, but is still performing within a range that you should be able to tolerate (between the threshold and four times the threshold). The performance isn't great, but users will be OK with it.
- **Frustrated** — The application is performing really poorly (more than four times the threshold), and users will be frustrated with the performance.

The APS score is a number between 0 and 10 that measures the network performance of an application:

$$\text{aps} = 10 * (\text{satisfied samples} + (\text{tolerated samples} / 2)) / \text{total samples}$$

Example: Calculating an APS

A threshold is configured for Network Delay as $T = 300$ msec for HTTP.

In one 10s period, 11 flows are sampled for HTTP with the following results:

- 5 flow samples are > 300 ms but < 1200 ms
- 6 flow samples are < 300 ms

The APS score is calculated as follows:

$$\text{aps} = 10 * (6 + 5/2) / 11 = 7.7$$

Because the appropriate threshold for an application is unique for each network environment, the Exinda appliance can monitor the traffic for an application and create a recommended set of thresholds that can be used for creating an APS. Recommended thresholds can be generated for existing APS objects if the current thresholds do not seem to be accurate. The Exinda appliance can also generate APS scores for non-transactional traffic between the client and server, such as Citrix XenApp Servers or Microsoft Remote Desktop.

Configure the Exinda appliance to monitor the important or business-critical applications running on the network through APS.

1. ["Create an Application Performance Score object" on page 113](#)
2. ["Generate recommended Application Performance Score thresholds" on page 114](#)
3. ["Review and modify the APS threshold values" on page 115](#)
4. ["View the Application Performance Score results" on page 102](#)

Capture Application Performance Metrics

The Exinda appliance can monitor several properties of the TCP flows specified applications, and collect

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 can also monitor a single metric value within the TCP flows for a specified application, known as Application Performance Metrics (APM). To understand what metrics are included in the APM, see ["How are the metrics calculated?" on page 95](#).

Create and monitor APMs to identify delays in communication between network objects.

1. ["Create an Application Performance Metric object" on page 106](#)
2. ["Send an alert when a threshold is crossed" on page 107](#)
3. ["View Application Performance Metrics" on page 108](#)

Create an Application Performance Metric object

Use Application Performance Metric (APM) objects to monitor the amount of time it takes to perform a specific transaction, or when certain events occur. The Exinda appliance can send an alert when the threshold of an APM metric is exceeded.

1. Click **Objects > Service Levels** and switch to the **Application Performance Metric** tab.
2. Click **Add New APM Object**.
3. Type a name for the metric.
4. In the **Metric** list, select the APM metric to monitor from the following list:

Metric	Description
bytes-lost	Bytes lost due to retransmissions.
network-delay	The time taken for data to traverse the network.
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.
server-delay	The time taken for a server to respond to a request.
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-connection-refused	The number TCP connections that are reset before the connection is established. (RST in SYN-SENT state)

Metric	Description
tcp-connections-started	The number of TCP connections initiated.
transaction-time	The total time for a transaction (network delay + server delay)

- In the **Application** list, select the application traffic for the APM to monitor.
- Select the internal network object to filter traffic.
- Select the external network object to filter traffic.
- In the **APM Threshold** field type the maximum number of milliseconds (ms) allowed for the metric to complete before an alert is triggered.
- As the Exinda appliance monitors the metric for the application, an average of the APM metric is calculated. In the **Alert Trigger Delay** list, select how long the average is calculated for the metric.

At the end of the selected time period, the APM for the application is checked to see if it exceeds the threshold. If the average is above the threshold, the Exinda appliance sends an alert. The calculation of the average is restarted at the beginning of each interval.

- If you want to be notified when the threshold is crossed, select the **Alert Enable** checkbox.
- Click **Apply Changes**.

The object is added to the list of configured APMs.

APM Name	Metric	Application Name	Network Object (Internal)	Network Object (External)	Alert Threshold	Alert Time Window	Alert Enabled	Edit	Delete
SMTP	server-delay	SMTP	Office	ALL	100 ms	1 Hour	<input checked="" type="checkbox"/>	Edit	Delete

Example

This following example shows an APM object that sends an alert when the server-delay metric for SMTP traffic matching the local internal network-object exceeds 100 ms over a period of 1 hour.

Add New APM Object

APM Name:

Metric:

Application:

Network Object - Internal:

Network Object - External:

APM Threshold:

Alert Trigger Delay:

Alert Enable:

Send an alert when a threshold is crossed

Add or modify the threshold settings that trigger alerts to be sent to administrators.

1. Click **Objects > Service Levels** and switch to the **Application Performance Metric** tab.
2. In the list of configured APMs, locate the APM to modify and click **Edit**.
3. In the **APM Threshold** field type the maximum number of milliseconds (ms) allowed for the metric to complete before an alert is triggered.
4. In the **Alert Trigger Delay** list, select how long the metric needs to remain above the threshold before the alert is sent.
5. Select the **Alert Enable** checkbox.
6. Click **Apply Changes**.

The object is added to the list of configured APMs.

APM Name	Metric	Application Name	Network Object (Internal)	Network Object (External)	Alert Threshold	Alert Time Window	Alert Enabled	Edit	Delete
SMTP	server-delay	SMTP	Office	ALL	100 ms	1 Hour	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>	<input type="button" value="Delete"/>

7. Click **System > Setup > Alerts**, and ensure that email and SNMP Trap alerts are enabled for APM.

Example

This following example shows an APM object that sends an alert when the server-delay metric for SMTP traffic matching the local internal network-object exceeds 100 ms over a period of 1 hour.

Add New APM Object

APM Name:

Metric:

Application:

Network Object - Internal:

Network Object - External:

APM Threshold:

Alert Trigger Delay:

Alert Enable:

View Application Performance Metrics

APM values are available in real time as well as summary form. This section describes how the Application Performance Metrics are accessed.

- ["Monitor the real time application response" on page 109](#)
- ["Monitor the real time TCP health" on page 110](#)
- ["View a network summary of application groups" on page 111](#)
- ["View TCP health" on page 100](#)
- ["View TCP efficiency" on page 112](#)

Monitor the 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. Click **Monitor > Real Time** and switch to the **Application Response** tab.

The following report is displayed:

Application Response						
Application Name	RTT (ms)	Network (ms)	Server (ms)	Transaction Delay (ms)	Transaction Count	Flows
Unclassified	420.80	3884.29	49.49	3905.79	2	102
HTTP	347.63	170.15	435.18	417.29	3	6
Replify	287.65	2403.08	0.00	2398.00	3	1
IMAP-SSL	206.36	558.94	37.31	520.19	35	1
HTTPS	182.18	55.04	1.10	55.68	7	9
slm	0.00	0.00	0.00	0.00	0	1
DNS	0.00	0.00	0.00	0.00	0	30
ssdp	0.00	0.00	0.00	0.00	0	2
ExindaCom	0.00	0.00	0.00	0.00	0	6
ICMP	0.00	0.00	0.00	0.00	0	13
PPTP	0.00	0.00	0.00	0.00	0	1
mDNS	0.00	0.00	0.00	0.00	0	2
IGMP	0.00	0.00	0.00	0.00	0	19
NTP	0.00	0.00	0.00	0.00	0	1

2. To change how often the table is refreshed, select an **Auto-Refresh Rate** from the list.

Display the report in the Exinda CLI

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

4. To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

5. 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) # █
```

Monitor the real time TCP health

The Real Time Host Health report shows the Retransmitted Bytes, Aborted Connections, Refused Connections, Ignored Connections and Flow Count for each internal and external host monitored by the Exinda appliance.

Display the report in the Exinda Web UI

1. Click **Monitor > Real Time** and switch to the **Host Health** tab.

The report is displayed:

Health					
Internal IP	Retransmitted (bytes)	Aborted	Refused	Ignored	Flows
192.168.80.66	0	0	0	0	1
172.16.1.240	0	1	0	0	13
172.14.1.10	0	0	0	0	90
172.16.0.160	0	0	0	0	2
172.16.0.115	0	0	0	0	4
192.168.100.252	0	0	0	0	1
192.168.0.207	0	0	0	0	1
172.16.1.149	0	0	0	0	4
172.16.1.242	0	0	0	0	1
172.16.0.63	0	0	0	0	2
192.168.90.180	0	0	0	0	1
192.168.0.35	0	0	0	0	1
192.168.0.178	0	0	0	4	5
192.168.0.83	0	0	0	0	1
0.0.0.0	0	0	0	0	1
172.16.0.179	0	0	0	0	1
192.168.0.145	0	0	0	0	1
192.168.0.209	0	0	0	0	1
192.168.0.114	0	0	0	0	1
192.168.20.115	0	0	0	0	1
172.16.0.252	0	0	0	0	1
172.16.0.114	0	0	0	0	2
192.168.60.12	0	0	0	0	1
172.16.0.119	0	1	0	0	1
192.168.0.171	0	0	0	0	78
192.168.0.54	0	0	0	0	1

Health					
External IP	Retransmitted (bytes)	Aborted	Refused	Ignored	Flows
189.47.235.215	0	0	0	0	1
182.237.13.217	0	0	0	0	1
114.27.0.33	0	0	0	0	1
209.162.180.188	0	0	0	0	1
46.146.120.14	0	0	0	0	1
77.124.175.104	0	0	0	0	1
190.50.180.129	0	0	0	0	1
203.2.192.124	0	0	0	0	1
95.25.240.132	0	0	0	0	1
81.182.22.28	0	0	0	0	1
92.112.162.185	0	0	0	0	1
116.48.3.56	0	0	0	0	1
78.37.161.202	0	0	0	0	1
114.143.5.78	0	0	0	0	1
183.98.3.27	0	0	0	0	1
98.82.49.248	0	0	0	0	1
87.242.31.140	0	0	0	0	1
239.255.255.250	0	0	0	0	1
190.44.86.103	0	0	0	0	1
85.228.237.61	0	0	0	0	1
98.226.208.222	0	0	0	0	1
239.255.255.100	0	0	0	0	1
211.206.58.76	0	0	0	0	1
122.106.153.166	0	0	0	0	1
70.55.149.222	0	0	0	0	2
96.244.84.2	0	0	0	0	1

2. To change how often the table is refreshed, select an **Auto-Refresh Rate** from the list.

Display the report in the Exinda CLI

1. Click **Tools > Console**.
2. Type the appliance username and password at the prompts.
3. To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The `hostname #` prompt is displayed.

- To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The `hostname (config)#` prompt is displayed.

- 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
```

View 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.

- Click **Monitor > Applications** and switch to the **Groups** tab.
- To expose Round trip time, Normalized Delays, Transaction Delays, and Efficiency statistics for each Application Group, click **Show Details**.

Top 30 Inbound Application Groups					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
Web	96259	116.599	76.11	8118.82	645
Social Networking	27639	36.786	181.54	4231.32	81
Other	11319	1.801	0.78	7.70	626
Exinda	3239	1.741	4.89	40.82	48
Streaming	737	0.970	264.86	366.70	2
Software Updates	691	0.898	105.13	183.42	3
Mail	1056	0.506	3.95	106.94	25
File Services	711	0.173	3.46	37.98	19
Thin Client	491	0.113	1.90	6.61	3
Voice	410	0.052	0.42	4.64	43
Interactive	142	0.034	1.54	5.05	17

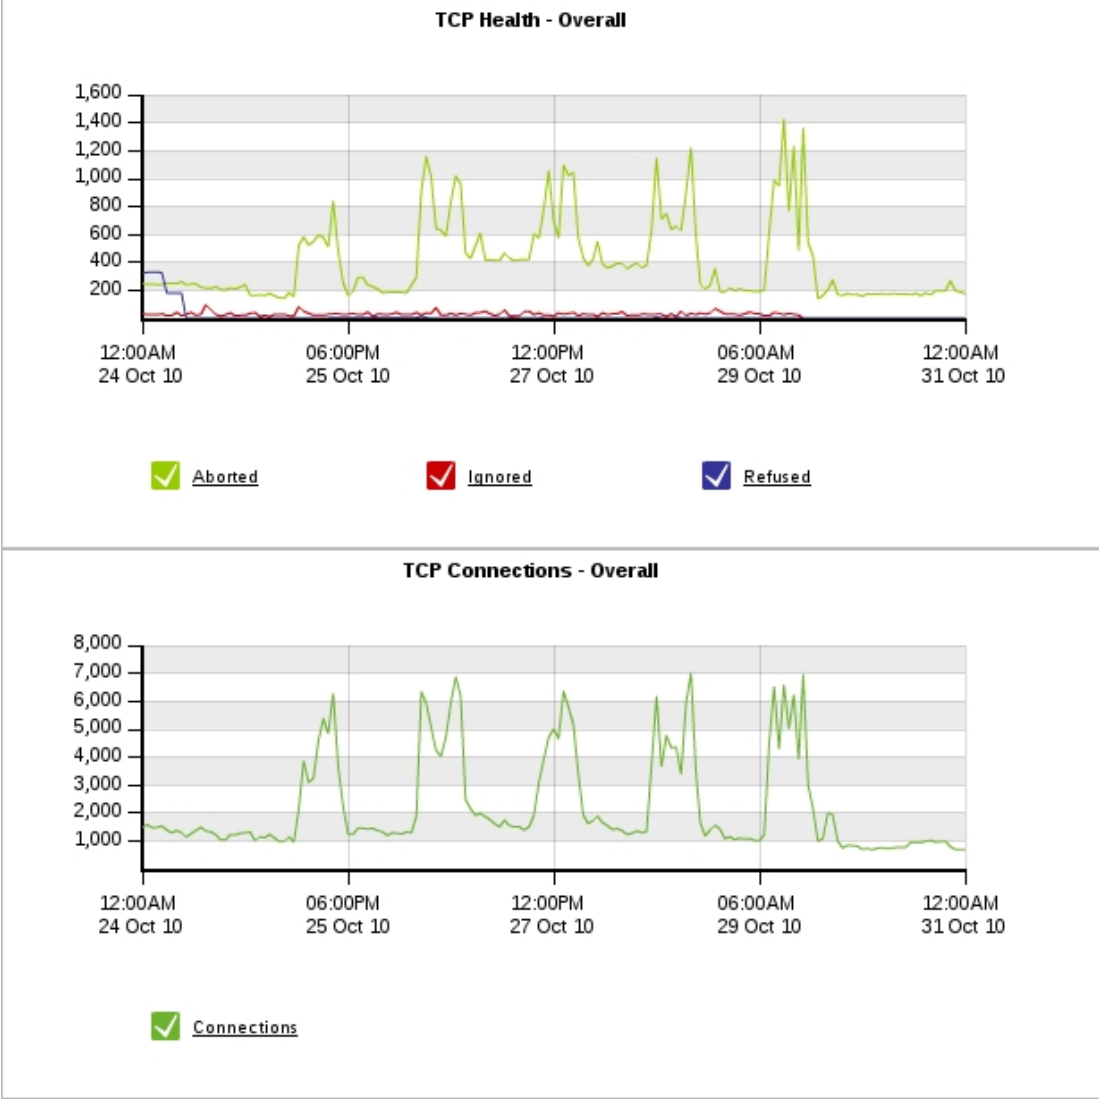
Top 30 Outbound Application Groups													
Name	Packets	Data (MB)	Throughput (kbps)		Flows	RTT (ms)	Normalized Delays (ms/kb)			Transaction Delays (ms)			Efficiency (%)
			Average	Max			Network	Server	Total	Network	Server	Total	
Web	53524	7.024	4.69	156.41	631	51	4280	309	4589	3942	717	4659	100.00
Social Networking	13556	1.267	6.25	79.05	81	81	290	32	322	1097	68	1165	100.00
Other	1949	0.362	0.92	9.58	159	48	511	341	852	1641	285	1926	100.00
Mail	802	0.362	2.96	4.28	22	163	5595	35	5630	24270	177	24447	98.84
File Services	753	0.167	3.10	23.26	22	29	9	6108	6117	5	1412	1417	100.00
Thin Client	581	0.065	1.09	2.14	3	92	11202	7284	18486	612	379	991	100.00
Voice	381	0.039	0.31	3.36	43	191	46216	1096018	1142234	24492	19440	43932	99.62
Streaming	371	0.027	7.31	10.35	2	18	10	503	513	13	698	711	100.00
Software Updates	399	0.026	3.00	4.77	3	0	1	13	14	30	9	39	100.00
Interactive	69	0.008	1.60	2.66	3	38	181	0	181	33	0	33	100.00

- To view the data for individual applications within a group, click the application group name.

View TCP efficiency

The TCP Efficiency Report shows the total efficiency of all TCP connections over time.

- Click **Monitor > Service Levels** switch to the **TCP Efficiency** tab.
- To display the total efficiency of all TCP connections over time,



Top 500 Applications				
	Connections	Aborted	Ignored	Refused
HTTP	327840	70627	126	4989
HTTPS	209951	59257	19	67

3. To change the time period that the report covers, select the **Range** from the list.
4. To view the data for individual applications within a group, click the application group name.

Create an Application Performance Score object

Create Application Performance Score (APS) objects for each of the important, business-critical applications running on the network to ensure that they are performing well. If an application reports a score below the acceptable threshold values, a notification is sent to the configured users.

1. Click **Objects > Service Levels** and switch to the **Application Performance Score** tab.
2. Click **Add New APS Object**.
3. Type a name for the score.
4. In the **Application** list, select the application traffic to monitor.
5. Specify the internal IP addresses that this APS object should measure. For example, specify ALL or a network object representing internal subnets.
6. Specify the external IP addresses that this APS object should measure. For example, specify ALL or a network object representing external application servers.
7. Select the **Alert Enable** checkbox.
8. In the **Alert Threshold** field type the maximum number of milliseconds (ms) allowed for the metric to complete before an alert is triggered.
9. In the **Alert Trigger Delay** list, select how long the metric needs to remain above the threshold before the alert is sent.
10. Set the thresholds for the APS metrics.
 - To generate a list of recommended thresholds for the APS metrics, see "[Generate recommended Application Performance Score thresholds](#)" on page 114.
 - To set specific scoring metrics without generating a baseline, deselect the Auto Baseline checkbox and enter the values for each metric.
11. 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.
12. Click **Add New APS Object**.

The object is added to the list of configured APS objects.
13. Click **System > Setup > Alerts**, and ensure that the Send Email and SNMP Trap alerts are enabled for APS.

Generate recommended Application Performance Score thresholds

Automatically create set of recommended threshold values for the Application Performance Score (APS) based on the traffic passing through the Exinda appliance over a specified period of time. If the configured thresholds are not performing as expected, run the auto baseline analysis to generate a new set of recommended values.

Note To ensure that the baseline values accurately reflect the usage of the application, start generating the baseline during a time period when users are actively generating traffic for the application.

1. Click **Objects > Service Levels** and switch to the **Application Performance Score** tab.
2. Locate the APS object in the list, and click **Edit**.

3. In the Baseline Info area, select the amount of time that the traffic will be analyzed to generate the APS baseline.

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.

4. Click **Start Baseline**.

Traffic is analyzed during the specified period, and a set of recommended 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.

To receive an email when the baseline period is over, ensure that the Exinda appliance is configured to send email notifications.

Note If no traffic is transferred during the selected auto baseline period, the Exinda appliance continues to analyze traffic for the next time period. For example, if no traffic is transferred during the one hour period, the traffic continues to be analyzed for one day. After the traffic has been analyzed for one week, and no traffic has been transferred, the auto baseline analysis stops.

Each time the auto baseline completes for a time period, and no traffic has been analyzed, an email notification is sent to the configured users.

Review and modify the APS threshold values

After the APS baseline analysis has completed, review the recommended values and adjust the threshold values as needed.

1. Click **Objects > Service Levels** and switch to the **Application Performance Score** tab.
2. Locate the APS object in the list, and click **Edit**.
3. In the Baseline Info area, ensure the status of the auto baseline is Completed.
4. In the Scoring Metrics area, review the recommended threshold values for the scoring metrics generated by the auto baseline analysis.

If the recommended threshold values are reasonable, click **Apply Changes** and proceed with Step 8.

Note The Network Loss metric is not calculated during the baseline analysis.

After the baseline analysis has successfully completed, the values in the Config column are editable.

5. Modify the threshold values as appropriate.
6. Click **Apply Changes**.
7. To write the changes to the configuration file, in the status bar click **Save**.

Config Status Unsaved changes (Save)

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. Click **Report** and switch to the **PDF Reports** tab.
2. Click **Add New PDF Report**.
3. In the Report Selection area select **APS, TCP Health**, and **TCP Efficiency**.
4. In the Report Details area, type a name for the report.
5. Specify how often the report will be generated.
6. Click **Add New Report**.
7. To generate the report, locate the report in the list and click **PDF**.

Chapter 10: Optimizer Configuration

The Optimizer delivers Quality of Service and Application Acceleration (x800 series only) mechanisms to improve application performance on the network. The intuitive, policy-based management helps match network behaviour to business objectives.

As an analogy, vehicle traffic on road systems can be dramatically slowed by disorganized control of traffic. Through intelligent traffic flow co-ordination techniques (such as allocating road lanes for particular types of vehicles), road traffic efficiency can be dramatically increased. Similarly, the efficiency of data flow on an IP network can be significantly improved using an intelligent optimization system powered by custom defined optimization rules.

As each network link has a fixed amount of bandwidth limiting the amount of data it may carry, optimum performance can be achieved by pre-sorting the data before it reaches the bottleneck links (typically the link between your network and the internet). Therefore "mission-critical" data is given priority over non-time sensitive data flow, such as SMTP mail or FTP traffic.

A Vehicle Traffic Analogy to Optimization

Using the vehicle analogy, if you have a number of people you need to get from point A to point B, via a freeway which passes over a bridge with a reduced number of lanes, the best way to get the most urgent passengers there first, is to:

- Sort the passengers into different vehicles so that the urgent passengers are grouped together in the fast cars;
- Identify these urgent car groups and allocate them to the first lanes;
- Similarly, sort the other groups of cars from fast (high priority to low priority);
- Give permission to certain slower groups to be allowed in the fast lanes, if and only, there is no fast car traffic;
- If there are certain groups which need a dedicated lane, then allocate the lane on the bridge as such.

By pre-determining these rules and queuing cars before they reach the bottleneck, it is possible to dramatically increase the traffic capacity of the road towards its theoretical maximum. To optimize the efficiency of your IP network, you need to carefully set the queuing or "optimizer" rules to meet the network requirements.

The optimizer provides:

- Rate shaping
- Traffic prioritization
- Application Acceleration
- Per Host QoS

- ToS/DSCP tagging
- Traffic blocking
- Time-based policies

Optimizer Policy Tree

Optimizer Policies are organized in a tree structure, and traffic is matched top-down, first by Virtual Circuits (VC), then by Policies. As traffic flows through the Exinda appliance, it is first matched to a Virtual Circuit (in order of Virtual Circuit number), then by a Policy within that Virtual Circuit (by order of Policy number).

Once traffic falls into a Virtual Circuit, it will never leave, so it must be captured by a policy within that Virtual Circuit. Similarly, once traffic is matched by a Policy, it never leaves. This means, more specific Virtual Circuits and Policies should be configured higher in order (towards the top) whereas more general Virtual Circuits and Policies should be last.

		Operations
Circuit 10 - Default (10000 kbps)		--Actions--
Virtual Circuit 10 - WAN inbound (10000 kbps from '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% (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:	Policy: ALL - Accelerate	Add To 'WAN inbound'
Create New Policy...		
Virtual Circuit 15 - WAN outbound (10000 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--
Order:	Policy: ALL - Accelerate	Add To 'WAN outbound'
Create New Policy...		
Create New Virtual Circuit...		
Create New Circuit...		

Circuits, Virtual Circuits and Policies can be manipulated by selecting various options from the drop-down to the right of the respective item.

Note System default Policies cannot be edited. Instead, they can be cloned, then customized.

The Optimizer service is controlled by clicking on the Optimizer links in the main toolbar.

Optimizer Status : On (Restart / Stop)

Note To view/change additional Optimizer configuration options, navigate to **System > System Setup > QoS Configuration** on the Web UI, advanced mode.

Create a new circuit

Circuits define physical connections to the WAN/Internet. If you have more than one bridge configured, you can bind different circuits to each bridge.

Typically, one Circuit would be created for each physical link to the WAN/Internet for which the Exinda appliance is placed in-line with. Circuits whose traffic can be filtered out by VCs, or by attaching it to a specific bridge, should be configured first; whereas the Internet Circuit (which contains a Catch All Virtual Circuit) must be configured last.

1. Click **Optimizer > Create New Circuit**.
2. In the **Number** field, type the priority order of the circuit, relative to other circuits.
3. Type a name for the circuit.
4. Identify the available **Inbound** and **Outbound** bandwidth of the circuit.

If the Circuit is synchronous, the inbound and outbound bandwidth values should be the same.

5. Specify the bridge or out-of-path interface to bind the Circuit to.
Bridge names, policy-based routing interfaces, and WCCP interfaces are available.
6. Click **Add New Circuit**.

Virtual Circuits

Virtual Circuits (VCs) are created within Circuits and are used to logically divide/partition the Circuit. For example, a Virtual Circuit may be configured for each branch office or, one Virtual Circuit for WAN data and one Virtual Circuit for Internet data. Each Virtual Circuit can contain different policies, which allows each Virtual Circuit to be treated differently.

Add New Virtual Circuit	
Virtual Circuit Number	10 . <input type="text" value="35"/>
Virtual Circuit Name	<input type="text"/>
Schedule	ALWAYS <input type="button" value="v"/>
Bandwidth Options	
Virtual Circuit Bandwidth	<input type="text"/> % <input type="button" value="v"/>
Oversubscription	<input checked="" type="radio"/> Automatic <input type="radio"/> Manual
Dynamic Virtual Circuit	<input type="checkbox"/>
Connection Options	
Connection Limit	<input type="text"/>
Filter Options	
VLAN Object	ALL <input type="button" value="v"/>
Network Object	---Network Objects--- <input type="button" value="v"/>
Application	ALL <input type="button" value="v"/>
Direction	Both <input type="button" value="v"/>
<input type="button" value="Add New Virtual Circuit"/> <input type="button" value="Cancel"/>	

Number	The order of the Virtual Circuit relative to other Virtual Circuits within the same Circuit.
Name	A logical name that represents the Virtual Circuit.
Schedule	A schedule that defines when the Virtual Circuit should be active.
Bandwidth	The maximum bandwidth used by the Virtual Circuit.
Oversubscription	Specify the behaviour when multiple Virtual Circuits are oversubscribed (the sum of the Virtual Circuit bandwidths exceed the parent Circuit bandwidth). The Exinda appliance can automatically calculate the oversubscription bandwidths or you can specify them manually. For more information, see the Virtual Circuit Oversubscription page.
Dynamic Virtual Circuit	This option allows you to specify per host QoS for each host that falls into this Virtual Circuit. For more information, see " Set a per-host limit on bandwidth usage " on page 123.

VLAN Object	VLAN tags or priorities used to filter out what traffic falls into the Virtual Circuit.
Network Object	Subnets/hosts/users/groups used to filter out what traffic falls into the Virtual Circuit.
Application	Application or Application Group used to filter out what traffic falls into the Virtual Circuit.
Direction	The traffic direction the Virtual Circuit should apply to.

The Filter Options are used to define what traffic falls into the Virtual Circuit. Filters can consist of VLAN Objects, Network Objects, Application/Application Groups and Direction. These filters are AND'd together.

Note A Virtual Circuit with a Network Object 'Email Server' and an Application Group 'Mail' will catch all Mail traffic to/from the Email Server.

Network Objects are typically used when Virtual Circuits are created for specific branch office locations or user group. Each branch office location or user group would be represented by a Static Network Object (typically 1 or more subnets) or a Dynamic Network Object (such as an Active Directory Group) and these Network Objects are defined in the Virtual Circuit. Only traffic to/from the subnets/users defined by the Network Object in the Virtual Circuit fall into the Virtual Circuit itself, all other traffic is evaluated by the next Virtual Circuit, and so on.

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. Another default Network Object, ALL, when used in a Virtual Circuit, will capture all traffic. This is called a Catch All Virtual Circuit and is typically the last (or only) configured Virtual Circuit.

The direction is used to ensure that the Virtual Circuit only captures traffic in a certain direction. This is useful for asynchronous circuits, as these generally require that at least 2 Virtual Circuits are defined, one for the inbound bandwidth and one for the outbound bandwidth.

If a network object is used in conjunction with a direction, the following rules apply:

Network Object	Direction	Captured Traffic
All	Both	All traffic, both inbound and outbound.
All	Inbound	All inbound traffic.
All	Outbound	All outbound traffic.
Not All	Both	Only inbound and outbound traffic to and from the subnets defined by the Network Object.
Not All	Inbound	Only inbound traffic to the subnets defined as 'internal' by the Network Object and from the subnets defined as 'external' by the Network Object.
Not All	Outbound	Only outbound traffic from the subnets defined as 'internal' by the Network

Network Object	Direction	Captured Traffic
		Object and to the subnets defined as 'external' by the Network Object.

Note The bandwidth used by a single Virtual Circuit must not exceed the parent Circuit bandwidth in either direction, but the sum of all Virtual Circuit bandwidths in either direction can exceed the parent Circuit bandwidth. This is called oversubscription and more information is available on the [Virtual Circuit Oversubscription](#) page.

Virtual Circuit Oversubscription

Over-subscription is where the sum of the Virtual Circuit bandwidths exceeds the sum of the parent Circuit bandwidth. This is the case in the following situation:

Circuit Bandwidth = 2Mbps

Virtual Circuit Bandwidth = 1Mbps

Virtual Circuit Bandwidth = 1Mbps

Virtual Circuit Bandwidth = 1Mbps

This means, the sum of the 3 Virtual Circuits is 3Mbps, but the Circuit bandwidth is only 2Mbps. The Virtual Circuits are oversubscribed.

There are 2 ways of dealing with oversubscribed Virtual Circuits, Automatic and Manual.

Oversubscription Automatic
 Manual

What happens in the **Automatic** case is that each Virtual Circuit will be guaranteed bandwidth on a pro-rata basis, 2/3Mbps (= .66Mbps) each and will be able to burst up to the specified bandwidth of 1Mbps (if this bandwidth is available). Under congestion (each Virtual Circuit under load) they will each share the Circuit's 2Mbps. If only 1 Virtual Circuit is under load, it will get the full 1Mbps. If only 2 Virtual Circuits are under load, they will both get 1Mbps.

The same pro-rata calculation is applied to all policies within an oversubscribed Virtual Circuit automatically.

In the **Manual** case you can specify a Guaranteed Bandwidth for each Virtual Circuit.

Oversubscription Automatic
 Manual

Guaranteed Bandwidth %

So using the example above, you can create the Virtual Circuits as follows:

Circuit Bandwidth = 2Mbps

Virtual Circuit Bandwidth = 1Mbps, Guarantee: 1Mbps

Virtual Circuit Bandwidth = 1Mbps, Guarantee: 512kbps

Virtual Circuit Bandwidth = 1Mbps, Guarantee: 512kbps

Note The sum of all the Virtual Circuit Guaranteed Bandwidths must not exceed the parent's Circuits bandwidth in both directions.

In this case, we've decided that the 1st Virtual Circuit will always receive the full 1Mbps if it needs it. The other 2 Virtual Circuits will receive 512kbps each. If the first Virtual Circuit is not using all of its 1Mbps, the other 2 Virtual Circuits can burst up to the full 1Mbps. The priority by which Virtual Circuits consume the burst bandwidth is determined by the priority of the policies within the Virtual Circuits.

Note There can be a mixture of Automatic and Manual Virtual Circuits in the same Circuit. The Guaranteed Bandwidths for each Manual Virtual Circuit are allocated first, then the Automatic calculations are made of the remaining bandwidth.

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. Click **Optimizer**.
2. Click **Create New Virtual Circuit**.
3. Type a name for the virtual circuit.
4. Type the amount of bandwidth to be used by the virtual circuit.
5. To enable Per Host QoS, select the **Dynamic Virtual Circuit** checkbox.
6. Set the amount of bandwidth (in KB per second or percentage of the virtual circuit bandwidth) that each host will receive in the **Per Host Bandwidth** field.

This bandwidth is guaranteed, so it will be available to each host, if required.

To have the amount of bandwidth each host receives calculated by dividing the Virtual Circuit guaranteed bandwidth by the number of active hosts, select **Automatically Share**.

7. Set the maximum amount of bandwidth (in KB per second or percentage of the virtual circuit bandwidth) that each host can burst to in the **Per Host Max Bandwidth** field.
If **No Bursting Allowed** is selected, each host only gets the bandwidth that they have been guaranteed.
8. Set the location of the hosts to allocate bandwidth to.
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.
9. Set the maximum number of hosts that can use the Dynamic Virtual Circuit.
If **Auto** is selected, the maximum number of hosts is calculated by assuming each host gets its guaranteed bandwidth.
If **Automatically Share** is selected, the maximum number of hosts is calculated by assuming each host is entitled to minimum bandwidth, which is 10kbps.

Any host that becomes active after the maximum number of hosts is exceeded do not fall into this Virtual Circuit.

- 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.

The following examples describe various Dynamic Virtual Circuit configurations.

<p>Name: Example 1</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: Auto</p> <p>Per User Max Bandwidth: 100%</p> <p>Host Location: Internal</p> <p>Max Hosts: Auto</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <p>If there is 1 user, they get the full 1024kbps.</p> <p>If there are 2 users, they each get 512kbps and can burst up to the full 1024kbps (if the other user is not using their guaranteed 512kbps).</p> <p>If there are 10 users, they each get 102kbps and can burst up to the full 1024kbps (if the other users are not using their guaranteed 102kbps).</p>
<p>Name: Example 2</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: 10%</p> <p>Per User Max Bandwidth: No</p> <p>Host Location: Internal</p> <p>Max Hosts: Auto</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p> <p>If there is 1 user, they get 102kbps and cannot burst.</p> <p>If there are 10 users, they each get 102kbps and cannot burst.</p> <p>If there are 100 users, the first 10 users each get 102kbps and cannot burst. The remaining 90 users will not match this VC.</p>
<p>Name: Example 3</p> <p>Bandwidth: 1024kbps</p>	<p>Internal Users is a Network Object that defines all hosts on the LAN side of the Exinda appliance.</p>

<p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: 64kbps</p> <p>Per User Max Bandwidth: 50%</p> <p>Host Location: Internal</p> <p>Max Hosts: 16</p>	<p>If there is 1 user, they get 64kbps and can burst up to 512kbps.</p> <p>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).</p> <p>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.</p>
<p>Name: Example 4</p> <p>Bandwidth: 1024kbps</p> <p>Direction: Both</p> <p>Network Object: Internal Users</p> <p>Application: Citrix</p> <p>Dynamic Virtual Circuits Enabled: Yes</p> <p>Per Host Bandwidth: 64kbps</p> <p>Per User Max Bandwidth: No</p> <p>Host Location: Internal</p> <p>Max Hosts: 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> <p>If there is 1 user, they get 64kbps for their Citrix traffic and cannot burst.</p> <p>If there are 16 users, they each get 64kbps for their Citrix traffic and cannot burst.</p> <p>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.</p>

Policies

Policies are used to filter specific traffic and apply one or more actions to that traffic. The following properties can be configured on all policies:

Policy Name:	<input type="text"/>
VC Policy Number:	<input type="text" value="110"/>
Schedule:	<input type="text" value="ALWAYS"/> ▼
Action:	<input type="text" value="Optimize"/> ▼
Policy Enabled:	<input checked="" type="checkbox"/>

Filter rule are used to define which traffic matches a policy. Use the form below to add or modify filter properties:

VLAN	Host	Direction	Host	ToS/DSCP	Application
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

VLAN	A pre-defined VLAN Object , which can select traffic based on 802.1Q VLAN ID and/or 802.1P VLAN priority tag.
Host / Direction / Host	A predefined Network Object (static or dynamic), which can select traffic to/from specific subnets or addresses.
ToS / DSCP 2	Select traffic based on DSCP/ToS marks in the IP header.
Application	Select traffic based on a predefined Application Object or Application Group .

One or all of these properties can be set to capture traffic for the policy. If each item is set to All , the policy will capture all traffic. This is called a Catch All Policy .

Note

- Selecting an entry in only one drop down, will default all the other items to All , for example, if you want to configure a Filter Rule for all HTTP traffic, you can select HTTP from the 'Application' drop down and save the policy without selecting anything else from the other drop downs.
- By default, 4 Filter Rules can be created per policy. If more are required, fill out the first 4, save the Policy, then edit the Policy and 4 more lines will become available.
- To delete individual Filter Rules, set all the fields for that Filter Rule to blank or use the [Optimizer | Policies](#) page on the Web UI, advanced mode.

If the Optimize action is selected, the following options will appear:

<input checked="" type="checkbox"/>	Guaranteed Bandwidth:	<input type="text"/>	%	<input type="text"/>
	Burst (Max) Bandwidth:	<input type="text"/>	%	<input type="text"/>
	Burst Priority:	<input type="text" value="1 (High)"/>		
<input checked="" type="checkbox"/>	Acceleration:	<input type="text" value="Acceleration"/>		
	WM Reduction Type:	<input type="text" value="Disk"/>		
<input checked="" type="checkbox"/>	ToS/DSCP Mark:	<input type="text"/>		
	VLAN Rewrite:	ID: <input type="text"/>	Priority: <input type="text"/>	

QoS	Allocate bandwidth to the policy so that QoS can be applied to any matching traffic. All bandwidth allocations are relative to the parent VC bandwidth. The following options are available as a percentage of the parent VC bandwidth or as an explicit value in kbps.
Guaranteed Bandwidth	<p>The minimum bandwidth made available for the policy. This is not reserved bandwidth if no traffic is matched to this policy, the guaranteed bandwidth is available for use by other policies.</p> <p>The guaranteed bandwidth of a single Policy must not exceed the parent Virtual Circuit bandwidth and the sum of all guaranteed bandwidths in each policy within a Virtual Circuit must not exceed the virtual Circuit bandwidth. In addition, the burst bandwidth must be greater than the guaranteed bandwidth, and less than or equal to the parent Virtual Circuit bandwidth.</p>
Burst (Max) Bandwidth	The maximum bandwidth made available for the policy. Once all guaranteed bandwidth is allocated, the remaining (excess) bandwidth is made available for burst usage. This value is the maximum burst bandwidth that will be allocated to a policy.
Burst Priority	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 for burst usage.
Acceleration	All traffic matching the policy will attempt to be accelerated. The following options define what kind of acceleration is applied to the matching traffic.
Acceleration	<p>Specify the type of acceleration to apply to the matching traffic. Available options are 'Acceleration' and 'Edge Cache'. 'Acceleration' is TCP-based Application Acceleration and is only available on x800 licensed appliances. Only outbound TCP traffic is accelerated.</p> <p>Edge Cache is Asymmetric Object Memory - a WAN reduction technology applied to outbound HTTP traffic. Edge Cache can be used with a single appliance.</p>
WM Reduction Type	Specify the type of WM reduction technology to apply to the matching traffic. If the 'Acceleration' option is selected above, the available options here are 'Disk', 'Compression' and 'None'. 'Disk' reduction uses the HDD in the Exinda appliance to store de-duplication patterns. It also uses 'Compression' reduction. 'Compression' reduction uses a network optimized LZ compression algorithm rather than de-duplication. 'None' means that the traffic will not attempted to be reduced but will still be accelerated.
Packet Marking	Mark individual packets matching this policy. The following packet marking options are available.
DSCP / ToS Mark	<p>Set a DSCP/ToS mark in the IP header of all packets matching the policy.</p> <p>For more information, refer to the DSCP / ToS How to Guide.</p>

VLAN Rewrite	Rewrite the 802.1Q VLAN ID and/or Priority only if an existing VLAN header is present. This is a packet based VLAN rewrite feature. Only packets matching this policy will be rewritten. Other packets that do not match this policy may be required to be rewritten in order for this feature to work (including non-IP packets such as ARP, which are not even processed by the Optimizer). Ensure that your topology supports this method of rewriting VLAN IDs before using this feature.
--------------	---

If the Ignore action is selected, traffic that matches the policy is still monitored, but will be ignored by the optimizer and pass through the appliance with no restriction on bandwidth. Typically this option is used for local traffic.

Note The "Ignore" option should not be used for policies within Dynamic Virtual Circuits.

If the "Discard" action is selected, traffic that matches the policy is dropped by the appliance. The following option is available:

Block Options: Discard only the first packet of a connection

The **Discard only the first packed of a connection** 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.

Note Any changes made to a Policy definition are global and will affect all Virtual Circuits that use that Policy.

Optimizer Policies

Polices can be created and edited independently of the Optimizer Policy Tree. The Optimizer | Policies page on the Web UI, advanced mode, lists all the system default and custom Policies. By default, there are no policies available. After running the Optimizer Wizard, a set of default policies will be installed.

The form at the top of the page allows you to create custom Policies.

Note See the [Optimizer Policy Tree | Policies](#) page for more information regarding creating custom Policies.

You can also use this page to delete individual Filter Rules from within Policies.

Note Any changes made to Policies will affect all instances of that policy if it is in use by more than one Virtual Circuit.

The following tables shows the policies that will be available to some of the default Application Groups after running the Optimizer wizard.

Num	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

The following matrix shows the combination of default Application Groups and policies (above) that will be made available as default policies.

Application Group	Policy Number												
	1	2	3	4	5	6	7	8	9	10	11	12	13
ALL		X					X	X					
Database						X	X	X				X	X
Files							X	X	X			X	X
Interactive	X							X	X	X			
Mail						X	X	X			X	X	
P2P			X	X	X								
Recreational			X	X	X	X	X						
Secure	X												

Application Group	Policy Number												
	1	2	3	4	5	6	7	8	9	10	11	12	13
Software Updates				X	X	X	X				X		
Streaming				X	X	X	X	X					
Thin Client	X							X	X	X			
Voice	X								X	X			
Web							X	X	X		X	X	X

Optimizer Wizard

The Optimizer Wizard is a convenient way to populate the Optimizer with some default Policies.

Caution Running the Optimizer Wizard will delete any existing Optimizer Policies and Optimizer Configuration.


The first 4 questions are always the same:

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.
Do you want to configure optimization policies?	Selecting YES will cause Questions 2 and 3 to appear.
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.
Do you want to apply QoS?	Selecting YES will apply traffic shaping.

Depending on your answers to Questions 2 and 3, the following scenarios are possible.

Scenario 1:


- Do you want to accelerate? **YES**
- Do you want to apply QoS? **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 optimization policies?	<input checked="" type="radio"/> Yes <input type="radio"/> No
Step 3: Do you want to accelerate? <i>Selecting YES will create policies that accelerate WAN applications. You must have another Exinda appliance on the WAN for this to work.</i>	<input checked="" type="radio"/> Yes <input type="radio"/> No
Step 4: Do you want to apply QoS? <i>Selecting YES will apply traffic shaping.</i>	<input checked="" type="radio"/> Yes <input type="radio"/> No
Step 5: 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 6: Enter inbound bandwidth (kbps)? <i>(MAX = 10240)</i>	<input type="text" value="10240"/> kbps
Step 7: Enter outbound bandwidth (kbps)? <i>(MAX = 10240)</i>	<input type="text" value="10240"/> kbps
Note that applying these settings will delete existing optimizer policies.	

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 enter the inbound and outbound bandwidths for this Exinda appliance.

Scenario 2:

- Do you want to accelerate? **NO**
- Do you want to apply QoS? **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 optimization policies? <input checked="" type="radio"/> Yes <input type="radio"/> No
Step 3:	Do you want to accelerate? <i>Selecting YES will create policies that accelerate WAN applications. You must have another Exinda appliance on the WAN for this to work.</i> <input type="radio"/> Yes <input checked="" type="radio"/> No
Step 4:	Do you want to apply QoS? <i>Selecting YES will apply traffic shaping.</i> <input checked="" type="radio"/> Yes <input type="radio"/> No
Step 5:	Do you want to apply an Enterprise or Service Provider QoS policy template? <i>Enterprise policies are strict in capping usage for P2P & recreational applications</i> <i>Service Provider policies are more generous with P2P & recreational traffic usage but these traffic groups are still bandwidth limited.</i> <input checked="" type="radio"/> Enterprise <input type="radio"/> Service Provider
Step 6:	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 7:	Enter inbound bandwidth (kbps) <i>(MAX = 10240)</i> <input type="text" value="10240"/> kbps
Step 8:	Enter outbound bandwidth (kbps) <i>(MAX = 10240)</i> <input type="text" value="10240"/> kbps
Note that applying these settings will delete existing optimizer policies.	

This will enable QoS (traffic shaping) only. You have the choice of the type of default Policy template to apply. There is a template better suited to Enterprise or one better suited to Service Providers. You will also need to select the WAN topology that best represents your deployment and also enter the inbound and outbound bandwidths for this Exinda appliance.

Scenario 3:

- Do you want to accelerate? **YES**
- Do you want to apply QoS? **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 optimization policies? <input checked="" type="radio"/> Yes <input type="radio"/> No
Step 3:	Do you want to accelerate? <i>Selecting YES will create policies that accelerate WAN applications. You must have another Exinda appliance on the WAN for this to work.</i> <input checked="" type="radio"/> Yes <input type="radio"/> No
Step 4:	Do you want to apply QoS? <i>Selecting YES will apply traffic shaping.</i> <input type="radio"/> Yes <input checked="" type="radio"/> No
Note that applying these settings will delete existing optimizer policies.	

This will enable Application Acceleration only.

ToS and DiffServ

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.

Used in conjunction with Exinda's superior classification techniques, including advanced layer 7 detection, users have complete control over how traffic is marked, and subsequently treated in the WAN cloud.

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 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; therefore, there are 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
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

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.

Match 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.

Add New VC Policy

Policy Name:

VC Policy Number:

Schedule:

Action:

Policy Enabled:

Guaranteed Bandwidth: %

Burst (Max) Bandwidth: %

Burst Priority:

Acceleration:

ToS/DSCP Mark:

VLAN	Host	Direction	Host	ToS/DSCP	Application
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text" value="< - >"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

Figure 1: 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.

Mark 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 will then be 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