

Clustering and HA


Exinda ExOS Version 6.3

© 2012 Exinda, Inc

Table of Contents

Part I Introduction	4
1 Using this Guide	4
2 Further Reading	5
Part II Overview	7
1 Terminology	7
Part III Getting Started	10
Part IV Cluster Configuration (Web UI)	12
Part V Cluster Configuration (CLI)	16
Part VI Cluster Operation	19
Part VII Cluster Failover	22

Part



1 Introduction

Clustering and HA

Exinda Firmware Version: 6.3

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.

1.1 Using this Guide

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

- Useful features, hints and important issues are called "notes" and they are identified in a light blue background.

Note: This is a note.

- Practical examples are presented throughout the manual for deeper understanding of specific concepts. These are called "examples" and are identified with a light green background.

This is an example.

- Warnings that can cause damage to the device are included when necessary. These are indicated by the word "caution" and are highlighted in yellow.

Caution: This is a caution.


1.2 Further Reading

In addition to this How to Guide, the following relevant user documentation is available and should be read in conjunction with this guide:

- Exinda User Manual
- Exinda Topologies Guide

Please visit <http://www.exinda.com> for more information.

Part



2 Overview

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.




Figure 1: A typical clustered Exinda deployment.

In this example, there are 2 physical links. An Exinda appliance is deployed between each switch and router, and a cable is connected between the 2 appliances for synchronization.


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

2.1 Terminology

- **Cluster Node:** An Exinda appliance that is a member of a cluster.
- **Cluster:** A group of Exinda appliances (cluster nodes) configured to operate as a single Exinda appliance.
- **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 Master:** The node responsible for synchronizing configuration changes with all other cluster nodes. Configuration changes should only be made from the cluster master.

- **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 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 it's management port.

Part



3 Getting Started

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


Table 1: Suggested cluster interface for each hardware series.

Where there are 2 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 2 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).

Note: Clustering is not available on the 1010, 1060 or 2060 Hardware Series.

Part

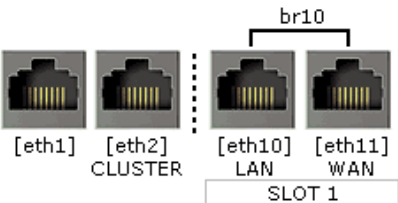


4 Cluster Configuration (Web UI)

Note: To configure clustering, navigate to the System | Network | IP Address on the Web UI, advanced mode.

The following example shows a typical cluster configuration on 2 appliances.

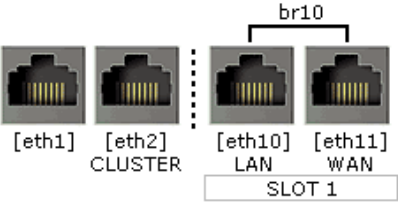
IP Address Configuration on Exinda 1:



Interface Settings	
eth1	Role: <input type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP
	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC
	Dynamic Addresses: fe80::224:e8ff:fe69:ef23/64
	Static Addresses: <input type="text" value="192.168.0.161"/> / <input type="text" value="24"/>
Comment: <input type="text" value="Management Interface"/>	
eth2	Role: <input checked="" type="checkbox"/> Cluster <input type="checkbox"/> Mirror <input type="checkbox"/> WCCP
	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC
	Dynamic Addresses: fe80::224:e8ff:fe69:ef24/64
	Static Addresses: <input type="text" value="192.168.1.1"/> / <input type="text" value="24"/>
Comment: <input type="text" value="Cluster Internal"/>	
br10 <input checked="" type="checkbox"/>	Autoconf: IPv4: <input type="checkbox"/> DHCP IPv6: <input type="checkbox"/> SLAAC
	Dynamic Addresses: fe80::2e0:edff:fe14:5e32/64
	Static Addresses: <input type="text"/> / <input type="text"/>
Comment: <input type="text" value="Bridge interface"/>	
Gateway Settings	
IPv4:	<input type="text" value="192.168.0.1"/>
IPv6:	<input type="text"/>
<input type="button" value="Apply Changes"/>	

Figure 2: IP Address configuration page on Exinda 1.

IP Address Configuration on Exinda 2:



Interface Settings

eth1	Role:	<input type="checkbox"/> Cluster	<input type="checkbox"/> Mirror	<input type="checkbox"/> WCCP
	Autoconf:	IPv4: <input type="checkbox"/> DHCP	IPv6: <input type="checkbox"/> SLAAC	
	Dynamic Addresses:	fe80::224:e8ff:fe69:ef23/64		
	Static Addresses:	192.168.0.162		/ 24
	Comment:	Management Interface		
eth2	Role:	<input checked="" type="checkbox"/> Cluster	<input type="checkbox"/> Mirror	<input type="checkbox"/> WCCP
	Autoconf:	IPv4: <input type="checkbox"/> DHCP	IPv6: <input type="checkbox"/> SLAAC	
	Dynamic Addresses:	fe80::224:e8ff:fe69:ef24/64		
	Static Addresses:	192.168.1.2		/ 24
	Comment:	Cluster Internal		
br10	Autoconf:	IPv4: <input type="checkbox"/> DHCP	IPv6: <input type="checkbox"/> SLAAC	
	Dynamic Addresses:	fe80::2e0:edff:fe14:5e32/64		
	Static Addresses:			/
	Comment:	Bridge interface		

Gateway Settings

IPv4: 192.168.0.1

IPv6:

Figure 3: IP Address configuration page on Exinda 2.

Cluster Master Settings

Interface: eth1

Master Address: 192.168.0.160 / 24

Figure 4: Cluster External (Master) configuration on Exinda 1 and Exinda 2

In the example above, Exinda 1 has a Management IP of 192.168.0.161 and Exinda 2 has a Management IP of 192.168.0.162. The Cluster External IP is configured as 192.168.0.160 on both appliances – regardless of which of these 2 appliances becomes the Cluster Master, it will be reachable on the 192.168.0.160 IP address.

The Cluster Internal IP on Exinda 1 is configured as 192.168.1.1 and on Exinda 2 as 192.168.1.2. 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.

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's best to use the Cluster Master IP when managing a cluster.

When logged into the Web UI of a cluster node, the role of the node will be shown in the header of the user interface as shown below.

Figure 5: The cluster role is displayed in the header when logged into the Web UI.


The System | Maintenance | Clustering page on the Advanced Web UI will show the status of all the nodes in the cluster, as shown below.

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

Figure 6: The status of all the nodes in the cluster.

It is also possible to reboot and shutdown other nodes in the cluster from this page.

Part



5 Cluster Configuration (CLI)

Configuration using the CLI is very similar to that of the Web UI.

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.

Next, 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.

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

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

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
```

Terminology

ID	The node's cluster assigned unique identifier.
Role	The current 'role' of node within the cluster (master or standby).
State	The node state (online or offline)
Host	The node's hostname.

External Address	The nodes's External (management) IP address.
Internal Address	The node's cluster Internal IP address.
Management IP	The clusters management (alias) address. The cluster is always reachable at this address as long as at least one node is online.

Part



VI

6 Cluster Operation

As part of normal cluster operations, the Cluster Master will synchronize parts of the system configuration to all other nodes in the cluster. Some configuration is specific to an individual appliance (e.g. IP addressing, 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|monitor|optimizer}
```

all

Both monitor and optimizer data are synchronized

monitor

synchronize monitor data only

optimizer

synchronize optimizer data only

Part



7 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.