

Load Balancing HPE Ezmeral Data Fabric

Version 1.1



Table of Contents

1. About this Guide	3
2. Loadbalancer.org Appliances Supported	3
3. Software Versions Supported	3
3.1. Loadbalancer.org Appliance	3
3.2. HPE Ezmeral Data Fabric	3
4. HPE Ezmeral Data Fabric	3
4.1. HPE Ezmeral Data Fabric Data Stores	3
5. Load Balancing HPE Ezmeral Data Fabric	4
5.1. Services to be Load Balanced	4
5.2. Virtual Service (VIP) Requirements	4
6. Deployment Concept	5
7. Load Balancer Deployment Methods	5
7.1. Layer 7 SNAT Mode	6
8. Loadbalancer.org Appliance – the Basics	6
8.1. Virtual Appliance	7
8.2. Initial Network Configuration	7
8.3. Accessing the Appliance WebUI	7
Main Menu Options	8
8.4. Appliance Software Update	9
Determining the Current Software Version	9
Checking for Updates using Online Update	9
Using Offline Update	10
8.5. Ports Used by the Appliance	10
8.6. HA Clustered Pair Configuration	11
9. Appliance Configuration for HPE Ezmeral Data Fabric	11
9.1. VIP 1 - DAG	11
Virtual Service (VIP) Configuration	11
Defining the Real Servers (RIPs)	12
9.2. VIP 2 - MOSS	12
Virtual Service (VIP) Configuration	12
Defining the Real Servers (RIPs)	13
9.3. Finalizing the Configuration	13
10. Testing & Verification	14
10.1. Accessing the DAGs and MOSS / Object Storage Nodes	14
10.2. Using System Overview	14
11. Technical Support	14
12. Further Documentation	14
13. Appendix	15
13.1. Configuring HA - Adding a Secondary Appliance	15
Non-Replicated Settings	15
Configuring the HA Clustered Pair	16
14. Document Revision History	18

1. About this Guide

This guide details the steps required to configure a load balanced HPE Ezmeral Data Fabric environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any HPE Ezmeral Data Fabric configuration changes that are required to enable load balancing.

For more information about initial appliance deployment, network configuration and using the Web User Interface (WebUI), please also refer to the [Administration Manual](#).

2. Loadbalancer.org Appliances Supported

All our products can be used with HPE Ezmeral Data Fabric. For full specifications of available models please refer to <https://www.loadbalancer.org/products>.

Some features may not be available or fully supported in all cloud platforms due to platform specific limitations. For more details, please refer to the "Main Differences to our Standard (Non-Cloud) Product" section in the appropriate cloud platform [Quick Start Guide](#) or check with Loadbalancer.org support.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.9.1 and later

Note

The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If you're using an older version, or the very latest, the screenshots presented here may not match your WebUI exactly.

3.2. HPE Ezmeral Data Fabric

- v7.0.0 and later

4. HPE Ezmeral Data Fabric

HPE Ezmeral Data Fabric is the industry-leading data platform for AI and analytics that solves enterprise business needs. Data fabric is a design concept that serves as an integrated layer (fabric) of data and connecting processes. It accelerates insights by automating ingestion, curation, discovery, preparation, and integration across data silos.

HPE Ezmeral Data Fabric offers cross-protocol data access via native S3, NFS, POSIX, REST, HDFS, and container storage interface (CSI) APIs. HPE Ezmeral Data Fabric emerges as a comprehensive solution, facilitating seamless integration, accessibility, and management of diverse data sources for organizations seeking enhanced efficiency and agility in their analytics processes.

4.1. HPE Ezmeral Data Fabric Data Stores

HPE Ezmeral Data Fabric File Store - HPE Ezmeral Data Fabric File Store is a distributed file system for data



storage, data management, and data protection. File Store supports mounting and cluster access via NFS and FUSE-based POSIX clients (basic, platinum, or PACC) and also supports access and management via HDFS APIs.

HPE Ezmeral Data Fabric Object Store - The HPE Ezmeral Data Fabric Object Store is a native object storage solution that efficiently stores objects and metadata for optimized access.

HPE Ezmeral Data Fabric Database - HPE Ezmeral Data Fabric Database is an enterprise-grade, high-performance, NoSQL database management system that you can use for real-time, operational analytics.

HPE Ezmeral Data Fabric Streams - HPE Ezmeral Data Fabric Streams brings integrated publish and subscribe messaging to the Data Fabric Converged Data Platform.

5. Load Balancing HPE Ezmeral Data Fabric

Note

It's highly recommended that you have a working HPE Ezmeral Data Fabric environment first before implementing the load balancer.

5.1. Services to be Load Balanced

This guide provides details on how to load balance the following:

Data Access Gateways - The HPE Ezmeral Data Fabric Data Access Gateway is a service that acts as a proxy and gateway for translating requests between lightweight client applications and the HPE Ezmeral Data Fabric cluster. The load balancer is used to distribute client requests across multiple Data Access Gateways.

Multithreaded Object Store (MOS) Servers - MOS Servers are high performance stateless gateways that accept S3 commands and translate them into low-level data fabric operations that are passed on to the storage nodes. The load balancer is used to distribute the object store load across multiple MOS servers.

5.2. Virtual Service (VIP) Requirements

To provide load balancing and HA for HPE Ezmeral Data Fabric, the following VIPs are required:

Ref.	VIP Name	Mode	VIP Port(s)	Real Server Port	Persistence Mode	Health Check
VIP 1	DAG	L7 SNAT (TCP)	443	5678	Source IP	Connect to Port
VIP 2	MOSS	L7 SNAT (TCP)	443	9000	None	Connect to Port

VIP 1 Notes

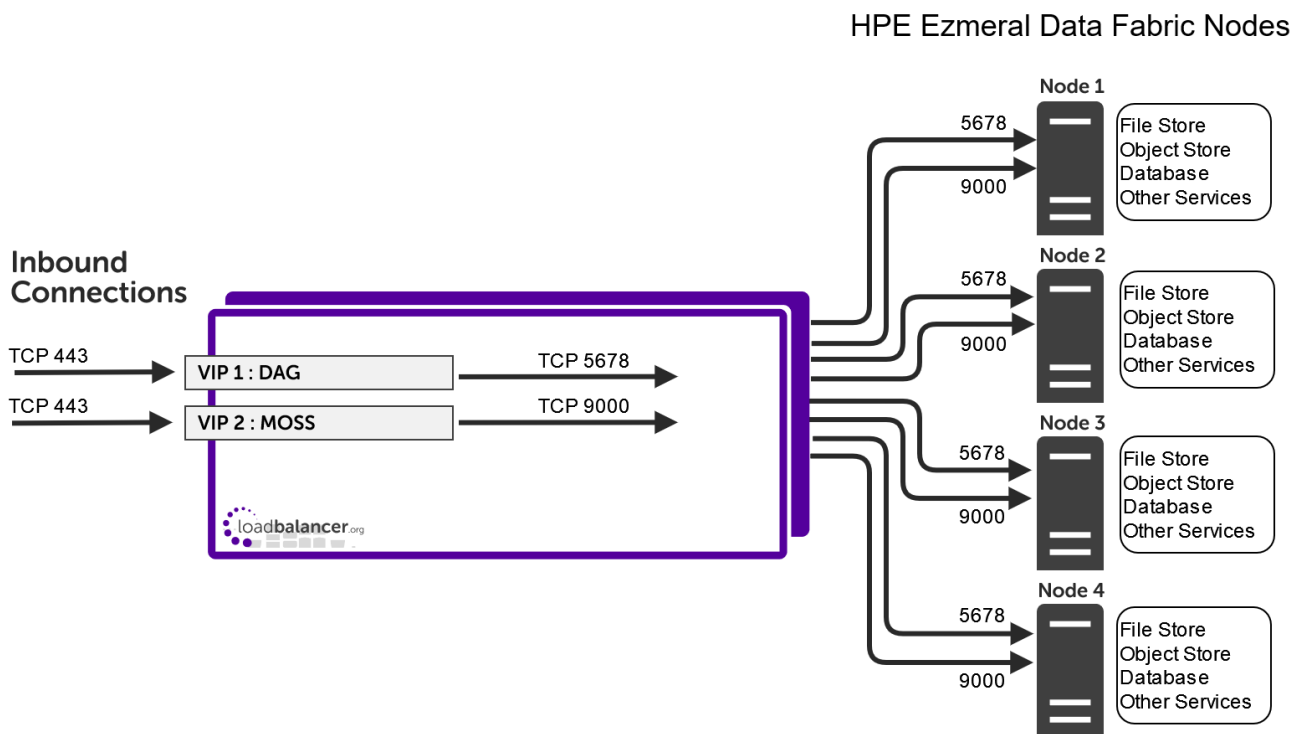
1. The VIP's port can be set to any convenient value, 443 is the suggested value and is used in this guide.
2. The default Real Server port for Data Access Gateway is 5678. This can be changed if required by modifying the `grpc.service.port` parameter in the `properties.cfg` file.
3. HTTPS access for Data Access Gateway is enabled by default. This is controlled by the `grpc.service.ssl.enabled` parameter in the `properties.cfg` file. For more details see [Administering the Data Access Gateway](#).



VIP 2 Notes

1. The VIP's port can be set to any convenient value, 443 is the suggested value and is used in this guide.
2. The default Real Server port for MOS Server is 9000. This can be changed if required by modifying the `moss.port=<port_number>` parameter in the `moss.conf` file.
3. HTTPS access for MOS Server must be configured. For more details see [HTTPS Access to Object Store](#).

6. Deployment Concept



VIP = **V**irtual **I**P Address

Note

The lab environment used to write this guide was based on the 4-node configuration shown above. For more information on cluster design considerations for production environments, refer to [Planning the Cluster](#).

Note

The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to the section [Configuring HA - Adding a Secondary Appliance](#) in the appendix for more details on configuring a clustered pair.

7. Load Balancer Deployment Methods

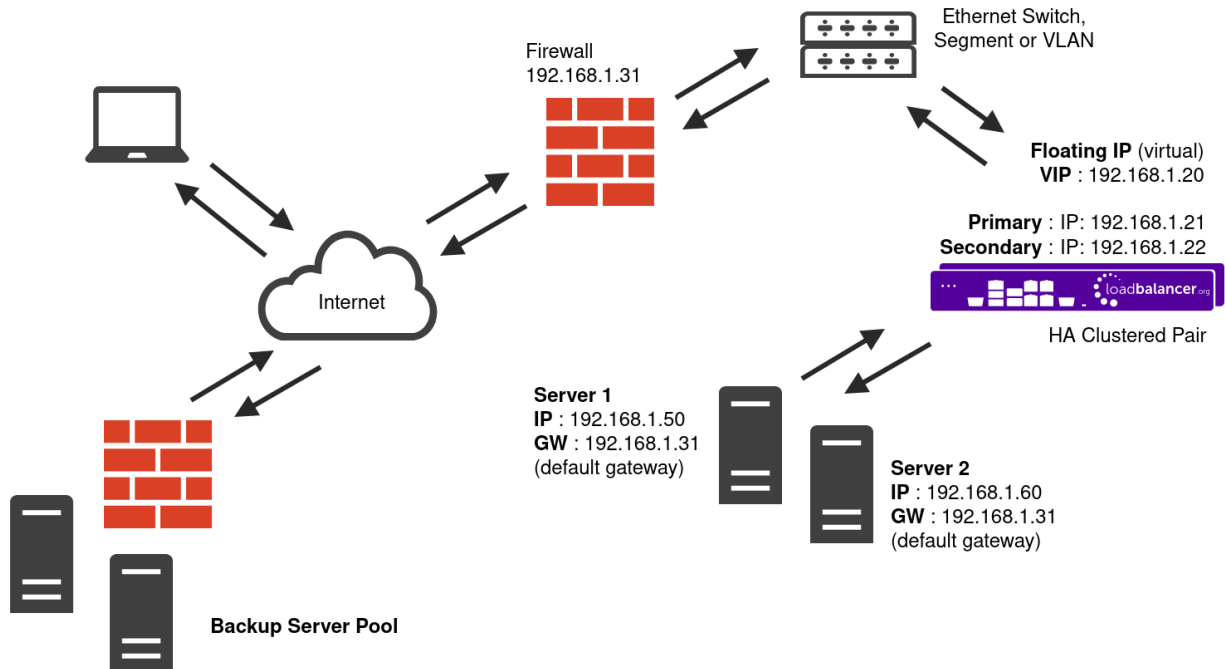
The load balancer can be deployed in 4 fundamental ways: *Layer 4 DR mode*, *Layer 4 NAT mode*, *Layer 4 SNAT mode*, and *Layer 7 SNAT mode*.

For HPE Ezmeral Data Fabric, layer 7 SNAT mode is recommended. This mode is described below and is used for the configuration presented in this guide.



7.1. Layer 7 SNAT Mode

Layer 7 SNAT mode uses a proxy (HAProxy) at the application layer. Inbound requests are terminated on the load balancer and HAProxy generates a new corresponding request to the chosen Real Server. As a result, Layer 7 is typically not as fast as the Layer 4 methods. Layer 7 is typically chosen when either enhanced options such as SSL termination, cookie based persistence, URL rewriting, header insertion/deletion etc. are required, or when the network topology prohibits the use of the layer 4 methods.



- Because layer 7 SNAT mode is a full proxy, any server in the cluster can be on any accessible subnet including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's own IP address by default, or any other local appliance IP address if preferred (e.g. the VIP address). This can be configured per layer 7 VIP. If required, the load balancer can be configured to provide the actual client IP address to the Real Servers in 2 ways. Either by inserting a header that contains the client's source IP address, or by modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. For more information on these methods please refer to [Transparency at Layer 7](#).
- Layer 7 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, **eth0** is normally used for the internal network and **eth1** is used for the external network although this is not mandatory.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

8. Loadbalancer.org Appliance – the Basics



8.1. Virtual Appliance

A fully featured, fully supported 30 day trial is available if you are conducting a PoC (Proof of Concept) deployment. The VA is currently available for VMware, Virtual Box, Hyper-V, KVM, XEN and Nutanix AHV and has been optimized for each Hypervisor. By default, the VA is allocated 2 vCPUs, 4GB of RAM and has a 20GB virtual disk. The Virtual Appliance can be downloaded [here](#).

Note

The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.

Note

Please refer to [Virtual Appliance Installation](#) and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.

Note

The VA has 4 network adapters. For VMware only the first adapter (**eth0**) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

8.2. Initial Network Configuration

After boot up, follow the instructions on the appliance console to configure the management IP address, subnet mask, default gateway, DNS servers and other network and administrative settings.

Important

Be sure to set a secure password for the load balancer, when prompted during the setup routine.

8.3. Accessing the Appliance WebUI

The WebUI is accessed using a web browser. By default, users are authenticated using Apache authentication. Users can also be authenticated against LDAP, LDAPS, Active Directory or Radius - for more information, please refer to [External Authentication](#).

Note

There are certain differences when accessing the WebUI for the cloud appliances. For details, please refer to the relevant [Quick Start / Configuration Guide](#).

1. Using a browser, navigate to the following URL:

<https://<IP-address-configured-during-the-network-setup-wizard>:9443/lbadmin/>

Note

You'll receive a warning about the WebUI's SSL certificate. This is due to the default self signed certificate that is used. If preferred, you can upload your own certificate - for more information, please refer to [Appliance Security Features](#).

Note

If you need to change the port, IP address or protocol that the WebUI listens on, please refer to [Service Socket Addresses](#).



2. Log in to the WebUI using the following credentials:

Username: loadbalancer

Password: <configured-during-network-setup-wizard>

Note To change the password, use the WebUI menu option: *Maintenance > Passwords*.

Once logged in, the WebUI will be displayed as shown below:

The screenshot displays the Loadbalancer WebUI interface. At the top left is the 'LOADBALANCER' logo. At the top right, there is a 'Enterprise VA Max' badge. Below the logo, there are status indicators: 'Primary | Secondary', 'Active | Passive', and 'Link'. A refresh icon and '15 Seconds' are also present. On the left side, there is a vertical navigation menu with options: System Overview, Local Configuration, Cluster Configuration, Maintenance, View Configuration, Reports, Logs, Support, and Live Chat. The main content area features a warning box: 'WARNING: YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS. Buy with confidence. All purchases come with a 90 day money back guarantee. Already bought? Enter your license key here' with a 'Buy Now' button. Below this is the 'System Overview' section, which includes a question 'Would you like to run the Setup Wizard?' with 'Accept' and 'Dismiss' buttons. A table header for 'VIRTUAL SERVICE' is visible, with columns for IP, PORTS, CONNS, PROTOCOL, METHOD, and MODE. Below the table, it states 'No Virtual Services configured.' Two line graphs are shown: 'Network Bandwidth' and 'System Load Average'. The Network Bandwidth graph shows RX (3k Min, 4k Avg, 32675k Total) and TX (6k Min, 7k Avg, 56693k Total) data. The System Load Average graph shows 1m, 5m, and 15m averages with their respective min, avg, and max values. A 'Memory Usage' graph is partially visible at the bottom.

3. You'll be asked if you want to run the Setup Wizard which can be used to configure layer 7 services. Click **Dismiss** if you're following a guide or want to configure the appliance manually or click **Accept** to start the wizard.

Main Menu Options



System Overview - Displays a graphical summary of all VIPs, RIPs and key appliance statistics

Local Configuration - Configure local host settings such as IP address, DNS, system time etc.

Cluster Configuration - Configure load balanced services such as VIPs & RIPs

Maintenance - Perform maintenance tasks such as service restarts and taking backups

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

Support - Create a support download, contact the support team & access useful links

Live Chat - Start a live chat session with one of our Support Engineers

8.4. Appliance Software Update

To ensure that the appliance(s) are running the latest software version, we recommend a software update check is performed.

Determining the Current Software Version

The software version is displayed at the bottom of the WebUI as shown in the example below:

Copyright © Loadbalancer.org Inc. 2002 – 2024
ENTERPRISE VA Max - v8.11.1

English ▼

Checking for Updates using Online Update

Note

By default, the appliance periodically contacts the Loadbalancer.org update server and checks for updates. An update check can also be manually triggered as detailed below.

1. Using the WebUI, navigate to: **Maintenance > Software Update**.
2. Select **Online Update**.
3. If the latest version is already installed, a message similar to the following will be displayed:

Information: Version v8.11.2 is the current release. No updates are available

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

Note

Do not navigate away whilst the update is ongoing, this may cause the update to fail.

6. Once complete (the update can take several minutes depending on download speed and upgrade version) the following message will be displayed:



Information: Update completed successfully.

- If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

Using Offline Update

If the load balancer does not have access to the Internet, offline update can be used.

Note

Please contact support@loadbalancer.org to check if an update is available and obtain the latest offline update files.

To perform an offline update:

- Using the WebUI, navigate to: **Maintenance > Software Update**.
- Select **Offline Update**.
- The following screen will be displayed:

Software Update

Offline Update

The following steps will lead you through offline update.

- Contact **Loadbalancer.org support** to obtain the offline update archive and checksum.
- Save the archive and checksum to your local machine.
- Select the archive and checksum files in the upload form below.
- Click *Upload and Install* to begin the update process.

Archive: No file chosen

Checksum: No file chosen

- Select the *Archive* and *Checksum* files.
- Click **Upload and Install**.
- If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

8.5. Ports Used by the Appliance

By default, the appliance uses the following TCP & UDP ports:

Protocol	Port	Purpose
TCP	22 *	SSH
TCP & UDP	53 *	DNS / GSLB
TCP & UDP	123	NTP



Protocol	Port	Purpose
TCP & UDP	161 *	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000 *	Gateway service (Centralized/Portal Management)
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

Note

The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to [Service Socket Addresses](#).

8.6. HA Clustered Pair Configuration

Loadbalancer.org recommend that load balancer appliances are deployed in pairs for high availability. In this guide a single unit is deployed first, adding a secondary unit is covered in the section [Configuring HA - Adding a Secondary Appliance](#) of the appendix.

9. Appliance Configuration for HPE Ezmeral Data Fabric

9.1. VIP 1 - DAG

Virtual Service (VIP) Configuration

- Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.

Virtual Service
[Advanced +]

Label ?

IP Address ?

Ports ?

Protocol
[Advanced +]

Layer 7 Protocol ?

Cancel
Update

- Define the *Label* for the Virtual Service as required, e.g. **DAG**.
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.110.100**.



4. Set the *Ports* field to the port that clients will connect to when accessing the service, e.g. **443**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the Virtual Service.

Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.

Label	<input type="text" value="DAG1"/>	?
Real Server IP Address	<input type="text" value="192.168.110.110"/>	?
Real Server Port	<input type="text" value="5678"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

Cancel
Update

2. Define the *Label* for the Real Server as required, e.g. **DAG1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.110.110**.
4. Set the *Real Server Port* field to **9000**.
5. Click **Update**.
6. Repeat these steps to add the remaining DAG nodes.

9.2. VIP 2 - MOSS

Virtual Service (VIP) Configuration

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.

Virtual Service		[Advanced +]
Label	<input type="text" value="MOSS"/>	?
IP Address	<input type="text" value="192.168.110.101"/>	?
Ports	<input type="text" value="443"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input style="border: 1px solid #ccc;" type="text" value="TCP Mode"/>	?

Cancel
Update

2. Define the *Label* for the Virtual Service as required, e.g. **MOSS**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.110.101**.
4. Set the *Ports* field to the port that clients will connect to when accessing the service, e.g. **9000**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the Virtual Service.
7. Now click **Modify** next to the newly created VIP.
8. Scroll to the *Persistence* section.
 - Set *Persistence Mode* to **None**.
9. Click **Update**.

Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.

Label	<input type="text" value="MOSS1"/>	?
Real Server IP Address	<input type="text" value="192.168.110.110"/>	?
Real Server Port	<input type="text" value="9000"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

2. Define the *Label* for the Real Server as required, e.g. **MOSS1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.110.110**.
4. Set the *Real Server Port* field to **9000**.
5. Click **Update**.
6. Repeat these steps to add the remaining MOSS nodes.

9.3. Finalizing the Configuration

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the "Commit changes" box at the top of the screen or by using the *Restart Services* menu option:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.

10. Testing & Verification

 **Note**


For additional guidance on diagnosing and resolving any issues you may have, please also refer to [Diagnostics & Troubleshooting](#).








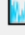







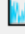




10.1. Accessing the DAGs and MOSS / Object Storage Nodes

First ensure that any DNS records that are used for access are updated so that the FQDNs resolve to the relevant VIP. Then verify that you're able to successfully access the load balanced services.

10.2. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the HPE Ezmeral Data Fabric nodes) and shows the state/health of each server as well as the state of the cluster as a whole. The example below shows that all nodes are healthy and available to accept connections:

System Overview  2024-04-24 14:25:27 UTC

	VIRTUAL SERVICE	IP	PORTS	CONN	PROTOCOL	METHOD	MODE	
	DAG	192.168.110.100	443	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONN			
	DAG1	192.168.110.110	5678	100	0	Drain	Halt	
	DAG2	192.168.110.111	5678	100	0	Drain	Halt	
	DAG3	192.168.110.112	5678	100	0	Drain	Halt	
	DAG4	192.168.110.113	5678	100	0	Drain	Halt	
	MOSS	192.168.110.101	443	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONN			
	MOSS1	192.168.110.110	9000	100	0	Drain	Halt	
	MOSS2	192.168.110.111	9000	100	0	Drain	Halt	
	MOSS3	192.168.110.112	9000	100	0	Drain	Halt	
	MOSS4	192.168.110.113	9000	100	0	Drain	Halt	

11. Technical Support

For more details about configuring the appliance and assistance with designing your deployment please don't hesitate to contact the support team using the following email address: support@loadbalancer.org.

12. Further Documentation

For additional information, please refer to the [Administration Manual](#).



13. Appendix

13.1. Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution. We recommend that the Primary appliance is fully configured first, then the Secondary appliance can be added to create an HA pair. Once the HA pair is configured, load balanced services must be configured and modified on the Primary appliance. The Secondary appliance will be automatically kept in sync.

 **Note**

For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the [documentation library](#)

The clustered HA pair uses Heartbeat to determine the state of the other appliance. Should the active device (normally the Primary) suffer a failure, the passive device (normally the Secondary) will take over.

Non-Replicated Settings

A number of settings are not replicated as part of the Primary/Secondary pairing process and therefore must be manually configured on the Secondary appliance. These are listed by WebUI menu option in the table below:

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	Interface IP addresses, bonding configuration and VLANs
Local Configuration	Routing	Default gateways and static routes
Local Configuration	System Date & time	Time and date related settings
Local Configuration	Physical – Advanced Configuration	Various appliance settings
Local Configuration	Portal Management	Portal management settings
Local Configuration	Security	Security settings
Local Configuration	SNMP Configuration	SNMP settings
Local Configuration	Graphing	Graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Backup & Restore	Local XML backups
Maintenance	Software Updates	Appliance software updates
Maintenance	Fallback Page	Fallback page configuration
Maintenance	Firewall Script	Firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

⚠ Important

Make sure that where any of the above have been configured on the Primary appliance, they're also configured on the Secondary.

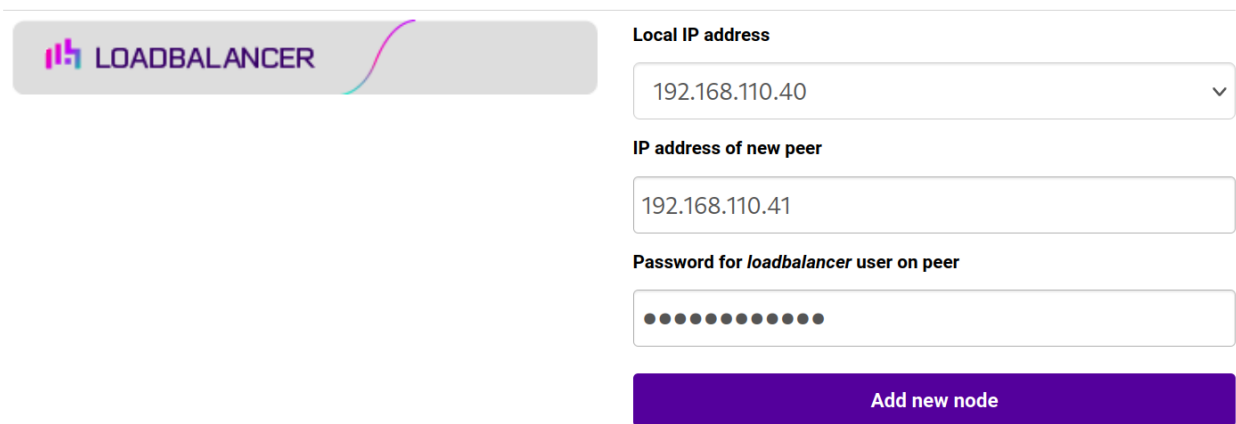
Configuring the HA Clustered Pair

📌 Note

If you have already run the firewall lockdown wizard on either appliance, you'll need to ensure that it is temporarily disabled on both appliances whilst performing the pairing process.

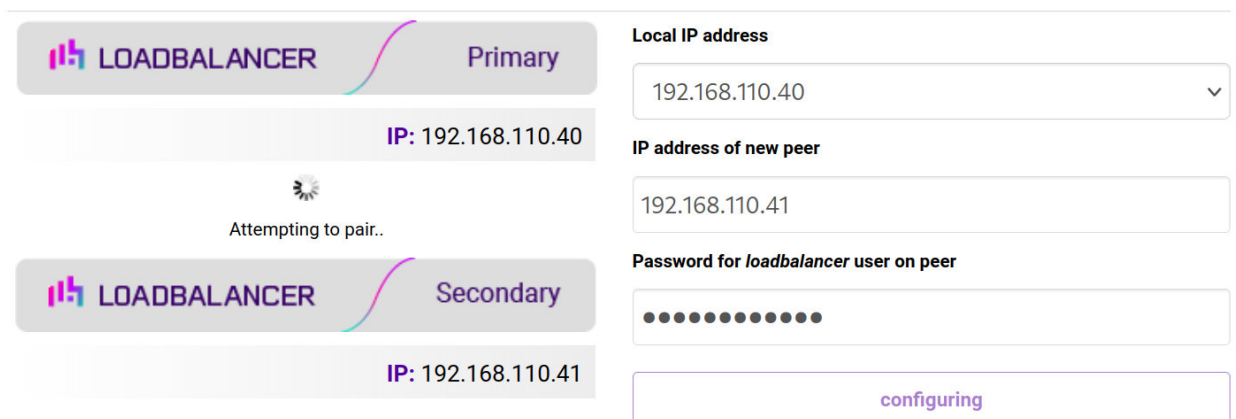
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: **Cluster Configuration > High-Availability Configuration**.

Create a Clustered Pair



3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

Create a Clustered Pair



6. Once complete, the following will be displayed on the Primary appliance:



High Availability Configuration - primary

The screenshot displays a configuration interface for a High Availability (HA) setup. It features two load balancer appliances arranged vertically. The top appliance is labeled 'LOADBALANCER Primary' and has the IP address 192.168.110.40. The bottom appliance is labeled 'LOADBALANCER Secondary' and has the IP address 192.168.110.41. To the right of these appliances is a prominent red button labeled 'Break Clustered Pair'.

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

Note

Clicking the **Restart Heartbeat** button on the Primary appliance will also automatically restart heartbeat on the Secondary appliance.

Note

For more details on configuring HA with 2 appliances, please refer to [Appliance Clustering for HA](#).

Note

For details on testing and verifying HA, please refer to [Clustered Pair Diagnostics](#).

14. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0	25 June 2024	Initial version		RJC
1.1	4 July 2024	Added Load Balancer Deployment Methods section	Standard requirement	RJC





Visit us: www.loadbalancer.org

Phone us: +44 (0)330 380 1064

Phone us: +1 833 274 2566

Email us: info@loadbalancer.org

Follow us: [@loadbalancer.org](https://twitter.com/loadbalancer.org)

About Loadbalancer.org

Loadbalancer.org's mission is to ensure that its clients' businesses are never interrupted. The load balancer experts ask the right questions to get to the heart of what matters, bringing a depth of understanding to each deployment. Experience enables Loadbalancer.org engineers to design less complex, unbreakable solutions - and to provide exceptional personalized support.

