LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing HPE Ezmeral Data Fabric Object Store

Version 1.0.1

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced HPE Ezmeral Data Fabric Object Store 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 Object Store. 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 Object Store

- All versions

4. HPE Ezmeral Data Fabric Object Store

HPE Ezmeral Data Fabric Object Store is HPE's native implementation of an object storage solution. It provides efficient and optimised S3-based access to data stored in HPE Ezmeral Data Fabric: HPE's industry-leading data platform.

5. Port Options for HPE Ezmeral Data Fabric Object Store

The **MOSS** (Multithreaded Object Store Server) nodes in an HPE Ezmeral Data Fabric Object Store deployment listen for incoming client connections on port 9000 by default. The MOSS nodes can be configured to provide *either* HTTP *or* HTTPS-based access.

5.1. Clients Use Port 9000

If incoming client connections will be expected to use the destination port 9000 then the load balancer should be

configured with a virtual service listening on port 9000.

5.2. Clients Use Port 80

If incoming client connections will be HTTP-based then it is possible to configure the load balancer with a virtual service listening on port 80 to allow clients to connect using the well-known port number for HTTP (as an alternative to clients having to connect using non-standard port 9000).

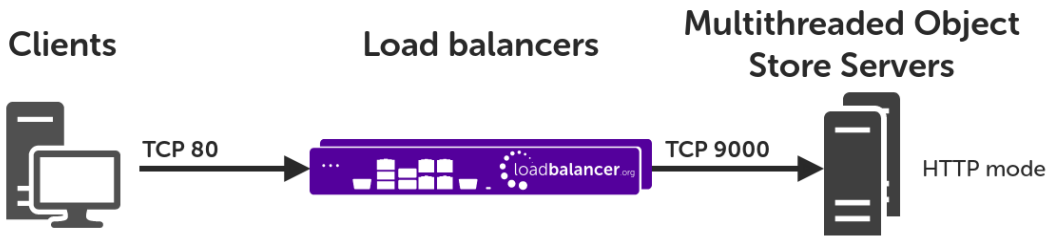

5.3. Clients Use Port 443

If incoming client connections will be HTTPS-based then it is possible to configure the load balancer with a virtual service listening on port 443 to allow clients to connect using the well-known port number for HTTPS (as an alternative to clients having to connect using non-standard port 9000).

5.4. Force to HTTPS

If the MOSS nodes are configured for HTTPS-based access then it is possible to configure the load balancer with a dummy virtual service listening on port 80 to redirect clients that attempt to connect using HTTP to reconnect using HTTPS. Instructions on how to configure this are provided in Section 9.4, "Configuring the Optional "Force to HTTPS" Service".

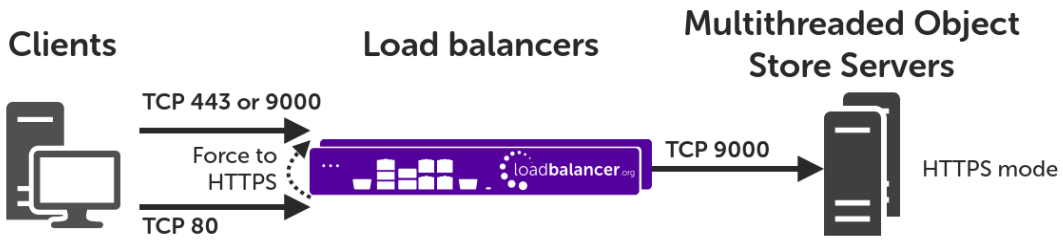

6. Load Balancing HPE Ezmeral Data Fabric Object Store

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

6.1. Persistence (aka Server Affinity)

HPE Ezmeral Data Fabric Object Store does not require session affinity at the load balancing layer. S3 requests are spread across all MOSS nodes as evenly as possible to try and distribute the load across the MOSS nodes as evenly as possible.

6.2. Virtual Service (VIP) Requirements

To provide load balancing and HA for HPE Ezmeral Data Fabric Object Store, a single VIP is required:

- MOSS Service

Optionally, an additional VIP may be required as follows:

- Force to HTTPS Redirect

6.3. Port Requirements

The following table shows the ports that are load balanced:

Ports	Protocols	Use
80, 443, 9000	TCP/HTTP(S)/S3	S3 Object Store Access

7. Deployment Concept

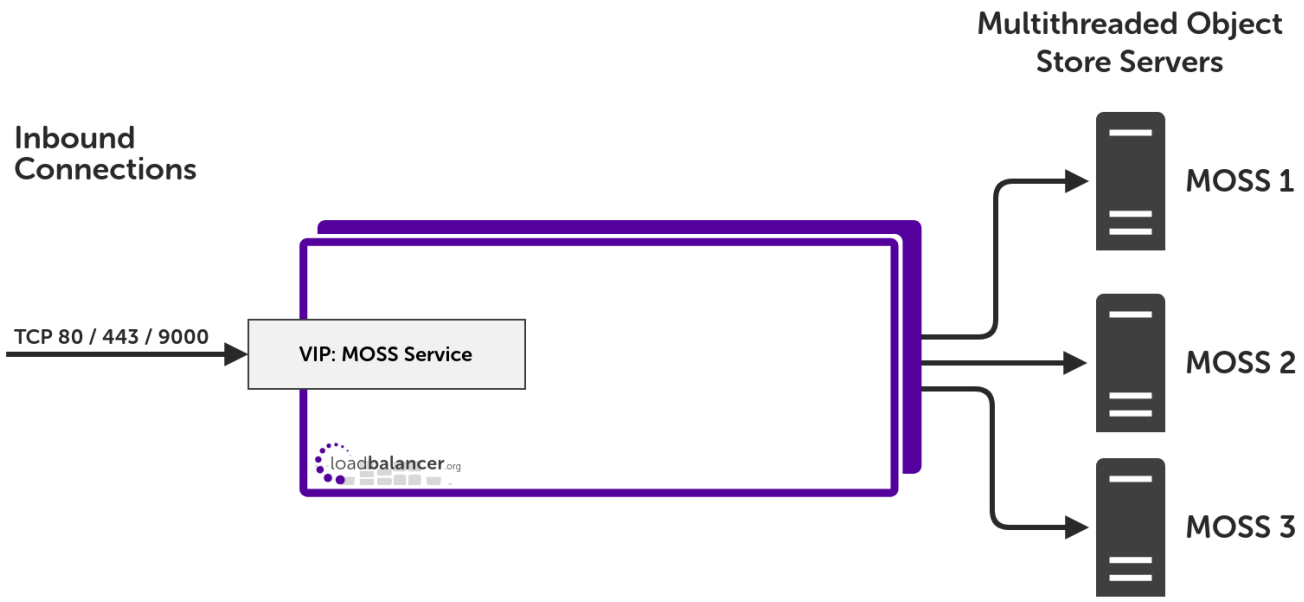

VIP = **V**irtual **IP** Address

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.

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:

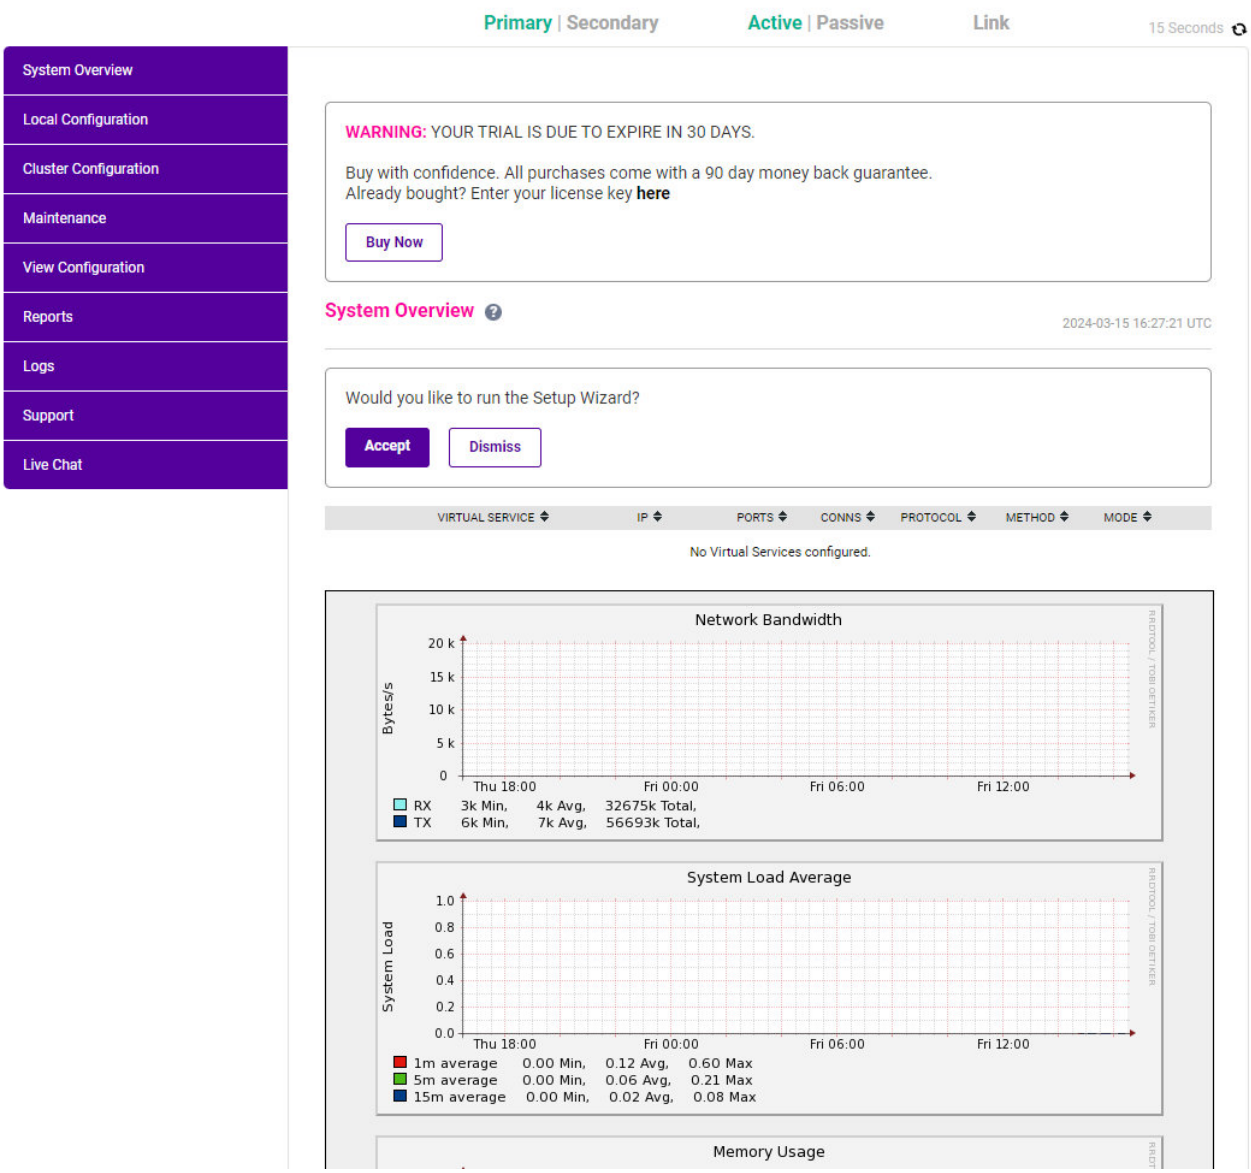

3. You'll be asked if you want to run the Setup Wizard which can be used to configure layer 7 services. Click **Dismiss** if you're following a guide or want to configure the appliance manually or click **Accept** to start the wizard.

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:

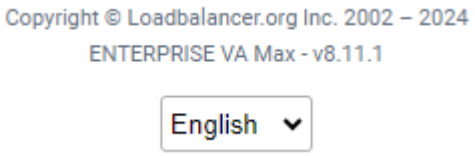

Checking for Updates using Online Update

> **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:

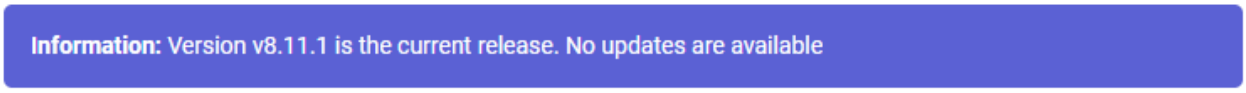

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

> **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:

7. 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:

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

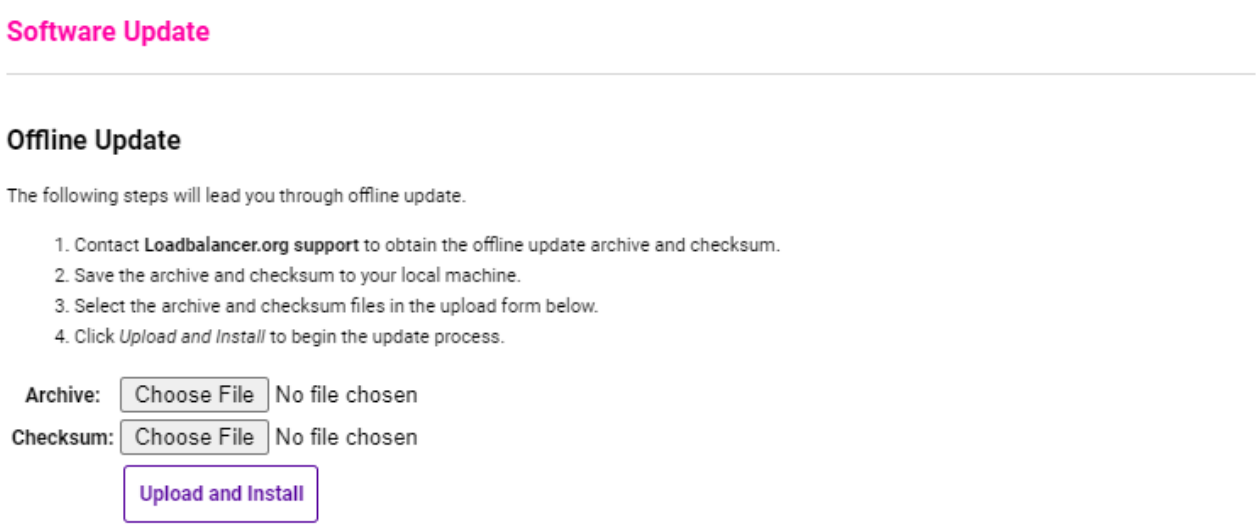

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

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
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

Protocol	Port	Purpose
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 Object Store

9.1. Enabling Multithreaded Load Balancing

> **Note** Multithreading is enabled by default for ***new*** load balancers starting from version 8.5.1 and does not require changing.
>
> ***If upgrading an older appliance*** then ensure that the multithreading configuration is set correctly, as described below.

The Loadbalancer.org appliance should be configured to actively use multiple CPU cores for the load balancing process. This is required to achieve the high level of performance and throughput required when load balancing an HPE Ezmeral Data Fabric Object Store deployment.

> **Note** A virtual host should be allocated a minimum of 4 vCPUs.

To enable multithreaded mode from the WebUI:

1. Navigate to *Cluster Configuration > Layer 7 - Advanced Configuration*.
2. Check the **Enable Multithreading** checkbox.
3. Check the **Default Number of Threads** checkbox.
4. Click **Update** to apply the changes.

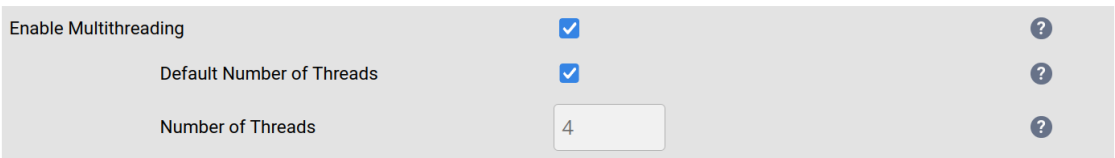

9.2. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **MOSS Service**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.85.150**.
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.

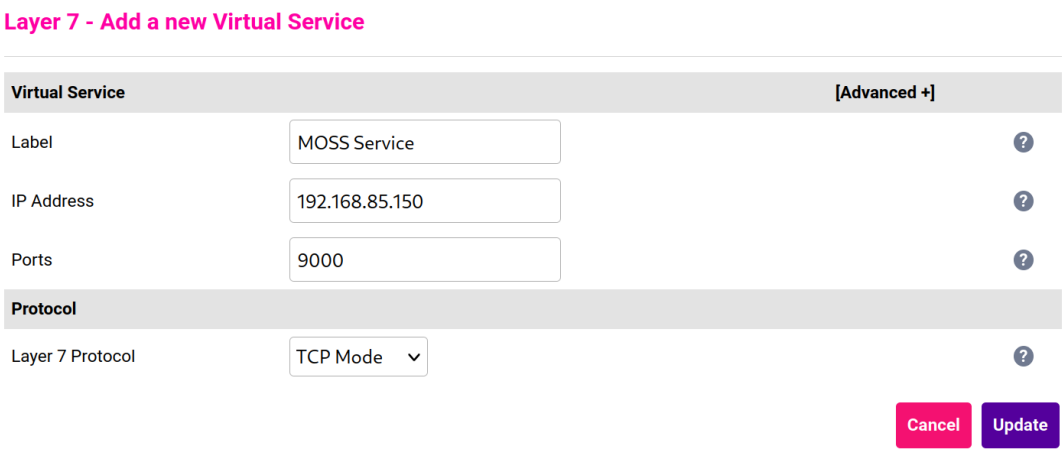

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Click **Update**.

9.3. 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.
2. Define the *Label* for the real server as required, e.g. **MOSS 1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.200**.
4. Set the *Real Server Port* field to **9000**.
5. Click **Update**.
6. Repeat these steps to add the remaining MOSS nodes.

Layer 7 Add a new Real Server - MOSS_Service

Label	MOSS 1	?
Real Server IP Address	192.168.85.200	?
Real Server Port	9000	?
Re-Encrypt to Backend	☐	?
Weight	100	?

Cancel Update

9.4. Configuring the Optional "Force to HTTPS" Service

> **Note** This step is **optional** and should be skipped if having a "force to HTTPS" service is not wanted or required.

If the MOSS nodes are configured for HTTPS-based access then it is possible to configure the load balancer with a dummy virtual service listening on port 80 to redirect clients that attempt to connect using HTTP to reconnect using HTTPS. The instructions on how to configure this are as follows:

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **HTTPS Redirect**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.85.150**.
4. Set the *Ports* field to **80**.
5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

Layer 7 - Add a new Virtual Service

Virtual Service		**[Advanced +]**
Label	HTTPS Redirect	?
IP Address	192.168.85.150	?
Ports	80	?
Protocol		
Layer 7 Protocol	HTTP Mode ⌄	?

Cancel Update

7. Click **Modify** next to the newly created VIP.
8. In the *Other* section click **Advanced** to expand the menu.
9. Set *Force to HTTPS* to **Yes** by clicking on the radio button.

10. Set the *HTTPS Redirect Port* to the port that the MOSS service VIP created in a previous step is listening on. This should be either 443 or 9000.
11. Click **Update**.

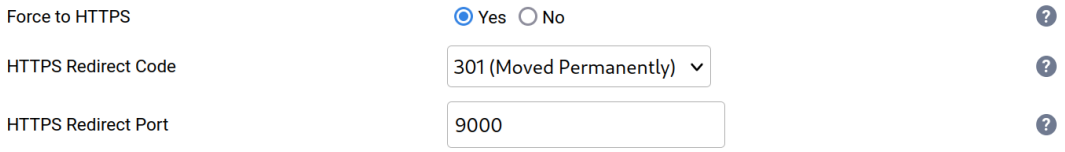

9.5. Finalizing the Configuration

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. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the Multithreaded Object Store Server nodes) and shows the state/health of each server as well as the state of the cluster as a whole. The example below shows a standard deployment where all MOSS nodes are healthy and available to accept connections:

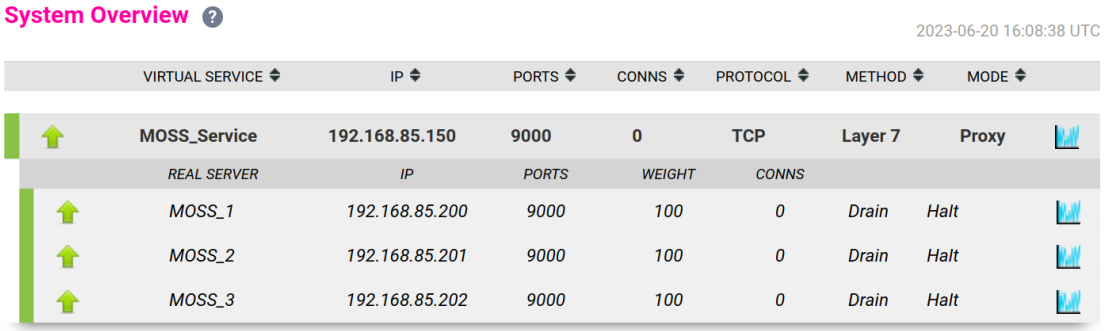

11. Technical Support

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*.

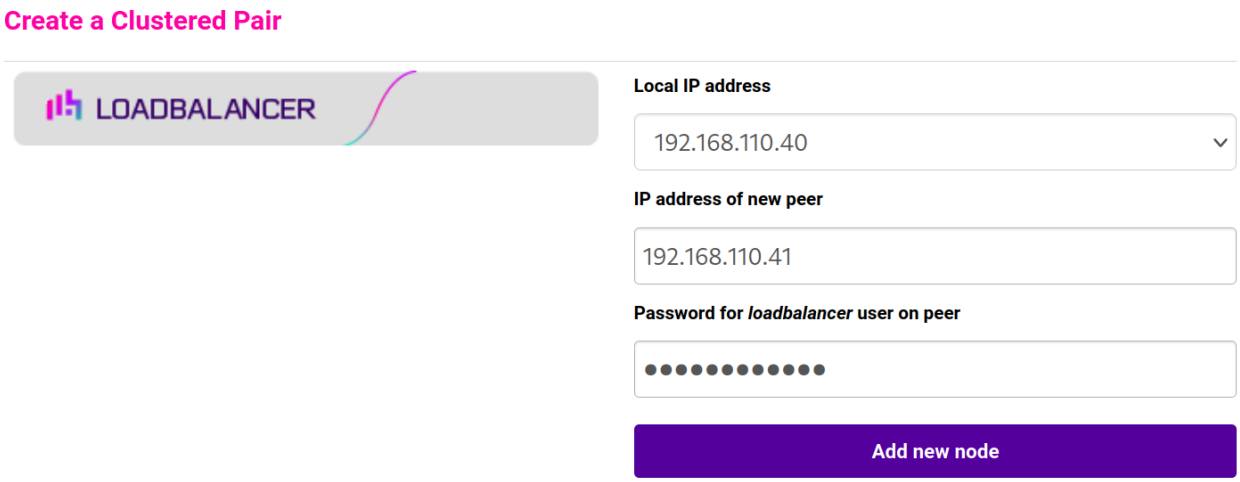

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

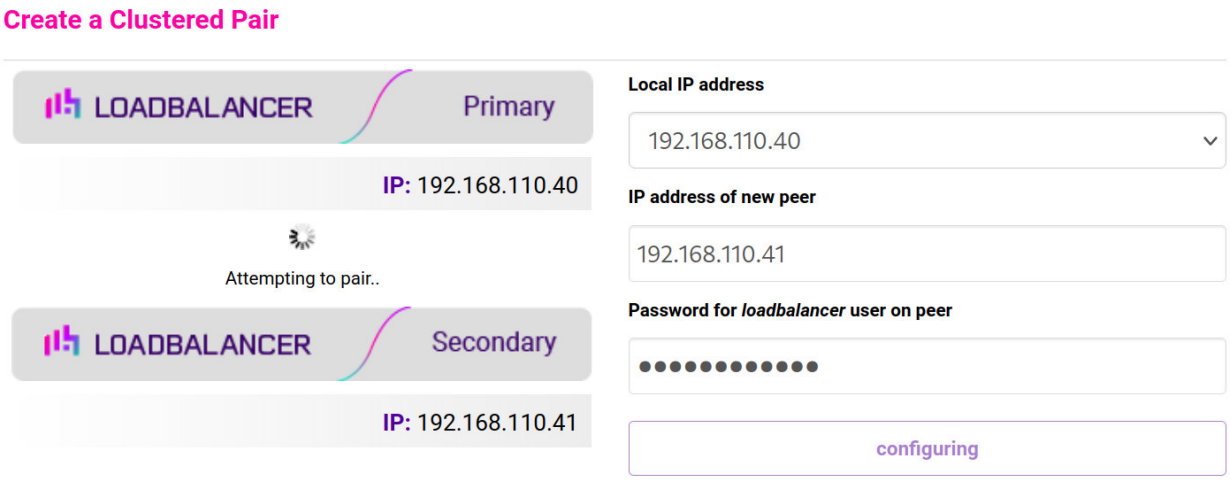

6. Once complete, the following will be displayed on the Primary appliance:

High Availability Configuration - primary

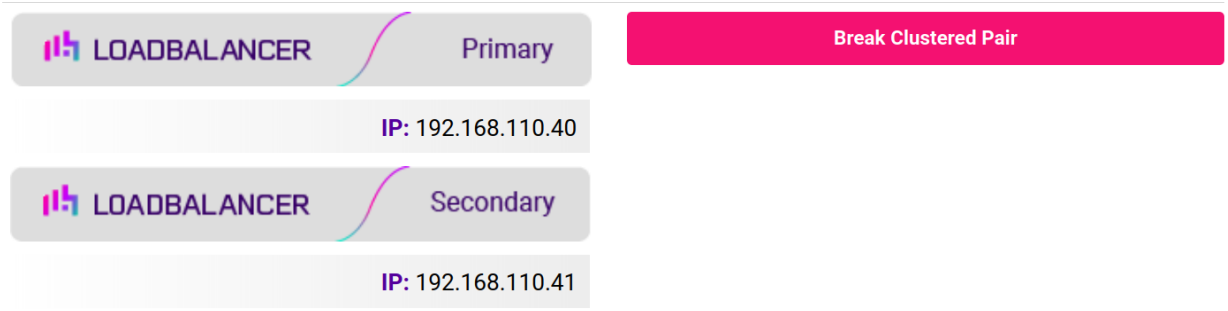

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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
1.0.0	21 June 2023	Initial version		AH
1.0.1	29 June 2023	Updated multithreading advice	New default option in the web user interface	AH

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

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.