LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing Evertz Mediator-X

Version 1.3.1

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced Evertz Mediator-X environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Evertz Mediator-X 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 for load balancing Evertz Mediator-X. For full specifications of available models please refer to https://www.loadbalancer.org/products. Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.3.8 and later

ℹ Note	The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If using an older software version, note that the screenshots presented here may not match the WebUI exactly.

3.2. Evertz Mediator

- Evertz Mediator-X

4. Evertz Mediator-X

Evertz Mediator-X unifies content acquisition, content processing, media management, production, playout, and delivery into a single, integrated environment. The unification of these services on a single platform delivers optimized media workflows and increased operational efficiency.

Built on over fifteen years of Mediator product development and deployment expertise, Mediator-X has a modern, scalable, infrastructure-agnostic architecture which can be deployed in public cloud, private cloud, or hybrid environments, enabling users to be flexible with their deployment strategies and to grow the platform wherever the business case dictates.

Evertz recommends Loadbalancer.org appliances to provide high availability and load balancing of the Mediator-X platform.

5. Load Balancing Evertz Mediator-X

ℹ Note	It's highly recommended that you have a working Evertz Mediator-X environment first before

implementing the load balancer.

5.1. Sizing, Capacity, and Performance for a Virtual Load Balancer Deployment

The Loadbalancer.org appliances can be deployed as **virtual appliances**.

For **small deployments** handling up to 300 concurrent connections/users, your virtual host should be allocated a minimum of 4 vCPUs, 4 GB of RAM, and 8 GB of disk storage.

For **large deployments** handling over 300 concurrent connections/users, your virtual host should be allocated a minimum of 8 vCPUs, 8 GB of RAM, and 8 GB of disk storage.

For **significantly larger deployments**, your Evertz representative will give you custom sizing and resource guidelines based on the expected load on your load balancers and your predicted usage profile.

5.2. Persistence (aka Server Affinity)

For the **layer 4 DR mode scenario**, each virtual service uses source IP address-based persistence.

For the **layer 7 load balancing scenario** (the configuration that adds TLS-based encryption), the persistence mode *X-Forwarded-For and Source IP* is used. This uses X-Forwarded-For HTTP headers as the primary persistence method, with source IP addresses used as a backup persistence method.

5.3. Virtual Service (VIP) Requirements

To provide load balancing and HA for Evertz Mediator-X, the following VIP is required:

- Mediator-X Global Access

The "Global" virtual service handles Mediator-X user interface traffic and API endpoint traffic. "Global" access to both services is provided using a single virtual service on the load balancer.

Additionally, a TLS/SSL termination service is required for the scenario that adds TLS-based encryption.

5.4. Port Requirements

The following table shows the ports that are load balanced:

Port	Protocols	Uses
80	TCP/HTTP	Mediator-X user interface access, Mediator-X API endpoint access
443	TCP/HTTPS	Mediator-X user interface access over TLS (optional)

5.5. TLS/SSL Termination

It is possible to configure a TLS/SSL termination service in front of the plaintext, port 80, HTTP-based *Mediator-Global* service. This enables inbound client connections to be secured using TLS. Connections from the load balancer to the Mediator-X servers remain as plaintext HTTP connections (not encrypted) on port 80. In this way, inbound client connections can be secured using encryption without needing to make any changes to the back end Mediator-X servers.

6. Deployment Concept

Evertz Mediator-X can be load balanced in two different ways:

- **Simple deployment**: Uses a single virtual service to load balance all of the port 80 traffic used by Mediator-X (the user interface traffic as well as the API endpoint traffic)
- **Deployment using TLS-based encryption**: An alternative deployment type that should only be used when there is the requirement to secure client connections using TLS-based encryption. Using this deployment type, clients can connect to the Mediator-X User Interface using HTTPS on port 443

6.1. Scenario 1 – Simple Deployment

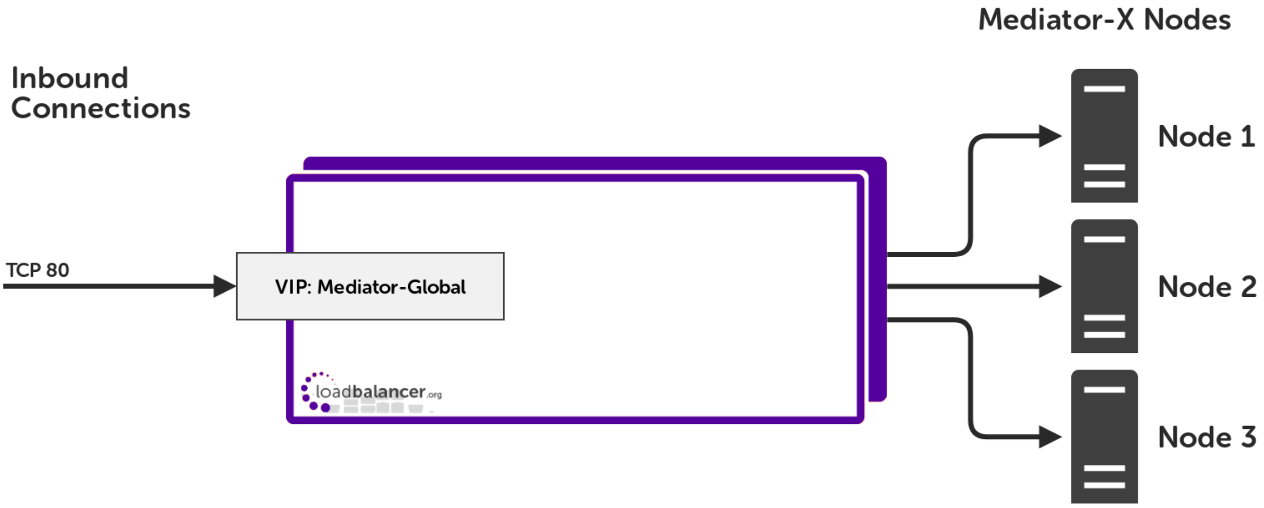

VIPs = **V**irtual **IP** Addresses

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 Configuring HA - Adding a Secondary Appliance for more details on configuring a clustered pair.

In this deployment, a single virtual service is used. The virtual service uses layer 4 DR mode, offering the greatest possible network speed and scalability.

Layer 4 DR mode is the load balancing method that has traditionally been used with Evertz Mediator deployments.

6.2. Scenario 2 – Deployment Using TLS-Based Encryption

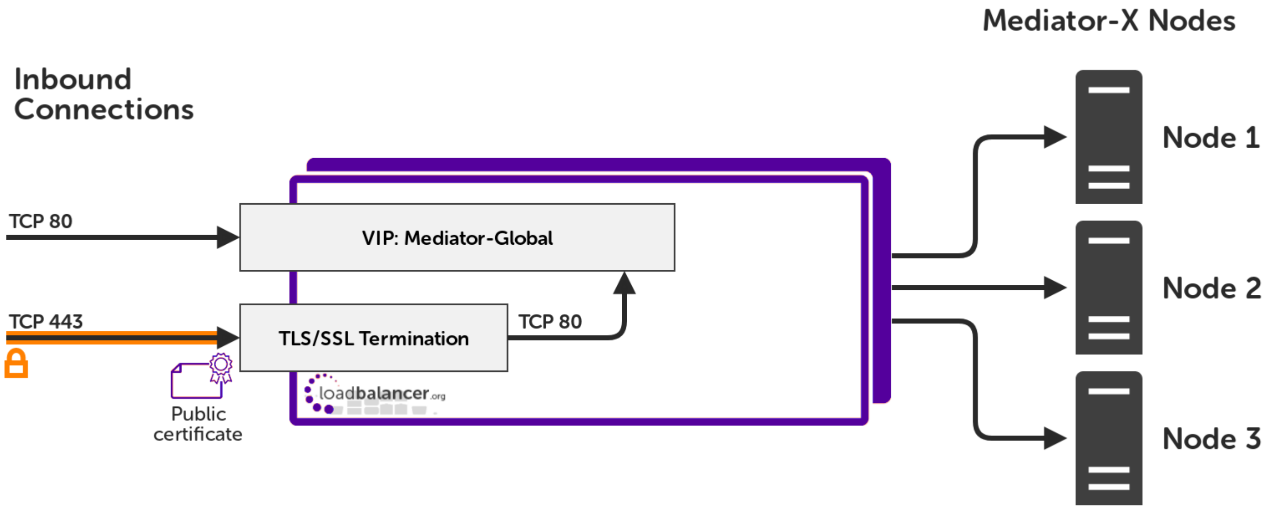

VIPs = **V**irtual **IP** Addresses

ℹ 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 Configuring HA - Adding a Secondary Appliance for more details on configuring a clustered pair.

In this deployment, one virtual service is used in addition to a TLS/SSL termination. The virtual service uses layer 7 SNAT mode.

This alternative deployment type allows for Mediator-X traffic to be secured using TLS, with clients sending encrypted traffic on port 443 instead of plaintext traffic on port 80.

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 Mediator-X, using layer 4 DR mode is recommended due to its raw throughput and huge scalability. It is also possible to use layer 7 SNAT mode, which allows adding TLS-based encryption for client traffic, however the performance of this set up is not as great as layer 4 DR mode. These load balancing modes are described below and are used for the configurations presented in this guide. For configuring using DR mode please refer to Appliance Configuration for Evertz Mediator-X – Using Layer 4 DR Mode (Scenario 1: Simple Deployment) and for configuring using layer 7 SNAT mode, which allows adding TLS-based encryption, refer to Appliance Configuration for Evertz Mediator-X – Using Layer 7 SNAT Mode (Scenario 2: Deployment Using TLS-Based Encryption).

7.1. Layer 4 DR Mode

One-arm direct routing (DR) mode is a very high performance solution that requires little change to your existing infrastructure.

ℹ Note — Kemp, Brocade, Barracuda & A10 Networks call this *Direct Server Return* and F5 call it *nPath*.

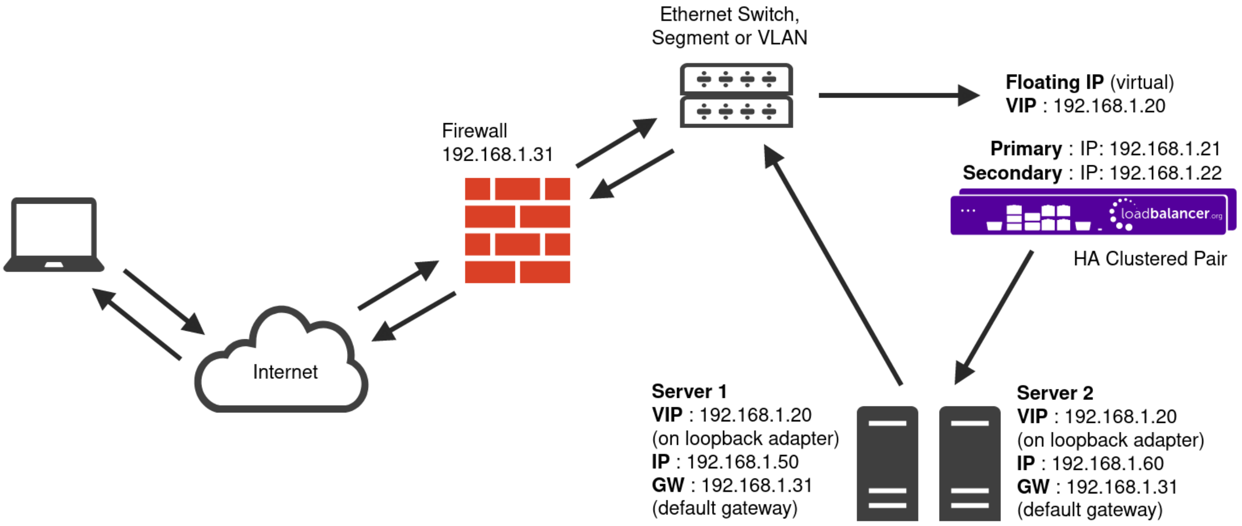

- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that you need to ensure that the Real Server (and the load balanced application) respond to both the Real Server's own IP address and the VIP.
- The Real Servers should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Servers in this way is referred to as ***Solving the ARP problem***. For more information please refer to DR Mode Considerations.
- On average, DR mode is 8 times quicker than NAT for HTTP, 50 times quicker for Terminal Services and much, much faster for streaming media or FTP.
- The load balancer must have an Interface in the same subnet as the Real Servers to ensure layer 2 connectivity required for DR mode to work.
- The VIP can be brought up on the same subnet as the Real Servers, or on a different subnet provided that the load balancer has an interface in that subnet.
- Port translation is not possible with DR mode, e.g. VIP:80 → RIP:8080 is not supported.
- DR mode is transparent, i.e. the Real Server will see the source IP address of the client.

7.2. 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.

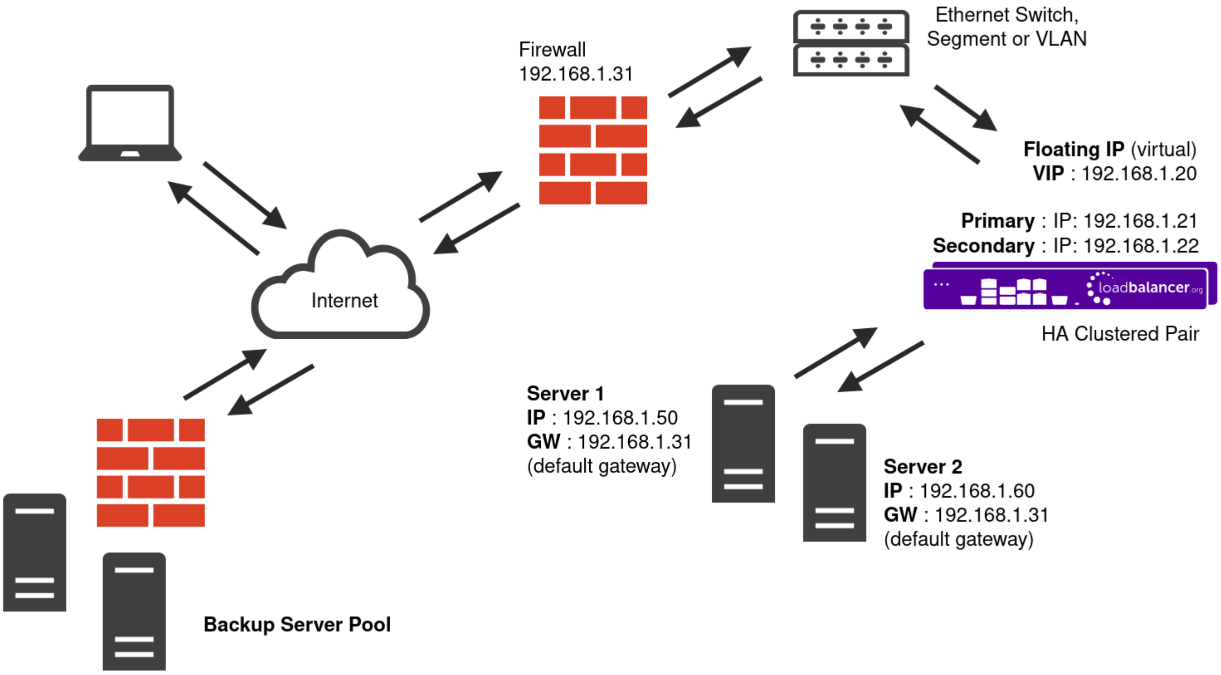
Ethernet Switch, Segment or VLAN
Firewall 192.168.1.31
Floating IP (virtual) VIP : 192.168.1.20
Primary : IP: 192.168.1.21
Secondary : IP: 192.168.1.22
loadbalancer.org
HA Clustered Pair
Internet
Server 1 IP : 192.168.1.50 GW : 192.168.1.31 (default gateway)
Server 2 IP : 192.168.1.60 GW : 192.168.1.31 (default gateway)
Backup Server Pool

- 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.

7.3. Our Recommendation

Where possible, we recommend that Layer 4 Direct Routing (DR) mode is used. This mode offers the best possible performance since replies go directly from the Real Servers to the client, not via the load balancer. It's also relatively simple to implement. Ultimately, the final choice does depend on your specific requirements and infrastructure.

If DR mode cannot be used, for example if the real servers are located in remote routed networks, then SNAT mode is recommended. SNAT is also recommended if TLS-based encryption is required for the HTTP aspect of the Mediator-X inbound client traffic.

If the load balancer is deployed in AWS or Azure, layer 7 SNAT mode must be used as layer 4 direct routing is not currently possible on these platforms.

8. Configuring Evertz Mediator-X for Load Balancing

Some changes must be made to the Mediator-X real servers in order for them to be correctly load balanced. **These changes need to be configured by an Evertz Deployment Team**. Contact your Evertz representative for further information.

1. On the Mediator-X deployment, connect to Node0 and log in as the root user. This can be done by executing the command **sudo su** and then entering the system specific shell access password.
2. Navigate to the directory **/srv/salt/pillar**. This can be done by executing the command **cd /srv/salt/pillar**.
3. Edit the file **system.sls**, for example using a text editor such as nano or vim: **nano system.sls**.
4. Find the **virtual_ips** parameter and add the virtual IP address that will be used for the load balanced deployment. This is the user facing IP address that all clients will connect to when accessing the load balanced Mediator-X services. The result should look like the following:

```
# Set the loadbalancer virtual ip or leave blank
virtual_ips : 10.0.1.50
```

5. Save and exit the system.sls file.
6. Run the salt command and call the ***state.highstate*** function, which will automatically apply the changed configuration across all Mediator-X nodes. The full command to execute is:

```
sudo salt "*" state.highstate
```

7. Run the salt command and call a function to restart the nginx service across all Mediator-X nodes. The full command to execute is:

```
sudo salt -G 'is_mediatorx:True' service.restart nginx
```

8.1. Additional Changes when Adding TLS-Based Encryption

When deploying using Scenario 2 – Deployment Using TLS-Based Encryption, there may be situations where it is desired (or required) to access certain resources over HTTPS that would otherwise only be served over plaintext HTTP. DASH manifests and the Mediator-X portal are two examples of such resources.

> ℹ Note
>
> For further Mediator-X-specific information beyond what is presented here, or for advice on whether other parts of the Mediator-X application to be configured to point to and use HTTPS, contact your Evertz representative.

DASH Manifests

To allow HTTPS-based access to take place:

1. In the Mediator-X settings navigate to *Browse media > Browse info*.

2. Ensure that the scheme of the *Browse http url* is **https**.

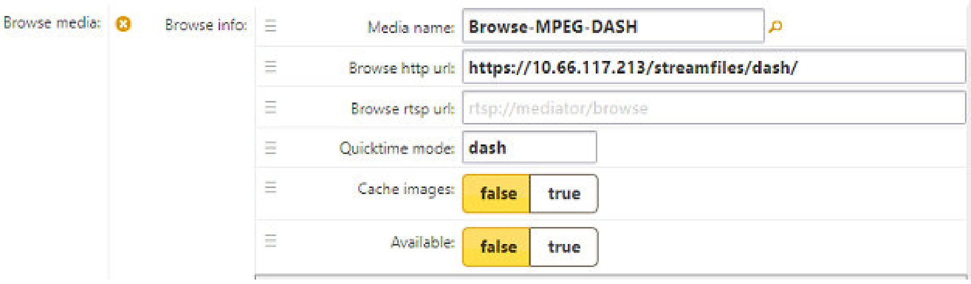

Browse media:

Browse info:

Media name: Browse-MPEG-DASH

Browse http url: https://10.66.117.213/streamfiles/dash/

Browse rtsp url: rtsp://mediator/browse

Quicktime mode: dash

Cache images: [✓ false] true

Available: [✓ false] true

Mediator-X Portal

To allow HTTPS-based access to take place:

1. In the Mediator-X settings navigate to *Edit system settings*.

2. Ensure that the scheme of the *Base ui url* is **https**.

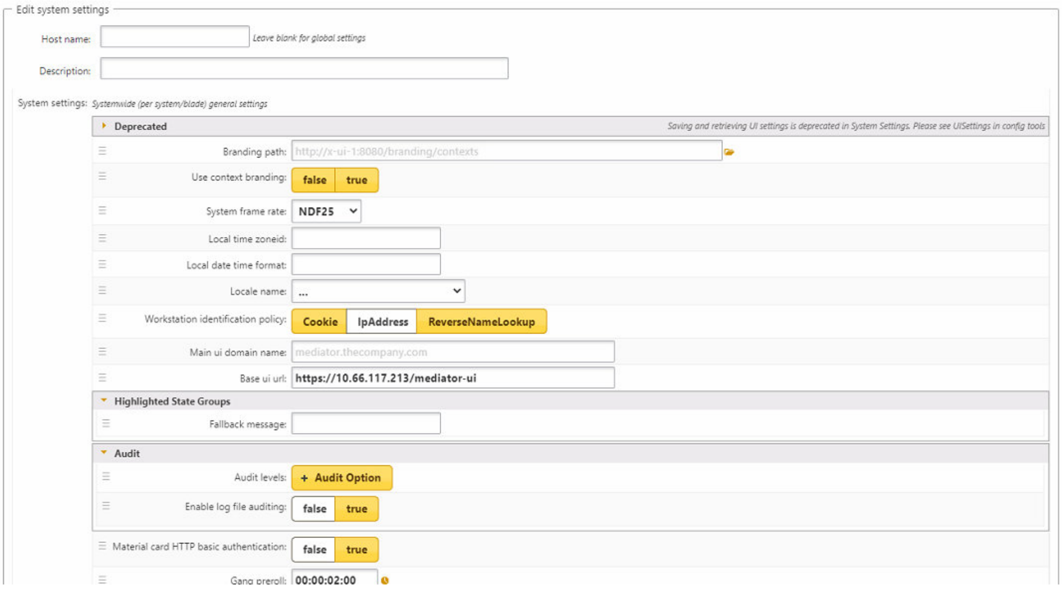

Edit system settings

Host name: Leave blank for global settings

Description:

System settings: Systemwide (per system/blade) general settings

Deprecated — Saving and retrieving UI settings is deprecated in System Settings. Please see UISettings in config tools

Branding path: http://x-ui-1:8080/branding/contexts

Use context branding: [✓ false] [✓ true]

System frame rate: NDF25

Local time zoneid:

Local date time format:

Locale name: ...

Workstation identification policy: [✓ Cookie] IpAddress [✓ ReverseNameLookup]

Main ui domain name: mediator.thecompany.com

Base ui url: https://10.66.117.213/mediator-ui

Highlighted State Groups

Fallback message:

Audit

Audit levels: + Audit Option

Enable log file auditing: false [✓ true]

Material card HTTP basic authentication: false [✓ true]

Gang preroll: 00:00:02:00

9. Loadbalancer.org Appliance – the Basics

9.1. 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.

ℹ 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.

ℹ 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.

9.2. Initial Network Configuration

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

(!) Important

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

9.3. 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.

ℹ Note

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

ℹ Note

A number of compatibility issues have been found with various versions of Microsoft Internet Explorer and Edge. The WebUI has been tested and verified using both Chrome & Firefox.

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

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

   ℹ Note

   You'll receive a warning about the WebUI's 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.

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

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

   ℹ Note

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

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

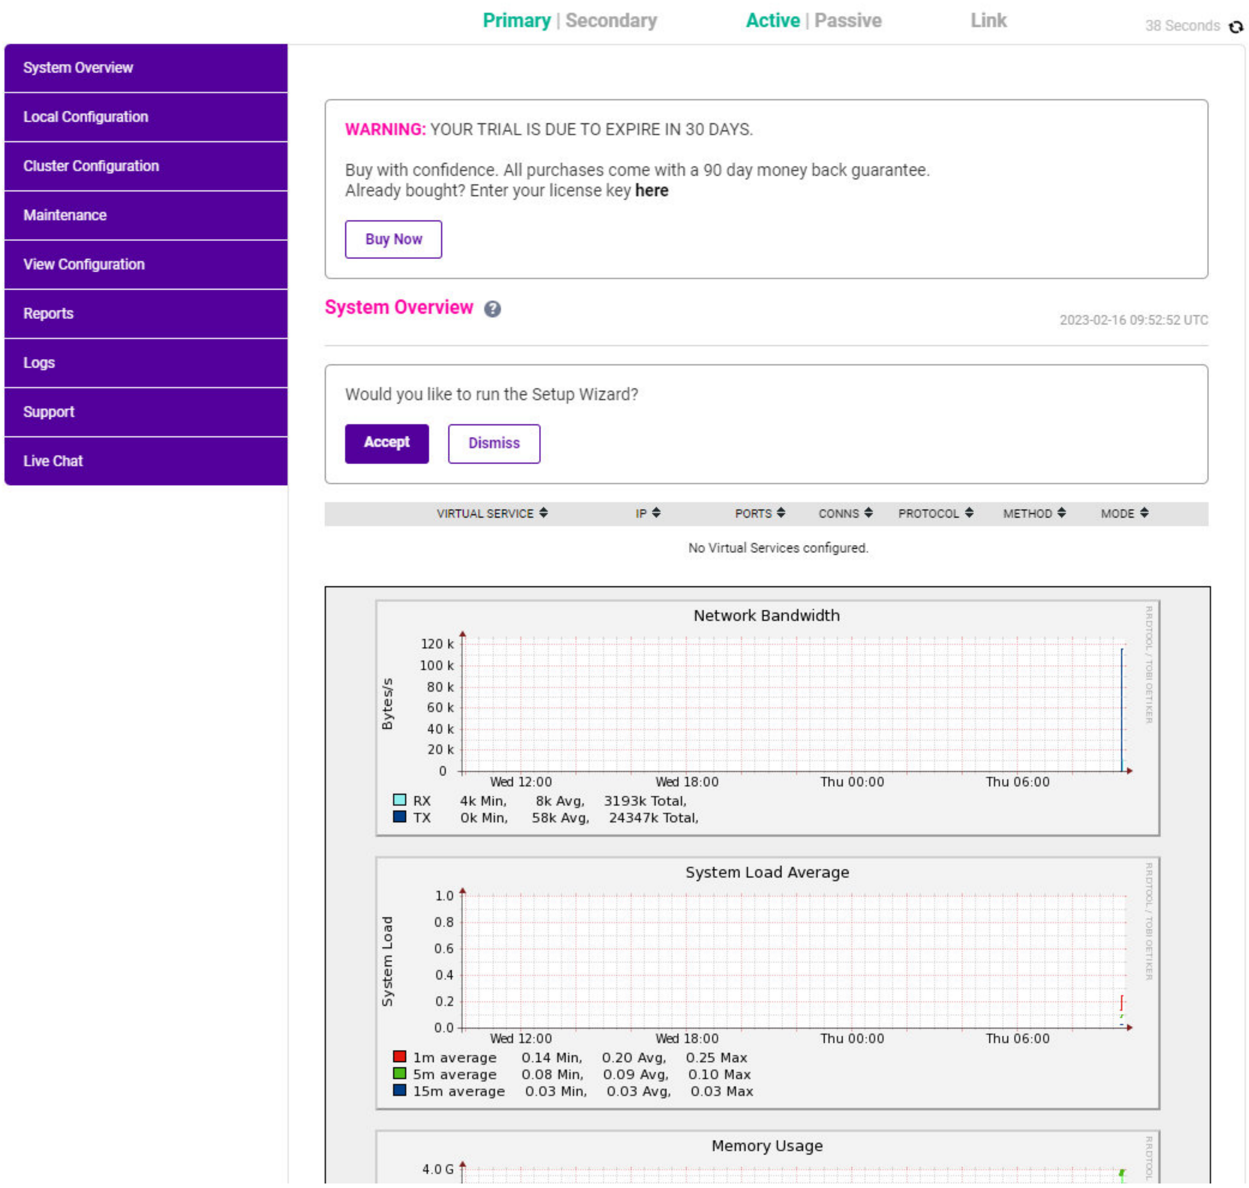

WARNING: YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS.

Buy with confidence. All purchases come with a 90 day money back guarantee.
Already bought? Enter your license key **here**

System Overview 2023-02-16 09:52:52 UTC

Would you like to run the Setup Wizard?

No Virtual Services configured.

```
□ RX   4k Min,    8k Avg,   3193k Total,
■ TX   0k Min,   58k Avg,  24347k Total,
```

```
🟥 1m average    0.14 Min,   0.20 Avg,   0.25 Max
🟩 5m average    0.08 Min,   0.09 Avg,   0.10 Max
■ 15m average    0.03 Min,   0.03 Avg,   0.03 Max
```

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

> Note   The Setup Wizard can only be used to configure Layer 7 services.

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

9.4. 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:

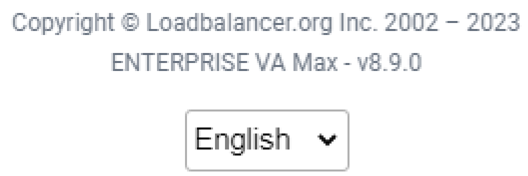

Copyright © Loadbalancer.org Inc. 2002 – 2023
ENTERPRISE VA Max - v8.9.0

Checking for Updates using Online Update

> ℹ 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:

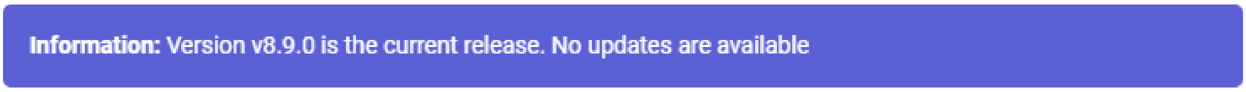

Information: Version v8.9.0 is the current release. No updates are available

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.

> ℹ 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:

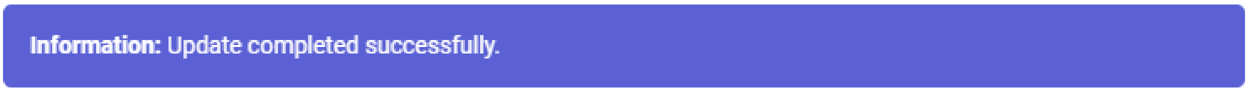

Information: Update completed successfully.

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.

> ℹ 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:

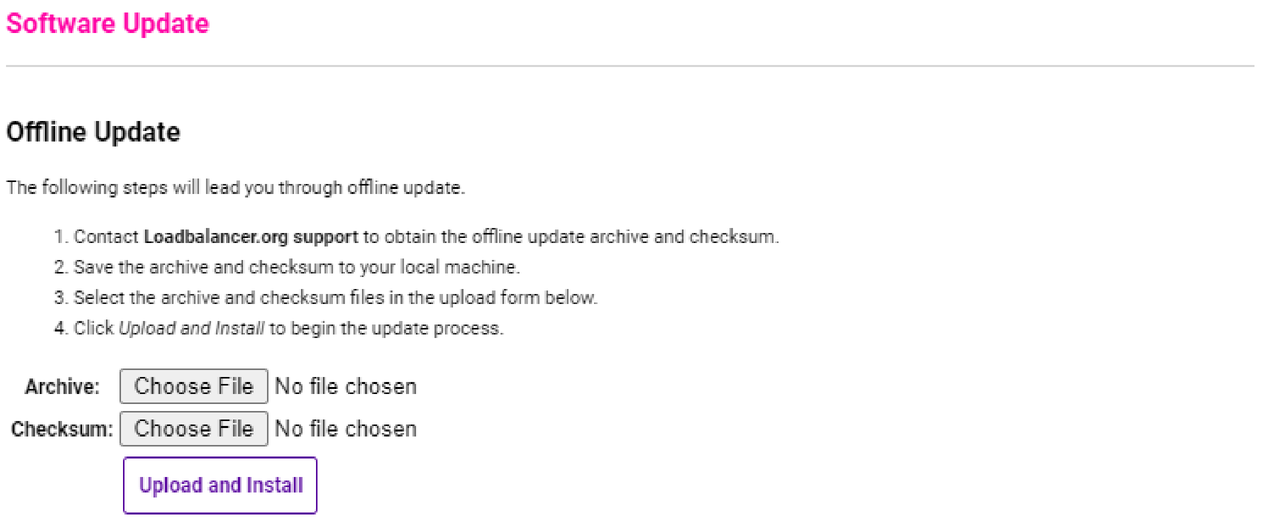

Software Update

Offline Update

The following steps will lead you through offline update.

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

Archive: Choose File No file chosen

Checksum: Choose File No file chosen

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.

9.5. 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
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	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback page
TCP	9443	WebUI - HTTPS

9.6. 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 Configuring HA - Adding a Secondary Appliance.

10. Appliance Configuration for Evertz Mediator-X – Using Layer 4 DR Mode (Scenario 1: Simple Deployment)

10.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **Mediator-Global**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.85.140**.
4. Set the *Ports* field to **80**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the virtual service.

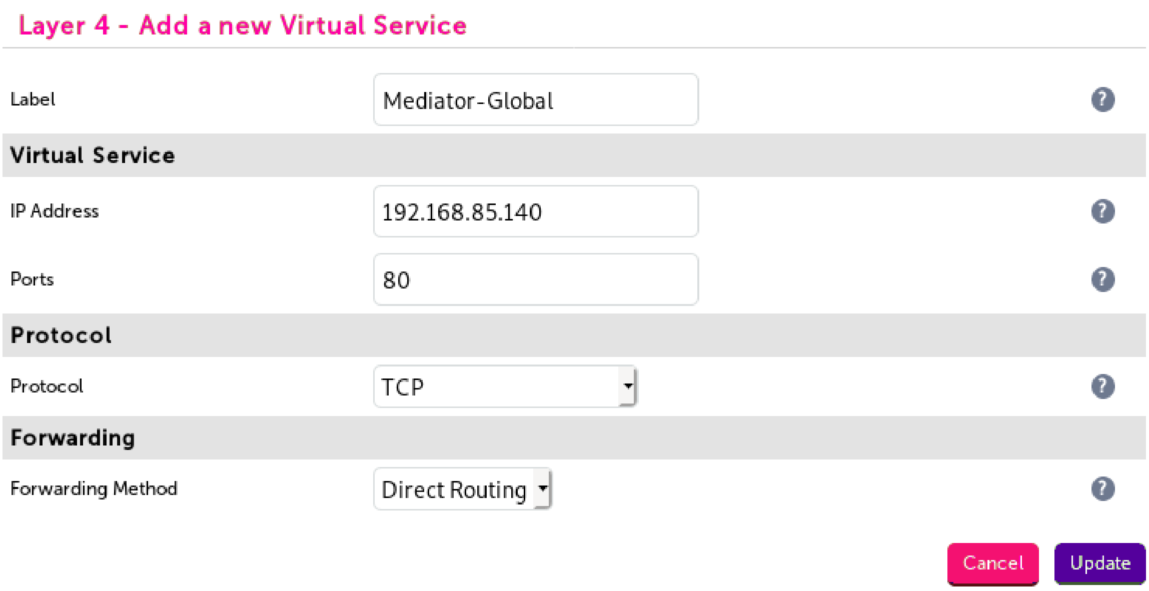

Layer 4 - Add a new Virtual Service

Label	Mediator-Global
Virtual Service	
IP Address	192.168.85.140
Ports	80
Protocol	
Protocol	TCP
Forwarding	
Forwarding Method	Direct Routing

8. Click **Modify** next to the newly created VIP.
9. Set the *Balance Mode* to **Weighted Round Robin**.
10. Ensure that the *Persistence Enable* checkbox is checked and that the *Timeout* is set to **300** (this should already be configured by default).
11. Set the *Health Checks Check Type* to **Negotiate**.
12. Set the *Check Port* to **80**.

13. Set the *Protocol* to **HTTP**.

14. Set the *Request to send* to **/mediator/main/loadBalancing/isExternallyAccessibleAPI**.

15. Ensure that the *Response expected* field is blank.

16. Click **Update**.

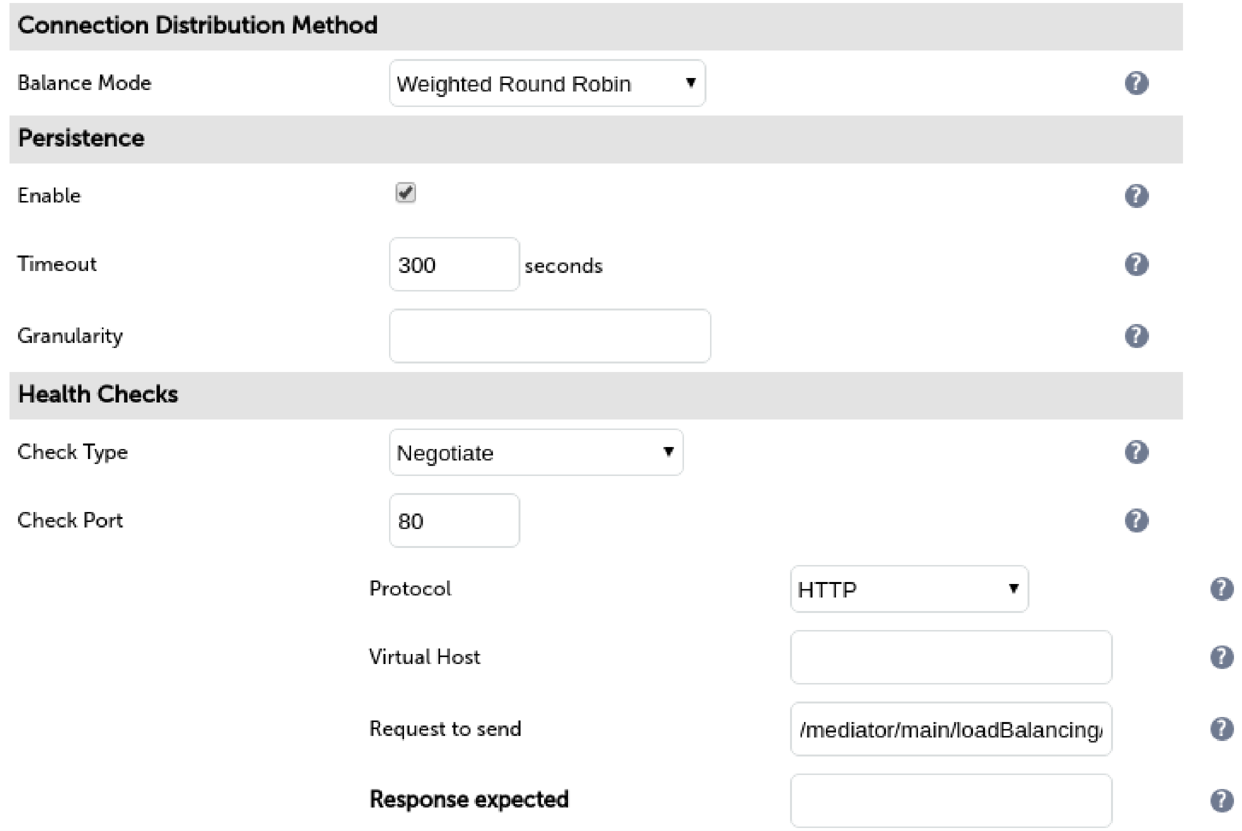

Connection Distribution Method

Balance Mode: Weighted Round Robin

Persistence

Enable: ☑

Timeout: 300 seconds

Granularity:

Health Checks

Check Type: Negotiate

Check Port: 80

Protocol: HTTP

Virtual Host:

Request to send: /mediator/main/loadBalancing/

Response expected:

10.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – 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. **node-04**.

3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.24**.

4. Click **Update**.

5. Repeat these steps to add additional Mediator-X servers as required.

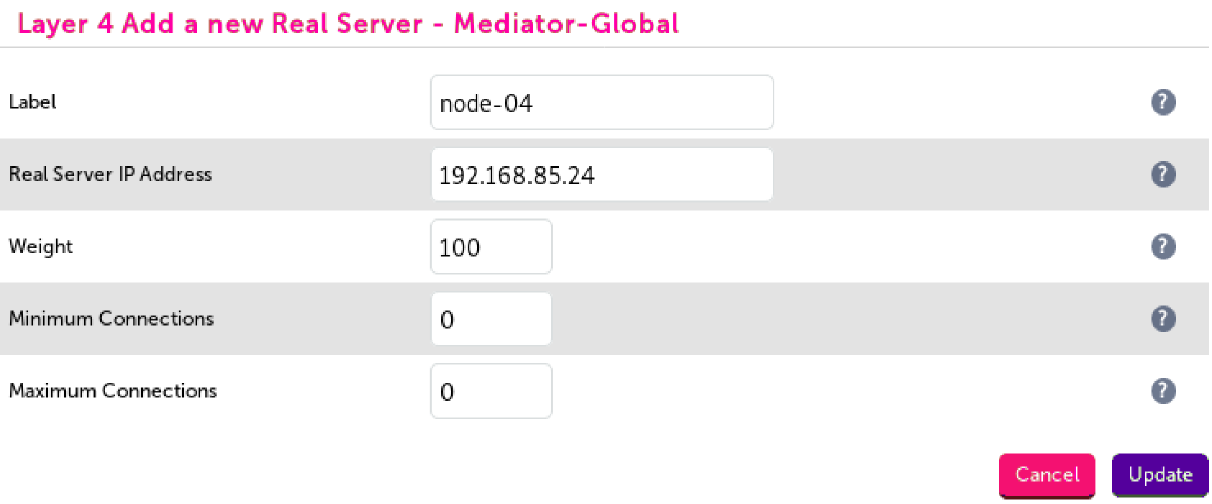

Layer 4 Add a new Real Server - Mediator-Global

Label	node-04
Real Server IP Address	192.168.85.24
Weight	100
Minimum Connections	0
Maximum Connections	0

11. Appliance Configuration for Evertz Mediator-X – Using Layer 7 SNAT Mode (Scenario 2: Deployment Using TLS-Based Encryption)

11.1. 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.

For the layer 7 load balancing scenario, 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 a Mediator-X deployment at layer 7.

> 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.

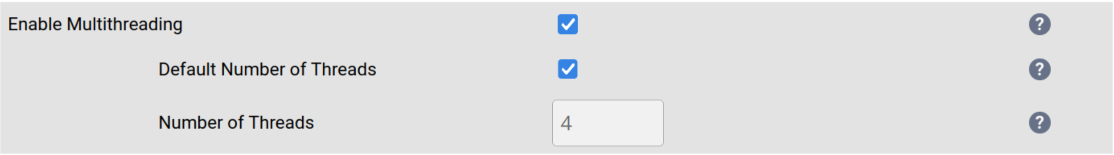

Enable Multithreading	☑
Default Number of Threads	☑
Number of Threads	4

11.2. 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. **Mediator-Global**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.85.140**.
4. Set the *Virtual Service Ports* field to **80**.
5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

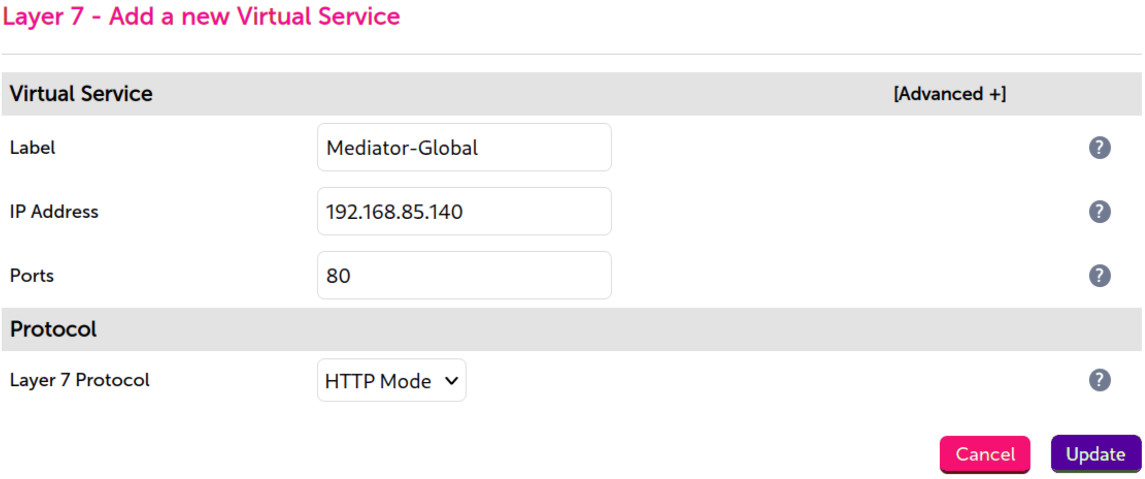

Layer 7 - Add a new Virtual Service

Virtual Service	[Advanced +]
Label	Mediator-Global
IP Address	192.168.85.140
Ports	80
Protocol	
Layer 7 Protocol	HTTP Mode

7. Click **Modify** next to the newly created VIP.
8. Set the *Balance Mode* to **Weighted Round Robin**.
9. Set *Persistence Mode* to **X-Forwarded-For and Source IP**.
10. Click the *Persistence* **Advanced** button to expand the menu.
11. Set *Persistence Timeout* to **5**.
12. Set *Health Checks* to **Negotiate HTTP (HEAD)**.
13. Set the *Request to send* to **/mediator/main/loadBalancing/isExternallyAccessibleAPI**.

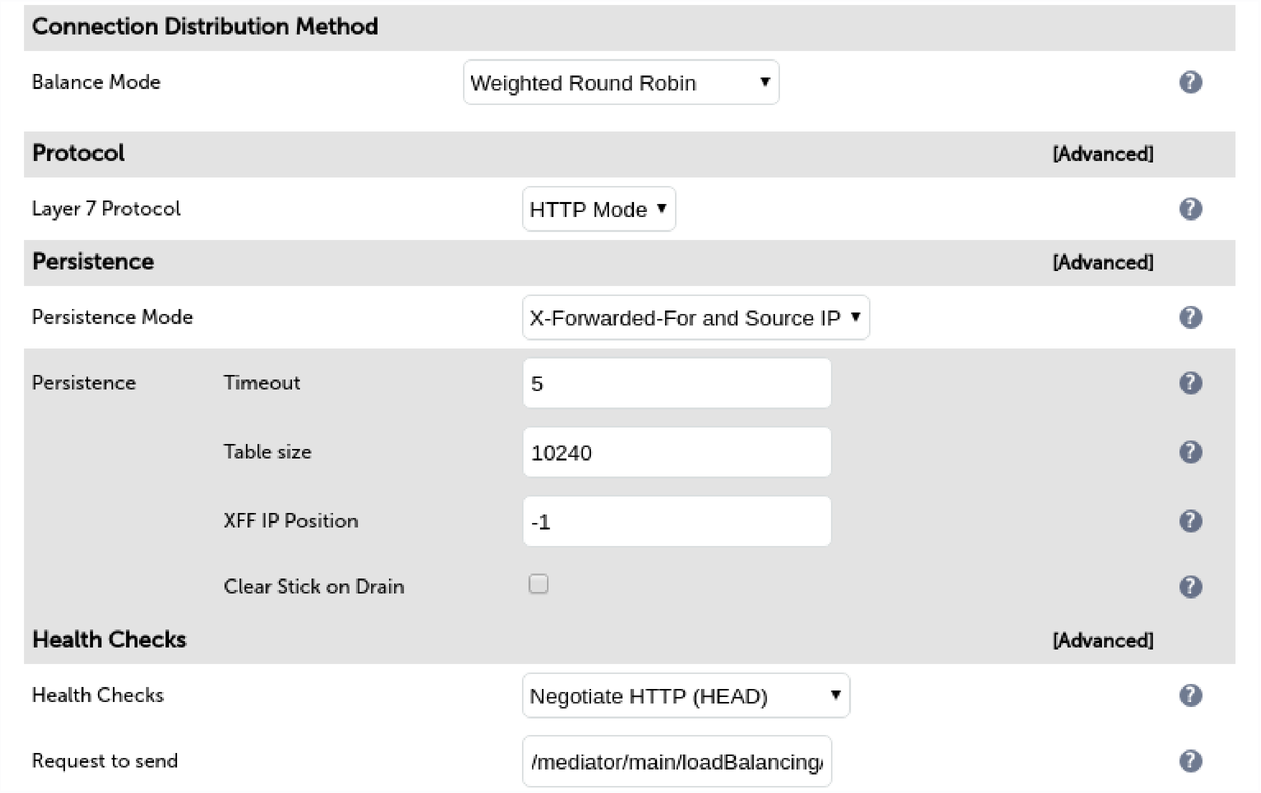

Connection Distribution Method		
Balance Mode		Weighted Round Robin ▾
Protocol		[Advanced]
Layer 7 Protocol		HTTP Mode ▾
Persistence		[Advanced]
Persistence Mode		X-Forwarded-For and Source IP ▾
Persistence	Timeout	5
	Table size	10240
	XFF IP Position	-1
	Clear Stick on Drain	□
Health Checks		[Advanced]
Health Checks		Negotiate HTTP (HEAD) ▾
Request to send		/mediator/main/loadBalancing

14. Click **Update**.

11.3. 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. Enter an appropriate name for the server in the *Label* field, e.g. **node-04**.

3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.24**.

4. Set the *Real Server Port* field to **80**.

5. Click **Update**.

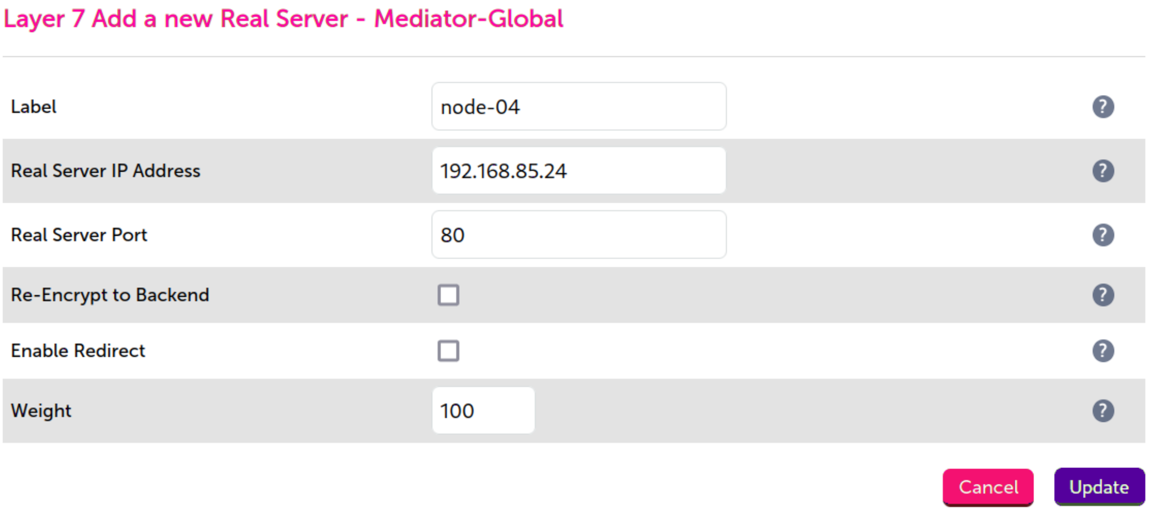

Layer 7 Add a new Real Server - Mediator-Global

Label	node-04
Real Server IP Address	192.168.85.24
Real Server Port	80
Re-Encrypt to Backend	□
Enable Redirect	□
Weight	100

6. Repeat these steps to add additional servers as required.

11.4. Setting Up the TLS/SSL Termination

Uploading a Certificate

An appropriate certificate must be present on the load balancer for TLS/SSL termination to work. Typically, a valid certificate is uploaded to the load balancer for use. The process for doing this is as follows:

1. Using the web user interface, navigate to *Cluster Configuration > SSL Certificate* and click on **Add a new SSL Certificate**.
2. Press the *Upload prepared PEM/PFX file* radio button.
3. Define the *Label* for the certificate as required, e.g. **Mediator-Certificate**.
4. Click on **Browse** and select the appropriate PEM or PFX style certificate.
5. If uploading a PFX certificate, enter the certificate's password in the *PFX File Password* field.
6. Click **Upload certificate**.

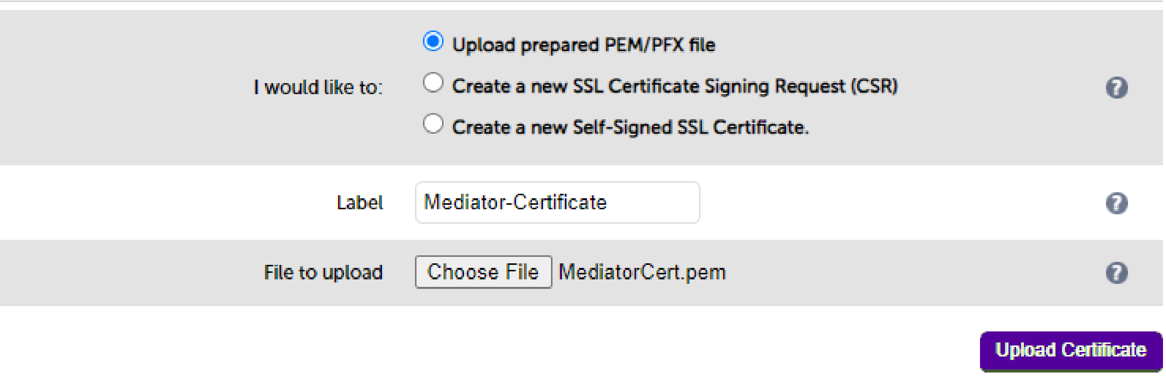

I would like to:
● Upload prepared PEM/PFX file
○ Create a new SSL Certificate Signing Request (CSR)
○ Create a new Self-Signed SSL Certificate.

Label: Mediator-Certificate

File to upload: MediatorCert.pem

For more information on creating PEM certificate files and converting between certificate formats please refer to Creating a PEM File.

In the absence of a valid certificate, it is also possible to create a certificate signing request (CSR) on the load balancer. A CSR can be submitted to a certificate authority for the issuance of a certificate. For more information on creating an CSR please refer to Generating a CSR on the Load Balancer.

Creating the TLS/SSL Termination

1. Using the web user interface, navigate to *Cluster Configuration > SSL Termination* and click on **Add a new Virtual Service**.
2. From the *Associated Virtual Service* drop-down list, select the **Mediator-Global** service which was created previously.
3. Set the *Virtual Service Port* field to **443**.
4. From the *SSL Certificate* drop-down list, select the appropriate certificate.
5. Click **Update** to create the TLS/SSL termination service.

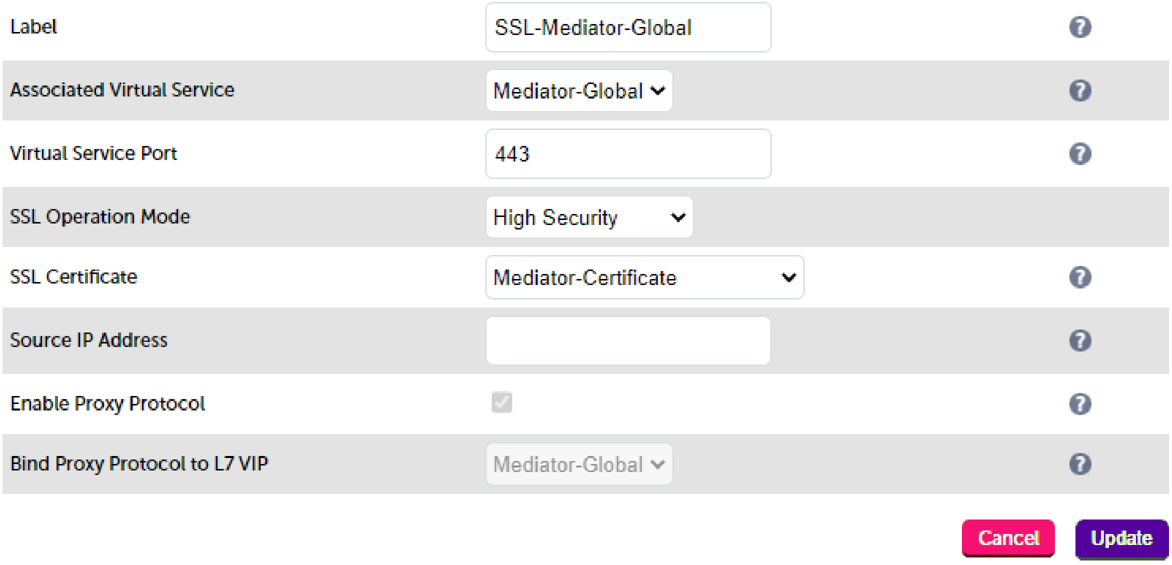

Label	SSL-Mediator-Global	?
Associated Virtual Service	Mediator-Global	?
Virtual Service Port	443	?
SSL Operation Mode	High Security	
SSL Certificate	Mediator-Certificate	?
Source IP Address		?
Enable Proxy Protocol	☑	?
Bind Proxy Protocol to L7 VIP	Mediator-Global	?

> Note
>
> If encountering issues accessing certain resources over HTTPS, refer to the earlier section Additional Changes when Adding TLS-Based Encryption.

11.5. Finalizing the Configuration

To apply the new settings, HAProxy and STunnel must both be reloaded. This can be done using the buttons 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**.
3. Click **Reload STunnel**.

12. Testing & Verification

> Note
>
> For additional guidance on diagnosing and resolving any issues you may have, please also refer to Diagnostics & Troubleshooting.

12.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the Mediator-X 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 four Mediator-X nodes are healthy and available to accept connections:

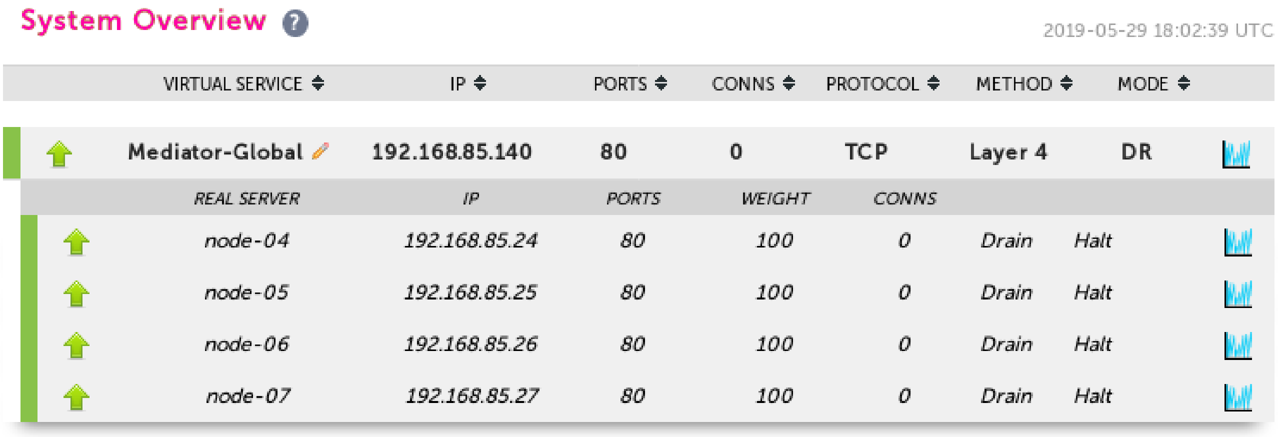

System Overview

2019-05-29 18:02:39 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
Mediator-Global	192.168.85.140	80	0	TCP	Layer 4	DR

REAL SERVER	IP	PORTS	WEIGHT	CONNS		
node-04	192.168.85.24	80	100	0	Drain	Halt
node-05	192.168.85.25	80	100	0	Drain	Halt
node-06	192.168.85.26	80	100	0	Drain	Halt
node-07	192.168.85.27	80	100	0	Drain	Halt

13. 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.

14. Further Documentation

For additional information, please refer to the Administration Manual.

15. Appendix

15.1. 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 configured first and then the Secondary should be added. Once the Primary and Secondary are paired, all load balanced services configured on the Primary are automatically replicated to the Secondary over the network using SSH/SCP.

> **Note**
>
> For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

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	All network settings including IP address(es), bonding configuration and VLANs
Local Configuration	Routing	Routing configuration including default gateways and static routes
Local Configuration	System Date & time	All time and date related settings
Local Configuration	Physical – Advanced Configuration	Various settings including Internet Proxy, Management Gateway, Firewall connection tracking table size, NIC offloading, SMTP relay, logging and Syslog Server
Local Configuration	Security	Appliance security settings
Local Configuration	SNMP Configuration	Appliance SNMP settings
Local Configuration	Graphing	Appliance graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Software Updates	Appliance software update management
Maintenance	Firewall Script	Appliance firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

Important Make sure that if these settings/updates have been configured on the Primary appliance, they're also configured on the Secondary appliance.

Adding a Secondary Appliance - Create an 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*.

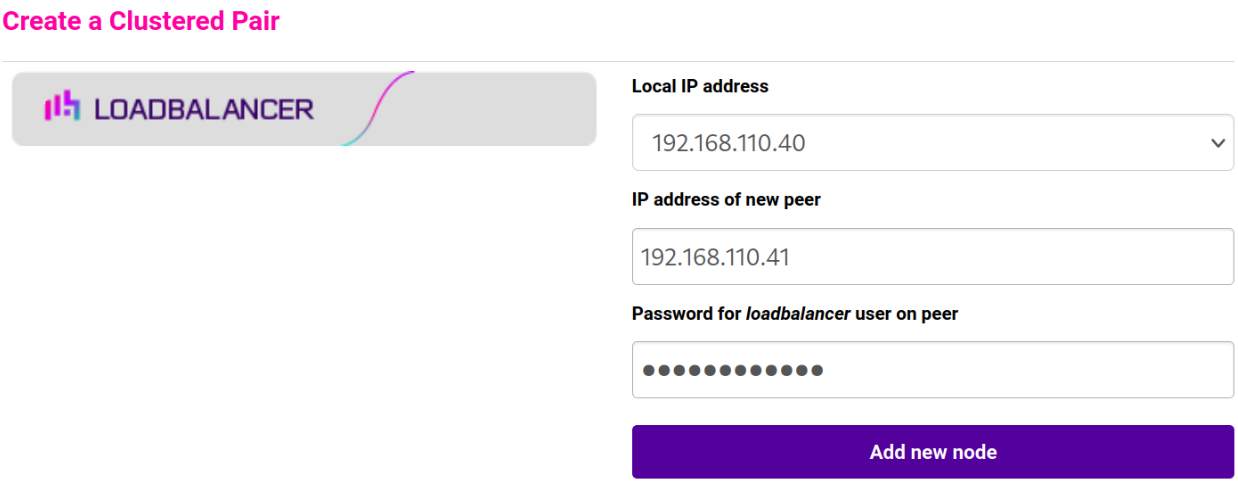
Create a Clustered Pair

Local IP address: 192.168.110.40

IP address of new peer: 192.168.110.41

Password for *loadbalancer* user on peer: ●●●●●●●●●●●●

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:

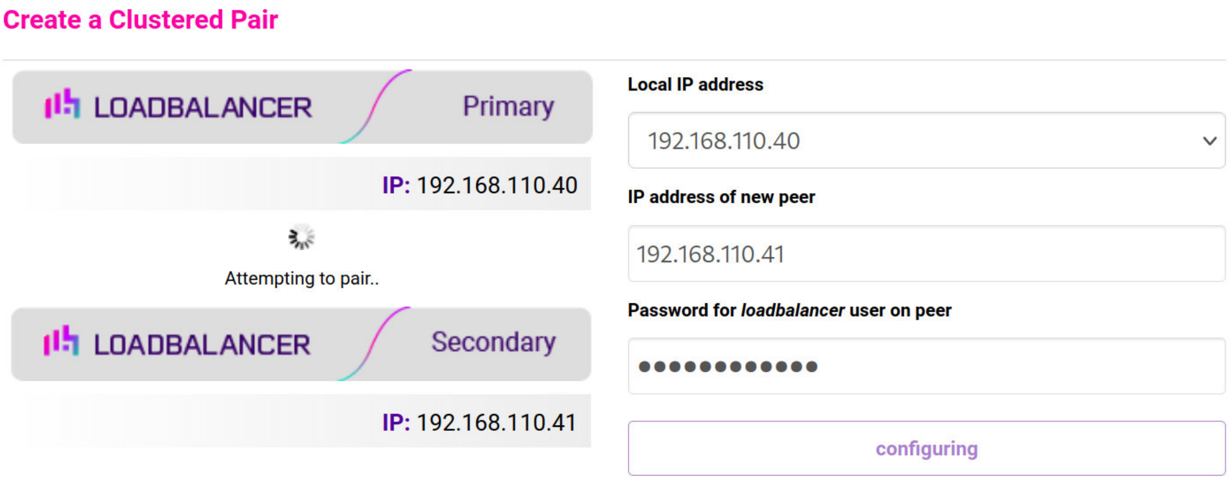
Create a Clustered Pair

Local IP address: 192.168.110.40

IP address of new peer: 192.168.110.41

Password for *loadbalancer* user on peer: ●●●●●●●●●●●●

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

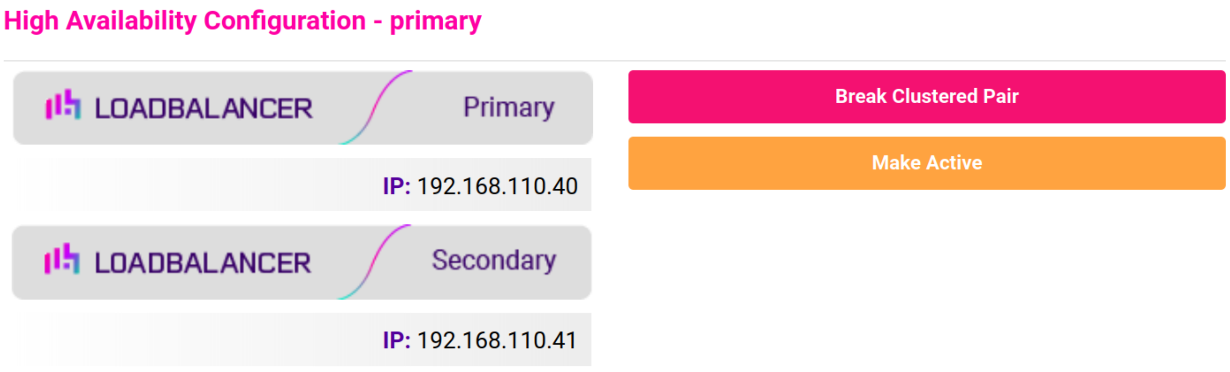

High Availability Configuration - primary

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.

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

> ℹ Note — For more details on configuring HA with 2 appliances, please refer to Appliance Clustering for HA.

> ℹ Note — For details on testing and verifying HA, please refer to Clustered Pair Diagnostics.

16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	30 May 2019	Initial version		AH
1.0.1	13 June 2019	Added section on enabling HAProxy multi-threaded mode Added a comment that the changes to the Mediator servers should be carried out by Evertz engineers Changed some terminology at the request of Evertz Changed the diagrams to reflect the new simplified configurations Changed the instructions and screenshots to reflect the new single virtual service configuration	Required updates	AH
1.1.0	24 July 2019	Styling and layout Changed the health checks at the request of Evertz	General styling updates Required updates	AH
1.1.1	1 August 2019	Made changes to section "Configuring Evertz Mediator-X for Load Balancing" at the request of Evertz	Required updates	AH
1.1.2	1 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.2.0	1 November 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.2.1	21 March 2022	Added new multithreading advice	Product change means multithreading is now enabled by default	AH
1.2.2	22 April 2022	Updated SSL related content to reflect latest software version	New software release	RJC

Version	Date	Change	Reason for Change	Changed By
1.2.3	6 July 2022	Added new advice on allowing Mediator-X resources to be accessible over HTTPS	Feedback from a customer deployment	AH
1.2.4	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.2.5	5 January 2023	Combined software version information into one section Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section Removed references to the colour of certain UI elements	Housekeeping across all documentation	AH
1.2.6	2 February 2023	Updated screenshots	Branding update	AH
1.2.7	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.3.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH
1.3.1	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.