

Load Balancing VMware App Volumes

Version 1.1.0



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. VMware App Volumes	3
4. VMware App Volumes	3
5. Load Balancing VMware App Volumes	4
5.1. Persistence (aka Server Affinity)	4
5.2. Port Requirements	4
5.3. TLS/SSL Termination	4
Server Certificates	4
5.4. TLS/SSL Offloading with Re-encrypt to the Backend (SSL Bridging)	4
5.5. Health Checks	4
6. Deployment Concept	4
6.1. Virtual Service (VIP) Requirements	5
7. Loadbalancer.org Appliance – the Basics	5
7.1. Virtual Appliance	5
7.2. Initial Network Configuration	6
7.3. Accessing the Appliance WebUI	6
Main Menu Options	7
7.4. Appliance Software Update	8
Determining the Current Software Version	8
Checking for Updates using Online Update	8
Using Offline Update	8
7.5. Ports Used by the Appliance	9
7.6. HA Clustered Pair Configuration	10
8. Appliance Configuration for VMware App Volumes	10
8.1. Configuring the Virtual Service (VIP)	10
8.2. Defining the Real Servers (RIPs)	11
8.3. Finalizing the Configuration	12
9. Additional Configuration Options & Settings	12
9.1. TLS/SSL Termination	12
Backend Encryption at the Virtual Service Level (Recommended)	12
Backend Encryption at the Real Server Level	13
Uploading a Certificate	13
Creating the TLS/SSL Termination Service (SSL Offloading)	14
Creating The TLS/SSL Termination Service (SSL Bridging)	14
9.2. Enabling 'Force to HTTPS'	15
9.3. Finalizing the Configuration	15
10. Testing & Verification	15
10.1. Using System Overview	15
11. Technical Support	16
12. Further Documentation	16
13. Appendix	17
13.1. Configuring HA - Adding a Secondary Appliance	17
Non-Replicated Settings	17
Adding a Secondary Appliance - Create an HA Clustered Pair	18
14. Document Revision History	20

1. About this Guide

This guide details the steps required to configure a load balanced VMware App Volumes environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any VMware App Volumes 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 VMware App Volumes. 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

- V8.6.1 and later

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. VMware App Volumes

- V2.x
- V3.x
- V4.x

4. VMware App Volumes

VMware App Volumes is a real-time application delivery system that enterprises can use to dynamically deliver and manage applications.

Applications are packaged and delivered by attaching a standard VMDK or VHD file to a virtual machine. You can centrally manage the applications with the App Volumes Manager, a Web-based interface that is integrated with Active Directory (AD) and vSphere. Administrators can assign, update, or remove applications to be delivered at the next user login without the need to modify the desktops or disrupt users while they are working.

Writable Volumes allow users to access their application data across sessions and devices.



5. Load Balancing VMware App Volumes

Note

It's highly recommended that you have a working VMware App Volumes environment first before implementing the load balancer.

5.1. Persistence (aka Server Affinity)

HTTP cookie persistence is required when load balancing VMware App Volumes.

5.2. Port Requirements

The following table shows the ports that are load balanced:

Port	Protocols	Use
80	TCP/HTTP	App Volumes HTTP traffic
443	TCP/HTTPS	App Volumes Secure HTTP traffic

5.3. TLS/SSL Termination

Server Certificates

TLS/SSL server certificates in either PFX or PEM file format can be uploaded to the Load balancer using the certificate upload. Alternatively, you can create a Certificate Signing Request (CSR) on the load balancer and send this to your CA to create a new certificate. For more information please refer to [Generating a CSR on the Load Balancer](#).

5.4. TLS/SSL Offloading with Re-encrypt to the Backend (SSL Bridging)

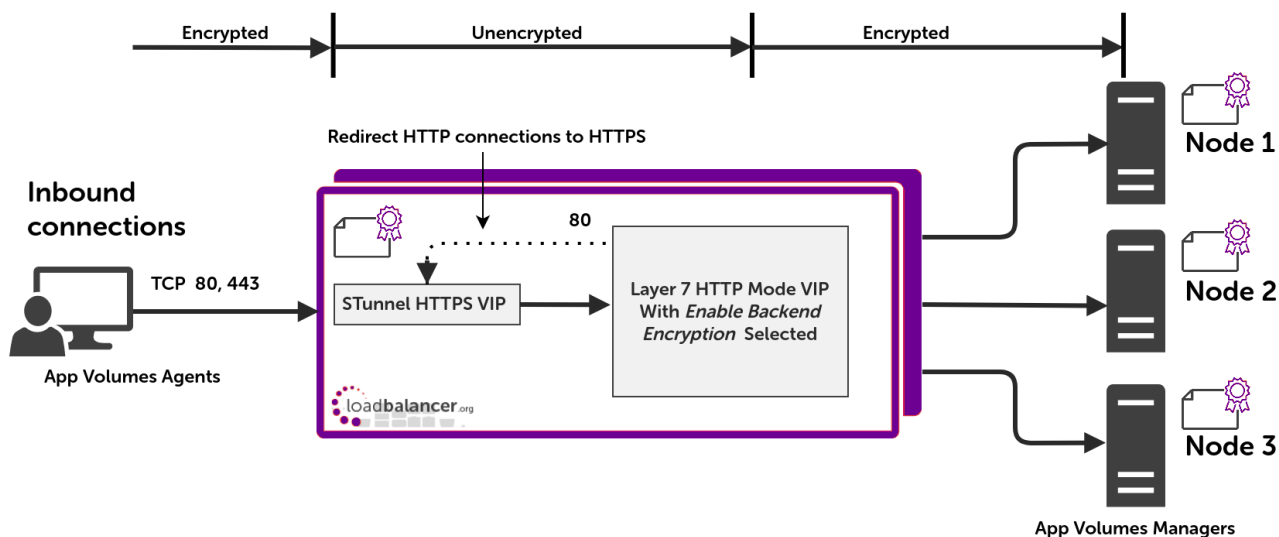
Terminating SSL on the load balancer is only necessary when using cookie based persistence for the primary protocol connections. Cookie based persistence is only needed when source IP address persistence cannot be used due to inline NAT/proxy devices hiding client source IP addresses. If SSL offload is used, the load balancer and the App Volumes servers **must** have the same certificate.

5.5. Health Checks

The load balancer is configured to check the health of each App Volumes server by periodically sending an HTTPS GET request to the location [/health_check](#) with the [Host](#) header appropriate for the FQDN in use. It will perform this HTTPS GET and expect a [200 OK](#) response. If it receives a response other than [200 OK](#) or doesn't get any response, that server will be marked as down and will not attempt to route client requests to it. It will continue to poll so that it can detect when it is available again.

6. Deployment Concept





VIPs = **V**irtual **I**P 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 the section [Configuring HA - Adding a Secondary Appliance](#) in the appendix for more details on configuring a clustered pair.

6.1. Virtual Service (VIP) Requirements

To provide load balancing and HA for VMWare App Volumes, a single VIP is required:

- VMWare App Volumes

7. Loadbalancer.org Appliance – the Basics

7.1. 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.




7.2. 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.

7.3. 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](#).

 **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/`

 **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>

 **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. 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, RIPv 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 & RIPv

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

7.4. 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

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 v8.9.0 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.

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.





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: No file chosen

Checksum: 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.

7.5. 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



7.6. 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.

8. Appliance Configuration for VMware App Volumes

8.1. 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. **App Volumes HTTP**.
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	<input type="text" value="App Volumes HTTP"/>	?
IP Address	<input type="text" value="192.168.85.150"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **HTTP Cookie**.
9. In the *Health Checks* section click **Advanced** to expand the menu.
10. Set *Health Checks* to **Negotiate HTTPS (GET)**.
11. Set *Request to send* to **/health_check**
12. Set *Check Port* to **443**.
13. Set *Host Header* to the FQDN of the App Volumes deployment, e.g. **app-volumes.example.com**

Persistence		[Advanced +]
Persistence Mode	HTTP Cookie	?
HTTP Cookie Name	SERVERID	?
Health Checks		[Advanced -]
Health Checks	Negotiate HTTPS (GET)	?
Request to send	/health_check	?
Response expected	Equals	?
Check Port	443	?
Username		?
Host Header	app-volumes.example.com	?
Password *		?

14. Under **SSL** check the **Enable Backend Encryption** checkbox.

SSL	
Enable Backend Encryption	<input checked="" type="checkbox"/> ?

15. Click **Update**.

8.2. Defining the Real Servers (RIPs)

- 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.
- Define the **Label** for the real server as required, e.g. **App Vol Srv 1**.
- Set the **Real Server IP Address** field to the required IP address, e.g. **192.168.85.200**.
- Set the **Real Server Port** field to **443**.
- Ensure that the **Re-Encrypt to Backend** checkbox is checked.
- Click **Update**.
- Repeat these steps to add additional App Volumes servers as required.

Layer 7 Add a new Real Server - App_Volumes_HTTP

Label	App Vol Srv 1	?
Real Server IP Address	192.168.85.200	?
Real Server Port	443	?
Re-Encrypt to Backend	<input checked="" type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	100	?



8.3. 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**.

9. Additional Configuration Options & Settings

9.1. TLS/SSL Termination

The load balancer can handle TLS/SSL termination in the following ways:

1. On the App Volumes servers, aka **SSL Pass-through**
2. On the load balancer, aka **SSL Offloading**
3. On the load balancer with re-encryption to the App Volumes servers, aka **SSL Bridging** (recommended)

In the "bridging" case, a TLS termination service utilizing stunnel (default and recommended) is configured on the appliance and a server certificate is uploaded and associated to the virtual service. Data is encrypted from the client to the load balancer and is also encrypted from the load balancer to the backend servers.

Notes on "SSL Bridging":

- This is similar to 'SSL Offload', the only difference is that the connection from the load balancer to the App Volumes servers is encrypted using the certificate located on the App Volumes server. This could be a self-signed certificate since no client connections are terminated here, only at the stunnel service.
- This mode can be enabled for the entire VIP and all associated App Volumes servers using the VIP option **Enable Backend Encryption** or per-App Volumes server using the **Re-Encrypt to Backend** option as detailed below.

Note

Performing TLS/SSL termination on the load balancer can be very CPU intensive. In most cases, for a scalable solution, terminating TLS/SSL on the App Volumes servers is the best option.

Backend Encryption at the Virtual Service Level (Recommended)

To enable re-encryption at the virtual service level:

1. Using the web user interface, navigate to **Cluster Configuration > Layer 7 – Virtual Services**.
2. Click **Modify** next to the App Volumes VIP.
3. Under **SSL** check the **Enable Backend Encryption** checkbox.



4. Click **Update**.

Backend Encryption at the Real Server Level

To enable re-encryption at the real server level, on a server-by-server basis:

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers*.
2. Click **Modify** next to the first App Volumes server.
3. Ensure that the **Re-Encrypt to Backend** checkbox is checked.
4. Click **Update**.
5. Repeat these steps to enable backend encryption for each App Volumes server as required.

Layer 7 Add a new Real Server - App_Volumes_HTTP

Label	<input type="text" value="App Vol Srv 1"/>	?
Real Server IP Address	<input type="text" value="192.168.85.200"/>	?
Real Server Port	<input type="text" value="443"/>	?
Re-Encrypt to Backend	<input checked="" type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

Uploading a Certificate

An appropriate server 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. **AppVolumes_Cert**.
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**.

Add a new SSL 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

AppVolumes_Cert

File to upload

Browse... No file selected.

Upload Certificate

Further information on creating PEM files and converting between certificate formats can be found in our *Administration Manual*.

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. Instructions on creating a CSR can be found in our *Administration Manual*.

Our *Administration Manual* can be found at: <https://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf>

Creating the TLS/SSL Termination Service (SSL Offloading)

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 App Volumes service that 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.

SSL Termination - Add a new Virtual Service

Label	SSL_App_Volumes_HTTP	?
Associated Virtual Service	App_Volumes_HTTP ▾	?
Virtual Service Port	443	?
SSL Operation Mode	High Security ▾	
SSL Certificate	appvolumes_cert ▾	?
Source IP Address		?
Enable Proxy Protocol	<input checked="" type="checkbox"/>	?
Bind Proxy Protocol to L7 VIP	App_Volumes_HTTP ▾	?

CancelUpdate

Creating The TLS/SSL Termination Service (SSL Bridging)



1. To configure "SSL Bridging" (TLS/SSL with re-encryption to the backend), follow the steps as per [Section 9.1.1, "Backend Encryption at the Virtual Service Level \(Recommended\)"](#) and enable re-encrypt to backend on the VIP.
2. Ensure that the real servers have port 443 defined with an SSL certificate installed.
3. Upload an SSL certificate to the load balancer as per the step above.
4. Create an SSL termination as per the steps above.
5. Observe that the system overview page now displays real servers with 're-encrypt to backend' padlock icons.

9.2. Enabling 'Force to HTTPS'

It is possible to force the use of TLS and disallow the use of plaintext HTTP. For security and privacy reasons, this is strongly recommended for any traffic that travels over the public internet. This is achieved by forcing all clients that connect using plaintext HTTP to reconnect using HTTPS.

1. Using the web user interface, navigate to **Cluster Configuration > Layer 7 – Virtual Services** and click on **Modify** next to the App Volumes VIP.
2. In the **Other** section click **Advanced** to expand the menu.
3. Set **Force to HTTPS** to **Yes**.
4. Click **Update**.

Other		[Advanced -]
Maximum Connections	<input type="text" value="40000"/>	?
Timeout	<input type="checkbox"/>	?
Set X-Forward-For header	<input checked="" type="checkbox"/>	?
Force to HTTPS	<input checked="" type="radio"/> Yes <input type="radio"/> No	?

9.3. Finalizing the Configuration

To apply the new settings, HAProxy and stunnel must both be reloaded as follows:

1. Using the WebUI, navigate to: **Maintenance > Restart Services** and click **Reload STunnel**
2. Using the WebUI, navigate to: **Maintenance > Restart Services** and 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 App Volumes 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 App Volumes nodes are healthy and available to accept connections:



	VIRTUAL SERVICE ⌵	IP ⌵	PORTS ⌵	CONNS ⌵	PROTOCOL ⌵	METHOD ⌵	MODE ⌵	
↑	App_Volumes_HTTP..	192.168.85.150	80	0	HTTP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	 App_Vol_Srv_1	192.168.85.200	443	100	0	Drain	Halt	
↑	 App_Vol_Srv_2	192.168.85.201	443	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 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

 **LOADBALANCER**

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41


Password for *loadbalancer* user on peer

••••••••••

Add new node

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

 **LOADBALANCER**

Primary

IP: 192.168.110.40

Attempting to pair..

 **LOADBALANCER**

Secondary

IP: 192.168.110.41

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41

Password for *loadbalancer* user on peer





••••••••••

configuring

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



High Availability Configuration - primary

 LOADBALANCER  Primary	Break Clustered Pair
IP: 192.168.110.40	Make Active
 LOADBALANCER  Secondary	
IP: 192.168.110.41	

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	25 September 2020	Initial creation		IBG
1.1.0	18 April 2023	Added instructions for enabling 'Force to HTTPS' Converted the document to AsciiDoc Significant updates to bring the document into line with current documentation format New document theme Modified diagram colours	'Force to HTTPS' was referenced but not previously documented step-by-step Document updates required moving it to the new documentation system Branding update	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](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.

