

Load Balancing Hitachi Content Platform

Version 1.2.0



Table of Contents

1. About this Brief	3
2. Loadbalancer.org Appliances Supported	3
3. Software Versions Supported	3
3.1. Loadbalancer.org Appliance	3
3.2. HCP	3
4. Hitachi Content Platform	3
5. Load Balancing HCP	3
5.1. Load Balancing & HA Requirements	4
5.2. Virtual Service (VIP) Requirements	4
6. Deployment Concept	4
7. Load Balancer Deployment Methods	4
7.1. Layer 7 SNAT Mode	5
8. Loadbalancer.org Appliance – the Basics	5
8.1. Virtual Appliance	6
8.2. Initial Network Configuration	6
8.3. Accessing the Appliance WebUI	6
8.3.1. Main Menu Options	8
8.4. Appliance Software Update	8
8.4.1. Online Update	8
8.4.2. Offline Update	9
8.5. Ports Used by the Appliance	9
8.6. HA Clustered Pair Configuration	10
9. Appliance Configuration for HCP	10
9.1. Layer 7 Timeouts	10
9.2. Virtual Services	10
9.2.1. VIP 1 - HCP-Data	10
9.2.2. VIP 2 - HCP-MQE	12
9.2.3. VIP 3 - HCP-NSB	13
9.3. Finalizing the Configuration	14
10. Optional Multi-Site Failover	14
11. Testing & Verification	15
11.1. Accessing the Application	15
11.2. Using System Overview	15
12. Technical Support	16
13. Further Documentation	16
14. Appendix	17
14.1. Configuring HA - Adding a Secondary Appliance	17
14.1.1. Non-Replicated Settings	17
14.1.2. Configuring the HA Clustered Pair	18
15. Document Revision History	20

1. About this Brief

This brief outlines the steps required to configure a load balanced Hitachi Content Platform (HCP) environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any HCP 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 HCP. For full specifications of available models please refer to <https://www.loadbalancer.org/products/enterprise>.

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

- V8.13.0 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. HCP

- v9.7 and later

4. Hitachi Content Platform

Hitachi Content Platform is a scalable, multi-purpose enterprise storage solution that delivers object storage, file services, and archival capabilities. Designed for cloud-scale deployments, HCP is engineered to provide robust performance, high availability, and comprehensive data management. The platform integrates seamlessly with load balancing solutions to distribute traffic across multiple nodes, ensuring optimal resource utilization and fault tolerance.

5. Load Balancing HCP

Note

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



5.1. Load Balancing & HA Requirements


Using a load balancer with HCP helps deliver maximum performance and availability by evenly distributing workload across all HCP nodes. It also ensures that requests are only sent to healthy nodes.

5.2. Virtual Service (VIP) Requirements

To provide load balancing and HA for HCP, the following VIPs are required:

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	HCP-Data	L7 SNAT (TCP)	443	Source IP	HTTPS (GET)
VIP 2	HCP-MQE	L7 SNAT (TCP)	8888	Source IP	HTTPS (GET)
VIP 3	HCP-NSB	L7 SNAT (TCP)	8000	Source IP	HTTPS (GET)

6. Deployment Concept



Note

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

7. Load Balancer Deployment Methods


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

For HCP, using layer 7 SNAT mode is recommended. This mode is described below and is used for the configurations presented in this guide. For configuring using layer 7 SNAT mode please refer to the section [Appliance Configuration for HCP](#).



7.1. Layer 7 SNAT Mode

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



- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network 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, **eth1** is typically used for client side connections and **eth0** is used for Real Server connections, although this is not mandatory since any interface can be used for any purpose.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

8. Loadbalancer.org Appliance – the Basics



8.1. Virtual Appliance

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

Note

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

Note

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

Note

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

8.2. Initial Network Configuration

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

Important

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

8.3. Accessing the Appliance WebUI

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

Note

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

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

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

Note

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

Note

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



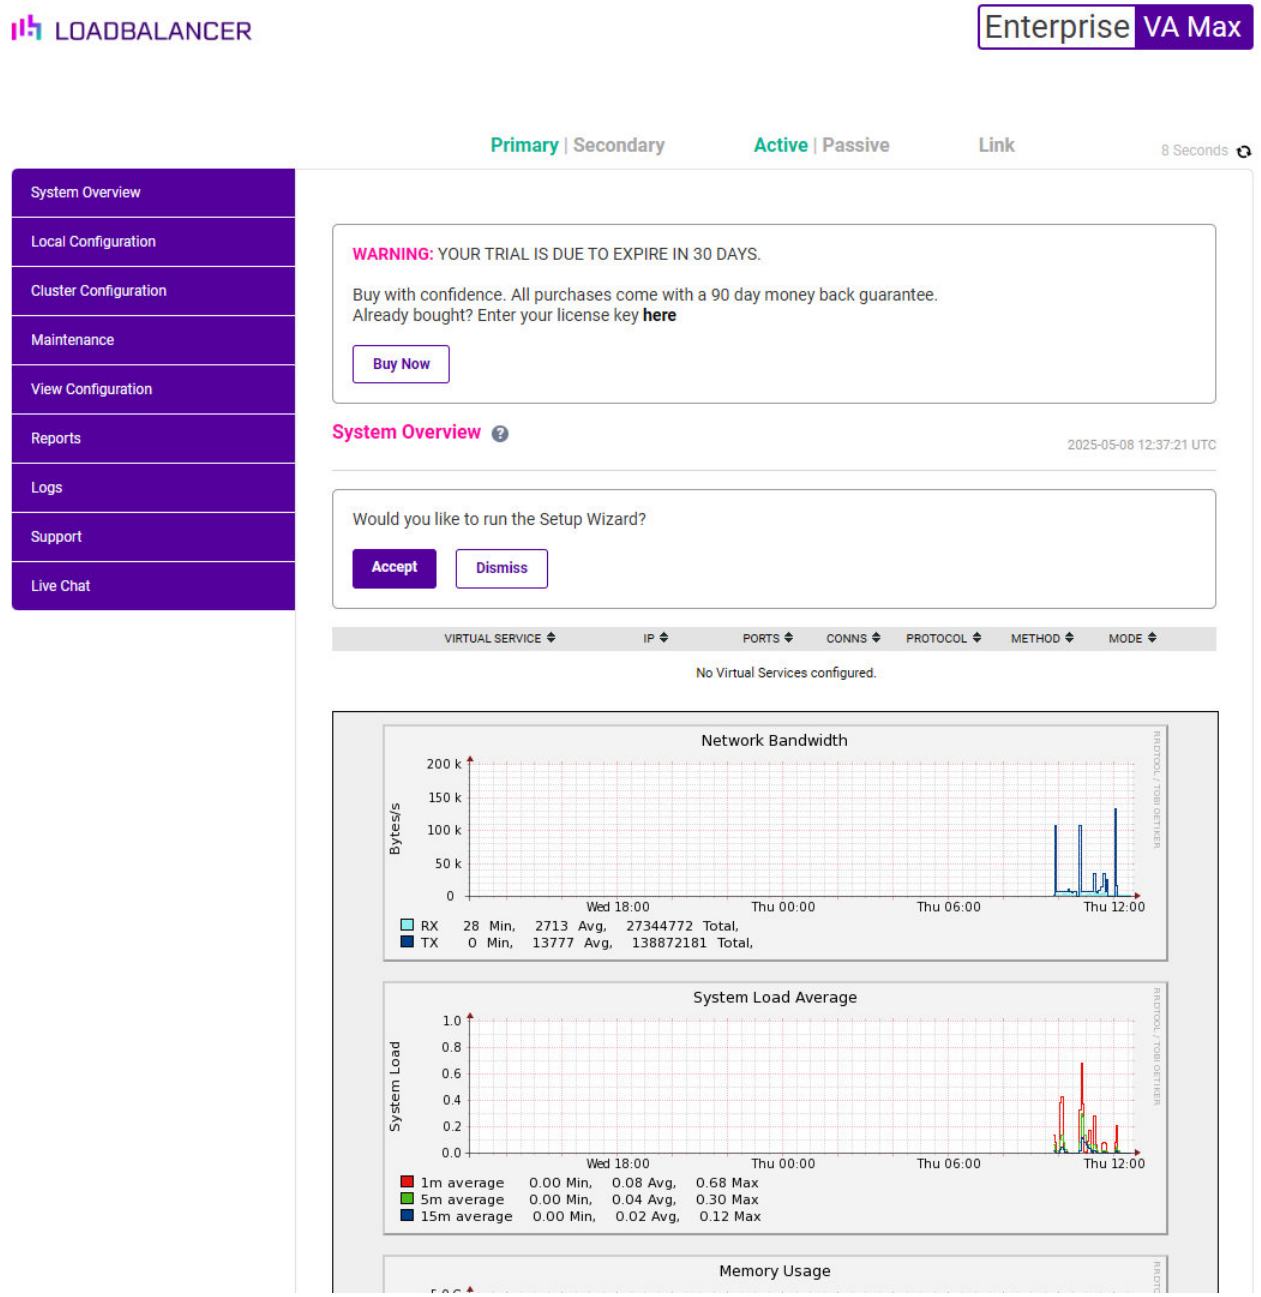
2. Log in to the WebUI using the following credentials:

Username: loadbalancer

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

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

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



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.



8.3.1. 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 creating 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

We recommend that the appliance is kept up to date to ensure that you benefit from the latest bug fixes, security updates and feature improvements. Both online and offline update are supported.

Note

For full details, please refer to [Appliance Software Update](#) in the Administration Manual.

Note

Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

8.4.1. Online Update

The appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. This is the default behavior and can be disabled if preferred. If an update is found, a notification similar to the example below will be displayed at the top of the WebUI:

Information: Update 8.13.5 is now available for this appliance.

Online Update

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.



Important

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

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

Information: Update completed successfully. Return to **system overview**.

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



8.4.2. Offline Update

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

To check for the latest version, please refer to our product roadmap page available [here](#). To obtain the latest offline update files contact support@loadbalancer.org.

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.

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	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 for ADC Portal comms



Protocol	Port	Purpose
TCP	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback server
TCP	9443	WebUI - HTTPS
TCP	25565	Shuttle service for ADC Portal comms

 **Note**

All ports listed above except port 123 (NTP) can be changed if required.

- To change the port used for heartbeat, refer to [Configuring High Availability](#)
- To change the port used for HAProxy replication, refer to [Layer 7 - Advanced Configuration](#)
- To change other ports, 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 HCP

9.1. Layer 7 Timeouts

Since transfers can be quite large, the client & server timeouts may need to be increased. By default, these are set to 42 seconds and 43 seconds respectively. If these timeouts are not enough, they can both be increased to a suitable value - for example 2 minutes using the WebUI menu option: **Cluster Configuration > Layer 7 - Advanced Configuration** and setting **Client Timeout & Real Server Timeout** to **2m** as shown below:

Connection Timeout	<input type="text" value="4000"/>	ms	
Client Timeout	<input type="text" value="2m"/>	ms	
Real Server Timeout	<input type="text" value="2m"/>	ms	

9.2. Virtual Services

9.2.1. VIP 1 - HCP-Data

Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="HCP-Data"/>	?
IP Address	<input type="text" value="10.30.20.143"/>	?
Ports	<input type="text" value="443"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

2. Enter a suitable *Label* (name) for the Virtual Service, e.g. **HCP-Data**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.30.20.143**.
4. Set *Ports* to **443**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update**.
7. Now click **Modify** next to the newly created VIP.
8. Scroll down to the *Health Checks* section.
 - Set the *Health Checks* to **HTTPS (GET)**.
 - Set the *Request to send* to **/node_status**
9. Click **Update**.

Configure the Associated Real Servers (RIPs)

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

Label	<input type="text" value="HCP-Node1"/>	?
Real Server IP Address	<input type="text" value="10.30.20.150"/>	?
Real Server Port	<input type="text" value="443"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

2. Enter a suitable *Label* (name) for the Real Server, e.g. **HCP-Node1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **10.30.20.150**.

4. Set the *Real Server Port* to **443**.
5. Click **Update**.
6. Repeat these steps to add additional Real Servers as required.

9.2.2. VIP 2 - HCP-MQE

Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="HCP-MQE"/>	?
IP Address	<input type="text" value="10.30.20.144"/>	?
Ports	<input type="text" value="8888"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

2. Enter a suitable *Label* (name) for the Virtual Service, e.g. **HCP-MQE**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.30.20.144**.
4. Set *Ports* to **8888**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update**.
7. Now click **Modify** next to the newly created VIP.
8. Scroll down to the *Health Checks* section.
 - Set the *Health Checks* to **HTTPS (GET)**.
 - Set the *Request to send* to **/node_status**
9. Click **Update**.

Configure the Associated Real Servers (RIPs)

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

Label	<input type="text" value="HCP-Node1"/>	?
Real Server IP Address	<input type="text" value="10.30.20.150"/>	?
Real Server Port	<input type="text" value="8888"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

Cancel
Update

2. Enter a suitable *Label* (name) for the Real Server, e.g. **HCP-Node1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **10.30.20.150**.
4. Set the *Real Server Port* to **8888**.
5. Click **Update**.
6. Repeat these steps to add additional Real Servers as required.

9.2.3. VIP 3 - HCP-NSB

Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="HCP-NSB"/>	?
IP Address	<input type="text" value="10.30.20.145"/>	?
Ports	<input type="text" value="8000"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input style="border: 1px solid #ccc;" type="text" value="TCP Mode"/>	?

Cancel
Update

2. Enter a suitable *Label* (name) for the Virtual Service, e.g. **HCP-NSB**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.30.20.145**.
4. Set *Ports* to **8000**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update**.
7. Now click **Modify** next to the newly created VIP.

8. Scroll down to the *Health Checks* section.
 - Set the *Health Checks* to **HTTPS (GET)**.
 - Set the *Request to send* to **/node_status**
9. Click **Update**.

Configure the Associated Real Servers (RIPs)

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

Label	<input type="text" value="HCP-Node1"/>	?
Real Server IP Address	<input type="text" value="10.30.20.150"/>	?
Real Server Port	<input type="text" value="8000"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

2. Enter a suitable *Label* (name) for the Real Server, e.g. **HCP-Node1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **10.30.20.150**.
4. Set the *Real Server Port* to **8000**.
5. Click **Update**.
6. Repeat these steps to add additional Real Servers as required.

9.3. Finalizing the Configuration

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

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

10. Optional Multi-Site Failover

If an HCP deployment is made up of two unique clusters at two sites, with replication taking place across a multi-site link, it is possible to set up a failover mechanism in the event that one HCP deployment should become unavailable.

For each Virtual Service:



1. Click **Modify** next to the VIP in question.
2. Under the **Fallback Server** section, set the **IP Address** to the VIP address of the HCP service at the other site, e.g. **192.168.1.1**.
3. Ensure that the **Port** field is empty.
4. Click **Update**.

Fallback Server		[Advanced +]
IP Address	<input type="text" value="192.168.1.1"/>	
Port	<input type="text"/>	

11. Testing & Verification

Note

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































11.1. Accessing the Application

The following tests can be used to verify the configuration:

Test	Action	Expected Result
S3 PUT/GET	Run the command: aws --endpoint-url https://hcp.example.com s3 cp file.txt s3://namespace/	Upload & download succeed; aws s3 ls lists object
Namespace Browser	Browse to: https://hcp.example.com:8000/	GUI loads, login succeeds
MQE query	Run the command: curl -k https://hcp.example.com:8888/query	Returns valid JSON/XML
Persistence check	Upload a multi-part object and watch the same RIP handle all parts	The client should remain connected to the same RIP

11.2. Using System Overview

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

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	HCP-Data	10.30.20.143	443	0	TCP	Layer 7	Proxy	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	HCP-Node1	10.30.20.150	443	100	0	Drain	Halt	
	HCP-Node2	10.30.20.151	443	100	0	Drain	Halt	
	HCP-Node3	10.30.20.152	443	100	0	Drain	Halt	
	HCP-Node4	10.30.20.153	443	100	0	Drain	Halt	
	HCP-MQE	10.30.20.144	8888	0	TCP	Layer 7	Proxy	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	HCP-Node1	10.30.20.150	8888	100	0	Drain	Halt	
	HCP-Node2	10.30.20.151	8888	100	0	Drain	Halt	
	HCP-Node3	10.30.20.152	8888	100	0	Drain	Halt	
	HCP-Node4	10.30.20.153	8888	100	0	Drain	Halt	
	HCP-NSB	10.30.20.145	8000	0	TCP	Layer 7	Proxy	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	HCP-Node1	10.30.20.150	8000	100	0	Drain	Halt	
	HCP-Node2	10.30.20.151	8000	100	0	Drain	Halt	
	HCP-Node3	10.30.20.152	8000	100	0	Drain	Halt	
	HCP-Node4	10.30.20.153	8000	100	0	Drain	Halt	

12. 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.

13. Further Documentation

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

14. Appendix

14.1. 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.

14.1.1. 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	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.

14.1.2. Configuring the HA Clustered Pair

(!) Important

During HA pairing, all WebUI users and passwords are synchronized from the Primary to the Secondary. After clustering completes (you will be logged out of the Secondary when this occurs), the Primary's credentials should be used to login to both nodes.

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
10.11.40.55

IP address of new peer
10.11.40.56

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: 10.11.40.55

Attempting to pair..

LOADBALANCER Secondary IP: 10.11.40.56

Local IP address: 10.11.40.55

IP address of new peer: 10.11.40.56

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 IP: 10.11.40.55

LOADBALANCER Secondary IP: 10.11.40.56

Break Clustered Pair

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

Note

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

Note

For more details on configuring HA with 2 appliances, please refer to [Configuring High Availability](#).

Note

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

15. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	14 July 2022	Initial version		AH
1.0.1	5 January 2023	<p>Combined software version information into one section</p> <p>Added one level of section numbering</p> <p>Added software update instructions</p> <p>Added table of ports used by the appliance</p> <p>Reworded 'Further Documentation' section</p>	Housekeeping across all documentation	AH
1.0.2	2 February 2023	Updated screenshots	Branding update	AH
1.0.3	7 March 2023	<p>Added the section "Finalizing the Configuration" to ensure HAProxy is explicitly reloaded</p> <p>Removed conclusion section</p>	<p>Provided clarity for reloading HAProxy post-configuration</p> <p>Updates across all documentation</p>	AH
1.1.0	24 March 2023	<p>New document theme</p> <p>Modified diagram colours</p>	Branding update	AH
1.2.0	19 June 2025	Various updates and improvements	Document maintenance	RJC





Visit us: www.loadbalancer.org

Phone us: +44 (0)330 380 1064

Phone us: +1 833 274 2566

Email us: info@loadbalancer.org

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

About Loadbalancer.org

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

