

Load Balancing Cloudian HyperFile

Version 1.2.0



Table of Contents

1. About this Guide	
2. Loadbalancer.org Appliances Supported	
3. Software Versions Supported	
3.1. Loadbalancer.org Appliance	
3.2. Cloudian HyperFile	
4. Cloudian HyperFile	
5. Load Balancing Cloudian HyperFile	4
5.1. Load Balancing & HA Requirements	4
5.2. Persistence (aka Server Affinity)	4
5.3. Virtual Service (VIP) Requirements	4
5.4. Port Requirements	4
6. Deployment Concept	4
7. Load Balancer Deployment Methods	
7.1. Layer 4 DR Mode	5
7.2. Layer 7 SNAT Mode	6
7.3. Our Recommendation	
8. Configuring Cloudian HyperFile for Load Balancing	
8.1. Configuring for Layer 4 DR Mode.	
8.2. Configuring for Layer 7 SNAT Mode	
9. Loadbalancer.org Appliance – the Basics	
9.1. Virtual Appliance	
9.2. Initial Network Configuration	
9.3. Accessing the Appliance WebUI	
Main Menu Options	10
9.4. Appliance Software Update	11
Determining the Current Software Version	
Checking for Updates using Online Update	11
Using Offline Update	11
9.5. Ports Used by the Appliance	
9.6. HA Clustered Pair Configuration	
10. Appliance Configuration for Cloudian HyperFile – Using Layer 4 DR Mode	13
10.1. Configuring the Virtual Service (VIP)	
10.2. Defining the Real Servers (RIPs)	
11. Appliance Configuration for Cloudian HyperFile – Using Layer 7 SNAT Mode	
11.1. Configuring the Virtual Service (VIP)	
11.2. Defining the Real Servers (RIPs)	
11.3. Finalizing the Configuration	
12. Testing & Verification	
12.1. Using System Overview.	
12.2. Layer 4 DR mode specific test	
13. Technical Support	
14. Further Documentation	
15. Appendix	
15.1. Configuring HA - Adding a Secondary Appliance	
Non-Replicated Settings	
Adding a Secondary Appliance - Create an HA Clustered Pair	
16. Document Revision History	21

1. About this Guide

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

8 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. Cloudian HyperFile

Version 3.6.1 and later

4. Cloudian HyperFile

Cloudian HyperFile is a scale-out NAS platform that provides file system protocols for clients and transparent data tiering to object storage (Cloudian HyperStore). Client applications write data to HyperFile and then HyperFile manages the underlying storage tiers, leveraging its native information lifecycle management (ILM) capabilities.

HyperFile provides capabilities including:

- Local data caching and tiering to Cloudian HyperStore object storage
- Bi-modal access to data (data tiered from HyperFile to object storage can be read through HyperFile's file protocols or directly through HyperStore's S3 interface)
- Integrated data protection via snapshots
- Active Directory / LDAP integration and user quotas
- Multi-controller configurations
- High availability (HA) configurations



• Write Once Read Many (WORM) support, together with compliance features such as auditing and so on

5. Load Balancing Cloudian HyperFile

8 Note

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

5.1. Load Balancing & HA Requirements

To allow a Cloudian HyperFile deployment to be load balanced, the HyperFile nodes must be deployed in a multicontroller configuration sharing an NFS volume.

5.2. Persistence (aka Server Affinity)

Source IP address persistence is required to successfully load balance Cloudian HyperFile. This is true for both the layer 4 DR mode and layer 7 load balancing scenarios described in this document.

5.3. Virtual Service (VIP) Requirements

To provide load balancing and HA for Cloudian HyperFile, a single VIP is used which covers all of the ports needed.

5.4. Port Requirements

The following table shows the ports that are load balanced:

Port	Protocols	Uses
111	TCP/RPC	Remote Procedure Call / portmap traffic (RPC)
1110	TCP/NFS	Cluster status service
2049	TCP/NFS	NFS daemon process (nfsd)
4045	TCP/NFS	Network lock manager process (nlockmgr)

Additional high ports, as well as the above mentioned ports using *UDP*, are used for NFS version 3 and below.

Note

As described later in this document, using * to cover *all* ports in a layer 4 setup is recommended for NFS version 3 and below.

6. Deployment Concept



VIPs = Virtual IP Addresses

8 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 Clustered Pair Configuration - Adding a Secondary Unit 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 Cloudian HyperFile, using either layer 4 DR mode or layer 7 SNAT mode is recommended. **If using NFS version 3 and below, layer 4 DR mode should be used** due to the wide range of ports that are used in these older versions of the NFS protocol.

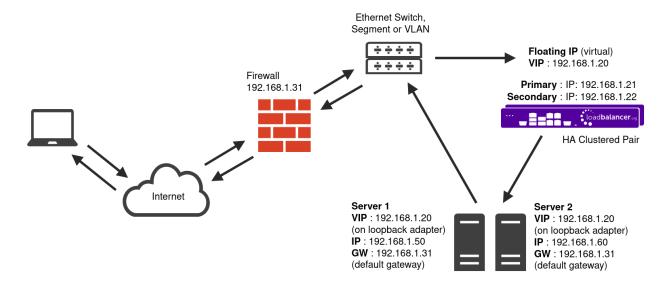
These modes are described below and are used for the configurations presented in this guide. For configuring using DR mode please refer to Appliance Configuration for Cloudian HyperFile - Using Layer 4 DR Mode, and for configuring using layer 7 SNAT mode refer to Appliance Configuration for Cloudian HyperFile - Using Layer 7 SNAT Mode.

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.

8 Note

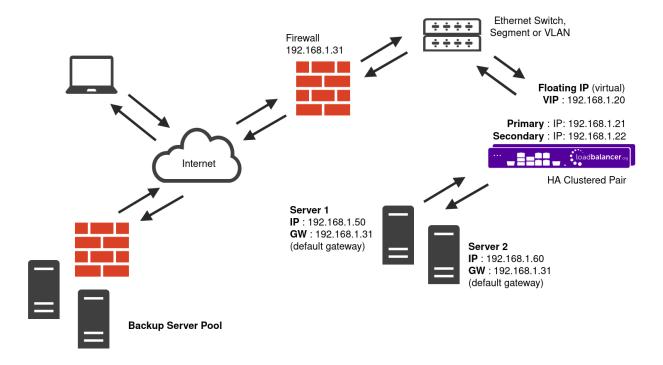
Kemp, Brocade, Barracuda & A10 Networks call this *Direct Server Return* and F5 call it *nPath*.



- 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.



- 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.

If the load balancer is deployed in AWS or Azure, layer 7 SNAT mode must be used as layer 4 direct routing is not



8. Configuring Cloudian HyperFile for Load Balancing

8.1. Configuring for Layer 4 DR Mode

(!) Important Layer 4 DR mode should be used if NFS version 3 and below is used.

For layer 4 DR mode to work, **every HyperFile node** must be configured so that its loopback adaptor owns the VIP address.

The change to the loopback adaptor should be set from the command line by writing a script to ensure that the change is persistent across reboots.

- The script should be put in the directory /etc/rc2.d and its filename must begin with a capital letter S. For example: /etc/rc2.d/Sloopbackscript
- An example script that can be used is presented below the example VIP address of 192.168.88.69 should be changed to match the VIP address being used:

```
#!/bin/sh

#
# This is to redirect ARP requests to the HyperFile VIP
#

ifconfig lo0:1 plumb
ifconfig lo0:1 192.168.88.69 netmask 255.255.255.255 up
```

8.2. Configuring for Layer 7 SNAT Mode

No changes are required on the HyperFile nodes for layer 7 SNAT mode.

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.

8 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.

8 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.

8 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.

8 Note

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

8 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/

8 Note

You'll receive a warning about the WebUl'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>

8 Note

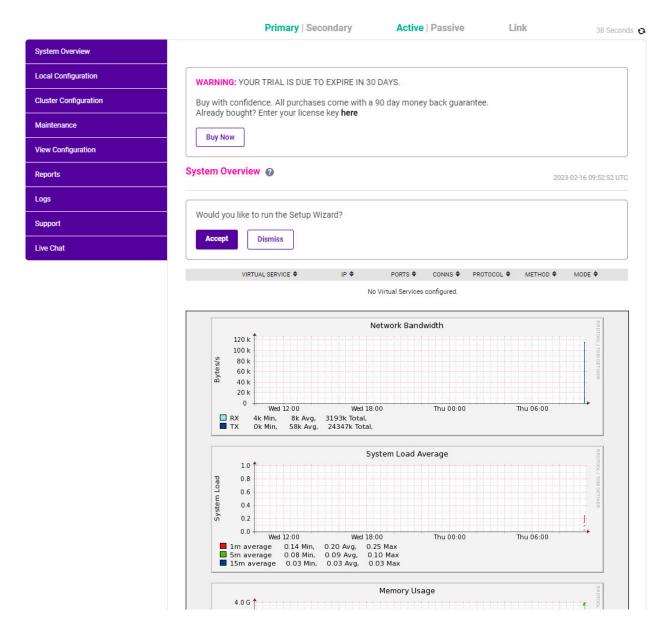
To change the password, use the WebUI menu option: *Maintenance > Passwords*.

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



IL LOADBALANCER





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



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

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.



8 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:

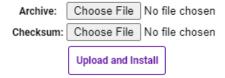
- 1. Using the WebUI, navigate to: *Maintenance > Software Update*.
- 2. Select Offline Update.
- 3. The following screen will be displayed:

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.
- Select the archive and checksum files in the upload form below.
- 4. Click Upload and Install to begin the update process.



- 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

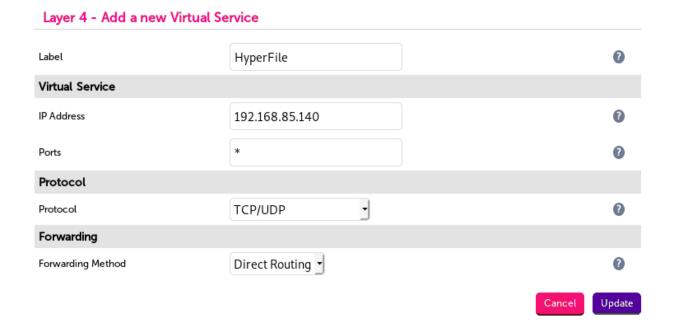
9.6. HA Clustered Pair Configuration

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 Clustered Pair Configuration - Adding a Secondary Unit.

10. Appliance Configuration for Cloudian HyperFile –Using Layer 4 DR Mode

10.1. Configuring the Virtual Service (VIP)

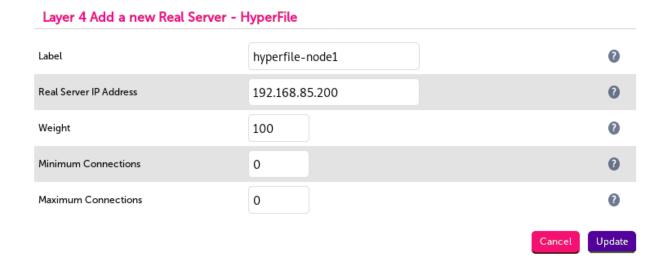
- 1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 Virtual Services* and click on **Add** a new Virtual Service.
- 2. Define the Label for the virtual service as required, e.g. HyperFile.
- 3. Set the Virtual Service IP Address field to the required IP address, e.g. 192.168.85.140.
- 4. Set the *Ports* field to * (this wildcard sets the VIP to use all ports).
- 5. Set the *Protocol* to **TCP/UDP**.
- 6. Leave the Forwarding Method set to Direct Routing.
- 7. Click **Update** to create the virtual service.



- 8. Click Modify next to the newly created VIP.
- 9. Ensure that the *Persistence Enable* checkbox is checked and that the *Timeout* is set to **300** (this should already be configured by default).
- 10. Set the Health Checks Check Type to Connect to port.
- 11. Set the Check Port to 2049.

10.2. Defining the Real Servers (RIPs)

- 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. **hyperfile-node1**.
- 3. Set the Real Server IP Address field to the required IP address, e.g. 192.168.85.200.
- 4. Click Update.
- 5. Repeat these steps to add additional HyperFile servers as required.



11. Appliance Configuration for Cloudian HyperFile –Using Layer 7 SNAT Mode

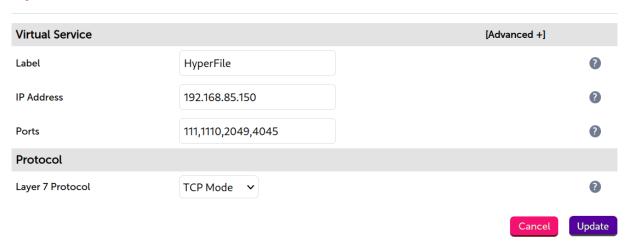


This load balancing method should not be used if NFS version 3 and below is to be used with HyperFile. Layer 4 DR mode should be used instead (see the previous section on how to set this up). This is because NFS versions 3 and below use additional high ports, as well as the standard ports but using UDP.

11.1. Configuring the Virtual Service (VIP)

- 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. **HyperFile**.
- 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 111,1110,2049,4045.
- 5. Set the Layer 7 Protocol to TCP Mode.
- 6. Click **Update** to create the virtual service.

Layer 7 - Add a new Virtual Service

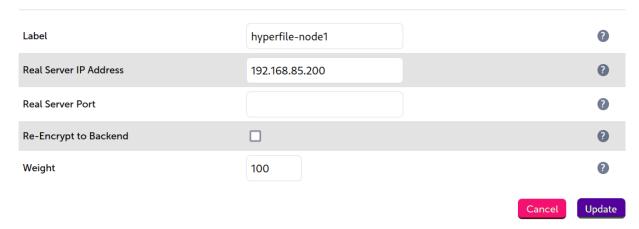


- 7. Click Modify next to the newly created VIP.
- 8 Set Persistence Mode to Source IP.
- 9. In the *Persistence* section click **Advanced** to expand the menu.
- 10. Set *Persistence Timeout* to **5** (the default units are minutes).
- 11. Set *Health Checks* to **Connect to port**.
- 12. In the *Health Checks* section click **Advanced** to expand the menu.
- 13. Set Check Port to 2049.
- 14. In the Other section click Advanced to expand the menu.
- 15. Check the **Timeout** checkbox.
- 16. Set *Client Timeout* to **5m** (the *m* is for minutes).
- 17. Set Real Server Timeout to 5m.
- 18. Click Update.

11.2. Defining the Real Servers (RIPs)

- 1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
- 2. Define the *Label* for the real server as required, e.g. **hyperfile-node1**.
- 3. Set the Real Server IP Address field to the required IP address, e.g. 192.168.85.200.
- 4. Click Update.
- 5. Repeat these steps to add additional HyperStore nodes as real servers as required.

Layer 7 Add a new Real Server - HyperFile



11.3. 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.

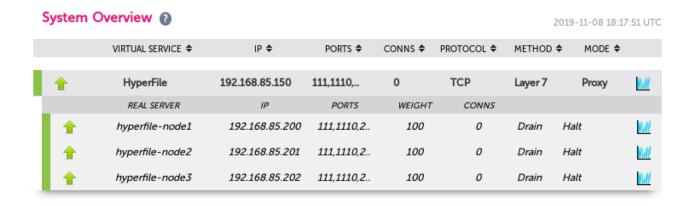
12. Testing & Verification

8 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 Cloudian HyperFile 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 three HyperFile nodes are healthy and available to accept connections:



12.2. Layer 4 DR mode specific test

If the layer 4 DR mode load balancing method has been used then an additional check can be performed to

confirm that the load balanced HyperFile deployment as a whole is functioning correctly.

After sending some test traffic to the virtual service, from the WebUI, navigate to *Reports > Layer 4 Current Connections*. Ensure that the test connections are not shown to be in the *SYN_RECV* state under the third column, 'state'. Successful connections are shown as *ESTABLISHED* like so:

Layer 4 Current Connections

Check Status

IPVS connection entries

pro expire state source virtual destination

IP 01:08 NONE 192.168.86.2:0 119.53.148.0:0 192.168.86.77:0

TCP 14:42 ESTABLISHED 192.168.86.2:669 192.168.86.69:2049 192.168.86.77:2049

If any of the connections are in the *SYN_RECV* state then it is very likely that the HyperFile nodes have not been correctly configured for layer 4 DR mode. Identify which nodes are affected, by looking at their IP address in the 'destination' column, and then refer to the section Configuring for Layer 4 DR Mode and ensure that all steps have been followed correctly.

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.

8 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.

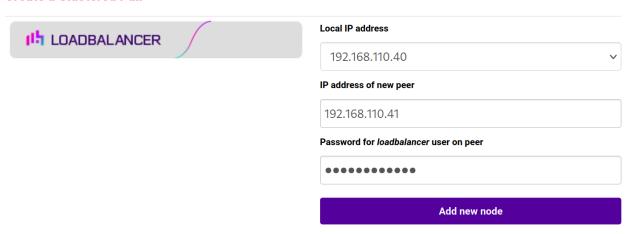
Adding a Secondary Appliance - Create an HA Clustered Pair

8 Note

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

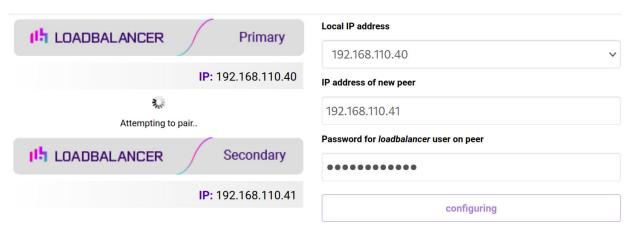
- 1. Deploy a second appliance that will be the Secondary and configure initial network settings.
- 2. Using the WebUI on the Primary appliance, navigate to: *Cluster Configuration > High-Availability Configuration*.

Create a Clustered Pair



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

Create a Clustered Pair



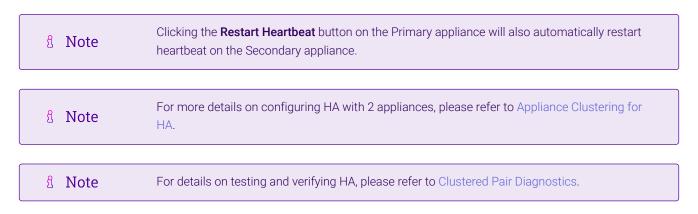
6. Once complete, the following will be displayed on the Primary appliance:



High Availability Configuration - primary



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.



16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	8 November 2019	Initial version		IG, AH
1.0.1	18 November 2019	Removed the instruction to change each node's default gateway to the VIP address in section 'Configuring Cloudian HyperFile for Load Balancing'	The step in question was not required and was removed for simplicity	АН
1.0.2	1 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	АН
1.1.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH,RJC,ZAC
1.1.1	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.1.2	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	АН
1.1.3	2 February 2023	Updated screenshots	Branding update	АН
1.1.4	7 March 2023	Removed conclusion section	Updates across all documentation	АН
1.2.0	24 March 2023	New document theme Modified diagram colours	Branding update	АН



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.

