

Load Balancing Panzura CloudFS

Version 1.2.0



Table of Contents

1. About this Guide	4
2. Loadbalancer.org Appliances Supported	4
3. Software Versions Supported	4
3.1. Loadbalancer.org Appliance	4
3.2. Panzura CloudFS	4
4. Panzura CloudFS	4
5. Load Balancing Panzura CloudFS	5
5.1. Load Balancing & HA Requirements	5
5.2. Persistence (aka Server Affinity)	5
5.3. Virtual Service (VIP) Requirements	5
5.4. Port Requirements	5
6. Deployment Concept	5
7. Load Balancer Deployment Methods	6
7.1. Layer 4 NAT Mode	6
7.2. Layer 7 SNAT Mode	9
7.3. Our Recommendation	10
8. Configuring Panzura CloudFS for Load Balancing	10
8.1. Configuring for Layer 4 NAT Mode	10
8.2. Configuring for Layer 7 SNAT Mode (recommended)	10
9. Loadbalancer.org Appliance – the Basics	10
9.1. Virtual Appliance	10
9.2. Initial Network Configuration	11
9.3. Accessing the Appliance WebUI	11
9.3.1. Main Menu Options	12
9.4. Appliance Software Update	13
9.4.1. Online Update	13
9.4.2. Offline Update	13
9.5. Ports Used by the Appliance	14
9.6. HA Clustered Pair Configuration	15
10. Appliance Configuration for Panzura CloudFS – Using Layer 4 NAT Mode	15
10.1. Configuring the SMB Virtual Service (VIP)	15
10.2. Defining the Real Servers (RIPs)	16
10.3. Configuring the NFS Virtual Service (VIP)	16
10.4. Defining the Real Servers (RIPs)	17
10.5. Finalizing the Configuration	18
10.5.1. Creating a floating IP for the Panzura CloudFS gateway address	18
11. Appliance Configuration for Panzura CloudFS – Using Layer 7 SNAT Mode	19
11.1. Configuring the SMB Virtual Service (VIP)	19
11.2. Defining the Real Servers (RIPs)	19
11.3. Configuring the NFS Virtual Service (VIP)	20
11.4. Defining the Real Servers (RIPs)	21
11.5. Finalizing the Configuration	21
12. Testing & Verification	21
12.1. Using System Overview	22
13. Technical Support	22
14. Further Documentation	22
15. Appendix	23
15.1. Configuring HA - Adding a Secondary Appliance	23

15.1.1. Non-Replicated Settings	23
15.1.2. Configuring the HA Clustered Pair	24
16. Document Revision History	26

1. About this Guide

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

- V8.9.1 and later

 **Note**

The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If you're using an older version, or the very latest, the screenshots presented here may not match your WebUI exactly.

3.2. Panzura CloudFS

- V7.1.9x and later

4. Panzura CloudFS

The Panzura Cloud File System (PCFS) is a distributed cloud file system used for storing application data that spans the globe, granting users in various geographical locations fast and consistent access to that data.

Panzura CloudFS capabilities include:

- Providing instances at each site that connect to each other and to a master data source. These instances can be deployed as either a virtual machine (VM), or as an in-cloud instance
- Frequently used files are cached at each office for fast access
- Files are kept constantly synchronized across all sites
- Users can access the same file at the same time



- The master copy is kept synchronized with the public or private cloud provider of your choice
- Byte-range global locking technology protects files from being accidentally overwritten

5. Load Balancing Panzura CloudFS

 Note

It's highly recommended that you have a working Panzura CloudFS environment first before implementing the load balancer. The Panzura instances should be configured on a Master/Subordinate configuration when being deployed behind a load balancer.

5.1. Load Balancing & HA Requirements

The function of the load balancer is to distribute inbound connections across a cluster of Panzura CloudFS nodes, to provide a highly available and scalable service. Two virtual services are used to load balance the different aspects of Panzura CloudFS.

5.2. Persistence (aka Server Affinity)

Persistence is not needed as the Panzura CloudFS Master and Subordinates synchronise configuration between themselves.

5.3. Virtual Service (VIP) Requirements

To provide load balancing for Panzura CloudFS, the following VIPs are required:

- **SMB**: for Windows print and file sharing cluster
- **NFS**: Network file system cluster

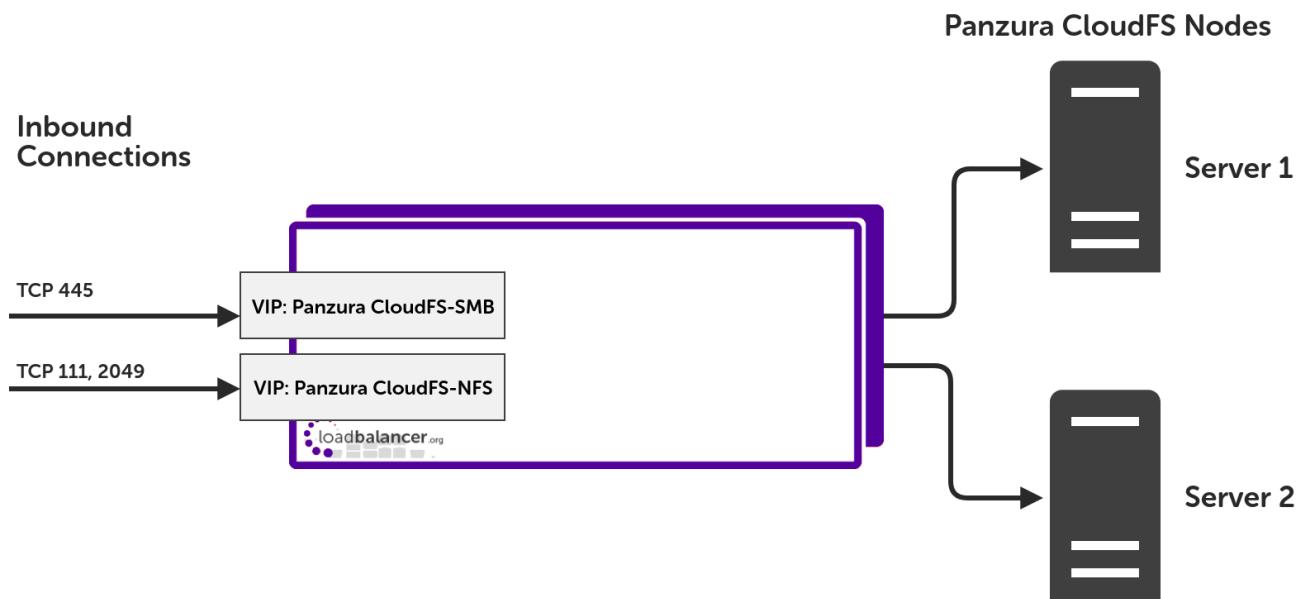
5.4. Port Requirements

The following table shows the ports that are load balanced:

Port	Protocols	Uses
111	TCP/RPC	Remote Procedure Call / portmap traffic (RPC)
445	TCP/SMB	Windows File and print sharing
2049	TCP/NFS	NFS daemon process (nfsd)

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 [Configuring HA - Adding a Secondary Appliance](#) for more details on configuring a clustered pair.

7. Load Balancer Deployment Methods

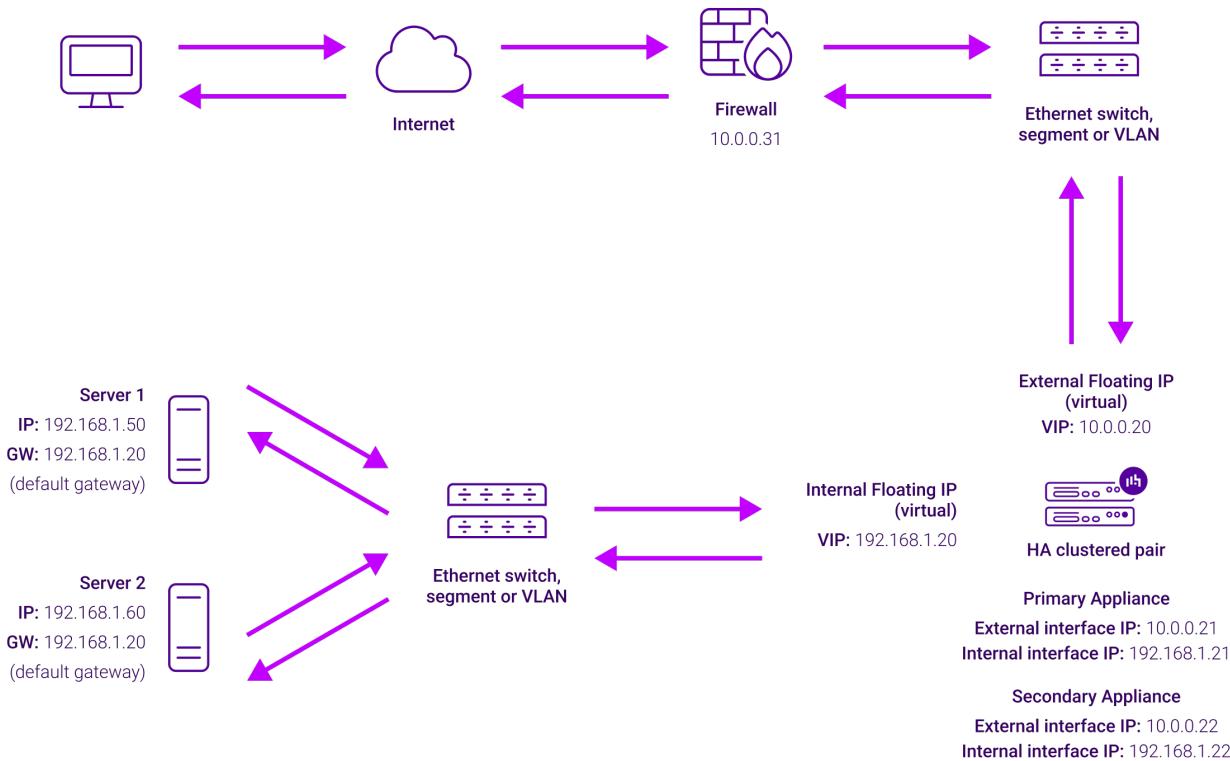
The load balancer can be deployed in 4 fundamental ways: *Layer 4 DR mode*, *Layer 4 NAT mode*, *Layer 4 SNAT mode*, and *Layer 7 SNAT mode*.

For Panzura CloudFS, using either layer 4 NAT mode or layer 7 SNAT mode is recommended. **Layer 4 DR mode is not recommended due to operating system restrictions on the Panzura CloudFS nodes.**

These modes are described below and are used for the configurations presented in this guide. For configuring using NAT mode please refer to [Appliance Configuration for Panzura CloudFS – Using Layer 4 NAT Mode](#), and for configuring using layer 7 SNAT mode refer to [Appliance Configuration for Panzura CloudFS – Using Layer 7 SNAT Mode](#).

7.1. Layer 4 NAT Mode

Layer 4 NAT mode is a high performance solution, although not as fast as layer 4 DR mode. This is because real server responses must flow back to the client via the load balancer rather than directly as with DR mode. The image below shows an example network diagram for this mode.



- The load balancer translates all requests from the Virtual Service to the Real Servers.
- NAT mode can be deployed in the following ways:
 - Two-arm (using 2 Interfaces)** (as shown above) - Here, 2 subnets are used. The VIP is located in one subnet and the load balanced Real Servers are located in the other. The load balancer requires 2 interfaces, one in each subnet.

Note

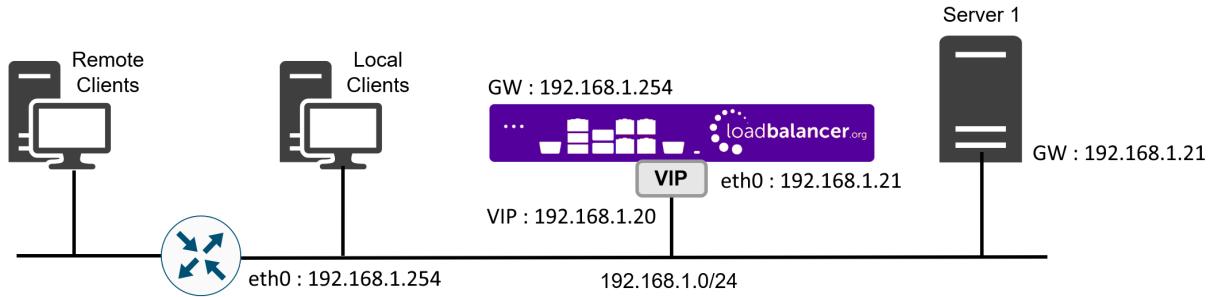
This can be achieved by using two network adapters, or by creating VLANs on a single adapter.

- Normally **eth0** is used for the internal network and **eth1** is used for the external network, although this is not mandatory since any interface can be used for any purpose.
- If the Real Servers require Internet access, **Auto-NAT** should be enabled using the WebUI menu option: *Cluster Configuration > Layer 4 - Advanced Configuration*, the external interface should be selected.
- The default gateway on the Real Servers must be set to be an IP address on the load balancer.

Note

For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can "float" (move) between Primary and Secondary appliances.

- Clients can be located in the same subnet as the VIP or any remote subnet provided they can route to the VIP.
- One-arm (using 1 Interface)** - Here, the VIP is brought up in the same subnet as the Real Servers.



- To support remote clients, the default gateway on the Real Servers must be an IP address on the load balancer and routing on the load balancer must be configured so that return traffic is routed back via the router.

Note

For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can "float" (move) between Primary and Secondary appliances.

- To support local clients, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer. For more information please refer to [One-Arm \(Single Subnet\) NAT Mode](#).
- If you want Real Servers to be accessible on their own IP address for non-load balanced services, e.g. RDP, you will need to setup individual SNAT and DNAT firewall script rules for each Real Server or add additional VIPs for this.
- Port translation is possible with Layer 4 NAT mode, e.g. VIP:80 → RIP:8080 is supported.
- NAT mode is transparent, i.e. the Real Server will see the source IP address of the client.

NAT Mode Packet re-Writing

In NAT mode, the inbound destination IP address is changed by the load balancer from the Virtual Service IP address (VIP) to the Real Server. For outbound replies the load balancer changes the source IP address of the Real Server to be the Virtual Services IP address.

The following table shows an example NAT mode setup:

Protocol	VIP	Port	RIP	Port
TCP	10.0.0.20	80	192.168.1.50	80

In this simple example all traffic destined for IP address 10.0.0.20 on port 80 is load-balanced to the real IP address 192.168.1.50 on port 80.

Packet rewriting works as follows:

- 1) The incoming packet for the web server has source and destination addresses as:



Source	x.x.x.x:34567	Destination	10.0.0.20:80
---------------	---------------	--------------------	--------------

2) The packet is rewritten and forwarded to the backend server as:

Source	x.x.x.x:34567	Destination	192.168.1.50:80
---------------	---------------	--------------------	-----------------

3) Replies return to the load balancer as:

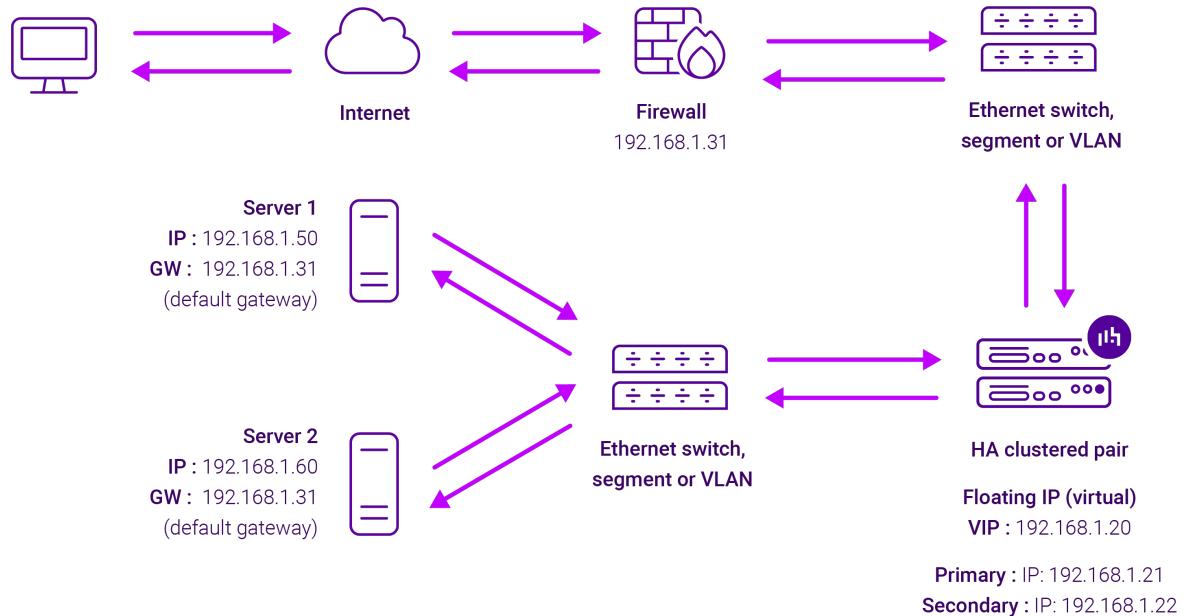
Source	192.168.1.50:80	Destination	x.x.x.x:34567
---------------	-----------------	--------------------	---------------

4) The packet is written back to the VIP address and returned to the client as:

Source	10.0.0.20:80	Destination	x.x.x.x:34567
---------------	--------------	--------------------	---------------

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. 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.

7.3. Our Recommendation

Where possible, we recommend that Layer 7 SNAT mode is used. This mode offers great performance with minimal to no changes required on the real servers and can be deployed in one-arm or two-arm mode. HAProxy is a high performance solution, but since it operates as a full proxy, it cannot perform as fast as the layer 4 solutions. Layer 7 SNAT mode is non-transparent by default, i.e. the Real Servers will see the source IP address of the load balancer.

8. Configuring Panzura CloudFS for Load Balancing

8.1. Configuring for Layer 4 NAT Mode

For layer 4 NAT mode to work, **every Panzura CloudFS node** must be configured so that its gateway points to a floating IP on the load balancer(s).

8.2. Configuring for Layer 7 SNAT Mode (recommended)

No changes are required on the Panzura CloudFS 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](#).

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.

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 servers and other network and administrative settings.

(1) 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](#).

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:



Primary | Secondary Active | Passive Link 8 Seconds

System Overview

Local Configuration

Cluster Configuration

Maintenance

View Configuration

Reports

Logs

Support

Live Chat

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](#)

Buy Now

System Overview 2025-05-08 12:37:21 UTC

Would you like to run the Setup Wizard?

Accept **Dismiss**

VIRTUAL SERVICE **IP** **PORTS** **CONN** **PROTOCOL** **METHOD** **MODE**

No Virtual Services configured.

Network Bandwidth

System Load Average

Memory Usage

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.



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

9.3.1. Main Menu Options

System Overview - Displays a graphical summary of all VIPs, RIPv4 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 & RIPv4

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

9.4. 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.

9.4.1. 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 8.13.2 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.

9.4.2. 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

Upload and Install

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 / GSLB
TCP & UDP	123	NTP
TCP & UDP	161 *	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000 *	Gateway service (Centralized/Portal Management)
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

Note

The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to [Service Socket](#)



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 Panzura CloudFS – Using Layer 4 NAT Mode

10.1. Configuring the SMB 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. **Panzura-SMB**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.86.140**.
4. Set the *Ports* field to **445**.
5. Set the *Protocol* to **TCP**.
6. Change the *Forwarding Method* to **NAT**.
7. Click **Update** to create the virtual service.

Layer 4 - Add a new Virtual Service

Label	Panzura-SMB	?
Virtual Service		
IP Address	192.168.86.140	?
Ports	445	?
Protocol		
Protocol	TCP	?
Forwarding		
Forwarding Method	NAT	?
		Cancel Update

8. Click **Modify** next to the newly created VIP.
9. Ensure that the *Persistence Enable* checkbox is unchecked.
10. Set the *Health Checks Check Type* to **Connect to port**.

11. Set the *Check Port* to **445**.

12. Click **Update**.

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. **Panzura1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **172.24.11.138**.
4. Click **Update**.
5. Repeat these steps to add additional Panzura servers as required.

Layer 4 Add a new Real Server - Panzura-SMB

Label	Panzura1	?
Real Server IP Address	172.24.11.138	?
Real Server Port		?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?
		Cancel Update

10.3. Configuring the NFS 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. **Panzura-NFS**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.86.140**.
4. Set the *Ports* field to **111, 2049**.
5. Set the *Protocol* to **TCP**.
6. Change the *Forwarding Method* to **NAT**.
7. Click **Update** to create the virtual service.



Layer 4 - Add a new Virtual Service

Label	Panzura-NFS	?
Virtual Service		
IP Address	192.168.86.140	?
Ports	111,2049	?
Protocol		
Protocol	TCP	?
Forwarding		
Forwarding Method	NAT	?
		Cancel Update

8. Click **Modify** next to the newly created VIP.
9. Ensure that the *Persistence Enable* checkbox is unchecked.
10. Set the *Health Checks Check Type* to **Connect to port**.
11. Set the *Check Port* to **2049**.
12. Click **Update**.

10.4. 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. **Panzura1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **172.24.11.138**.
4. Click **Update**.
5. Repeat these steps to add additional Panzura servers as required.



Layer 4 Add a new Real Server - Panzura-NFS

Label	Panzura1	?
Real Server IP Address	172.24.11.138	?
Real Server Port		?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?

Cancel **Update**

10.5. Finalizing the Configuration

For layer 4 NAT mode to work **every Panzura CloudFS** node must be configured so that its gateway points to a floating IP of the load balancer.

10.5.1. Creating a floating IP for the Panzura CloudFS gateway address

1. Using the web user interface, navigate to *Cluster Configuration > Floating IPs*.

FLOATING IPs		
192.168.111.40	Delete	
192.168.111.42	Delete	
New Floating IP		

Add Floating IP

2. Specify the new floating IP.
3. Click **Add Floating IP**.

Note

When using a clustered pair, ensure that the Secondary also has a static IP address assigned that's in the same subnet as the floating IP being added. Failure to do so will result in heartbeat issues during a failover.

Note

Floating IPs are not deleted automatically when Virtual Services are removed or the IP address is changed, this must be done manually.



11. Appliance Configuration for Panzura CloudFS – Using Layer 7 SNAT Mode

11.1. Configuring the SMB 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. **Panzura-SMB**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.86.140**.
4. Set the *Ports* field to **445**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	Panzura-SMB	?
IP Address	192.168.86.140	?
Ports	445	?
Protocol		
Layer 7 Protocol	TCP Mode	?
		Cancel Update

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Set *Health Checks* to **Connect to port**.
10. In the *Other* section click **Advanced** to expand the menu.
11. Check the **Timeout** checkbox.
12. Set *Client Timeout* to **5m** (the *m* is for minutes).
13. Set *Real Server Timeout* to **5m**.
14. 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. **Panzura1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **172.24.11.138**.
4. Click **Update**.
5. Repeat these steps to add additional Panzura CloudFS nodes as real servers as required.

Layer 7 Add a new Real Server

Label	Panzura1	?
Real Server IP Address	172.24.11.138	?
Real Server Port		?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	100	?
<input type="button" value="Cancel"/> <input type="button" value="Update"/>		

11.3. Configuring the NFS 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. **Panzura-NFS**.
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 **111, 2049**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	Panzura-NFS	?
IP Address	192.168.86.140	?
Ports	111,2049	?
Protocol		
Layer 7 Protocol	TCP Mode	?
<input type="button" value="Cancel"/> <input type="button" value="Update"/>		

7. Click **Modify** next to the newly created VIP.



8. Set **Persistence Mode** to **None**.
9. Set **Health Checks** to **Connect to port**.
10. In the **Other** section click **Advanced** to expand the menu.
11. Check the **Timeout** checkbox.
12. Set **Client Timeout** to **5m** (the *m* is for minutes).
13. Set **Real Server Timeout** to **5m**.
14. Click **Update**.

11.4. 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. **Panzura1**.
3. Set the **Real Server IP Address** field to the required IP address, e.g. **172.24.11.138**.
4. Click **Update**.
5. Repeat these steps to add additional Panzura CloudFS nodes as real servers as required.

Layer 7 Add a new Real Server

Label	Panzura1	?
Real Server IP Address	172.24.11.138	?
Real Server Port		?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	100	?
<input type="button" value="Cancel"/> <input type="button" value="Update"/>		

11.5. Finalizing the Configuration

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

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

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 Panzura CloudFS 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:

System Overview 									2019-12-11 12:44:31 UTC
VIRTUAL SERVICE 		IP 	PORTS 		CONNNS 	PROTOCOL 	METHOD 	MODE 	
	Panzura-SMB	192.168.86.140	445	0	TCP	Layer 7	Proxy		
	REAL SERVER	IP	PORTS	WEIGHT	CONNNS				
	Panzura1	172.24.11.138	445	100	0	Drain	Halt		
	Panzura2	172.24.11.139	445	100	0	Drain	Halt		
	Panzura-NFS	192.168.86.140	111,2049	0	TCP	Layer 7	Proxy		
	REAL SERVER	IP	PORTS	WEIGHT	CONNNS				
	Panzura1	172.24.11.138	111,2049	100	0	Drain	Halt		
	Panzura2	172.24.11.139	111,2049	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 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.

15.1.1. 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



① Important

Make sure that where any of the above have been configured on the Primary appliance, they're also configured on the Secondary.

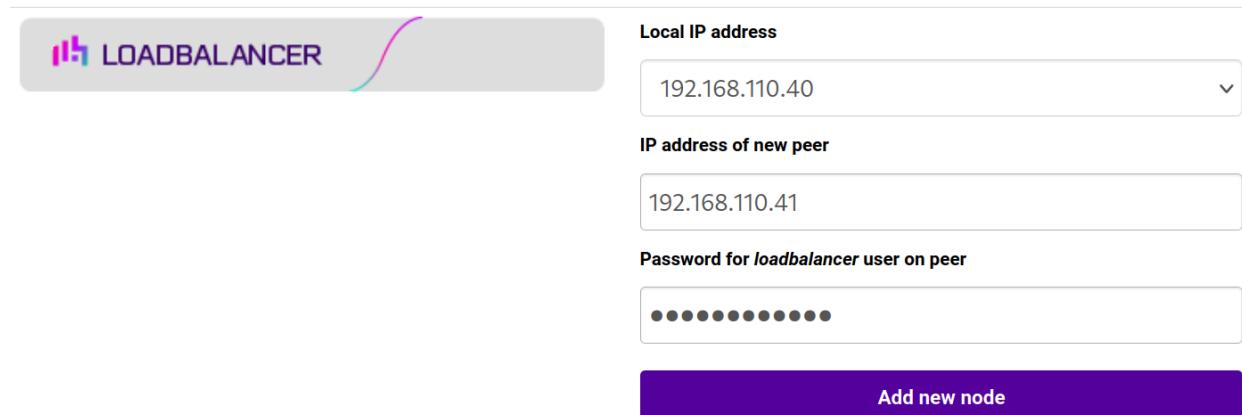
15.1.2. Configuring the HA Clustered Pair

ℹ 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

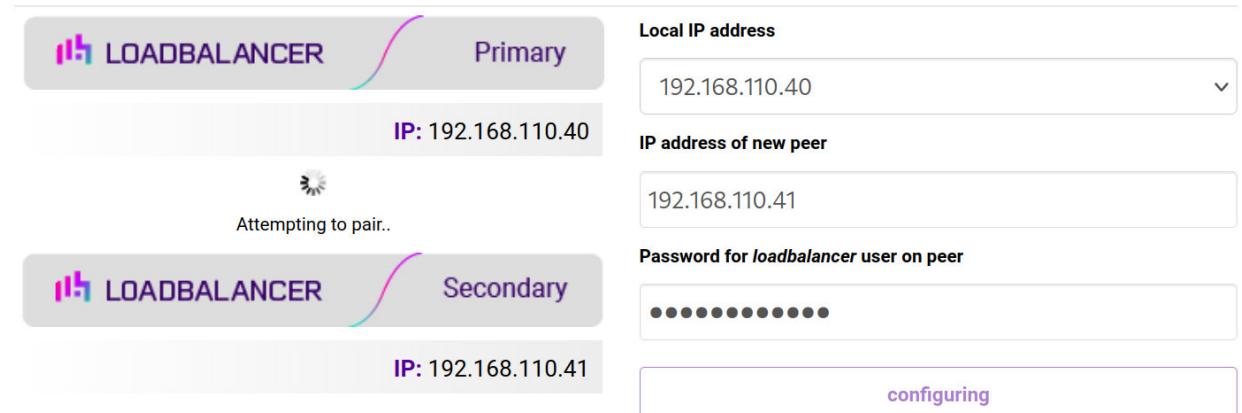
IP address of new peer
192.168.110.41

Password for *loadbalancer* user on peer
••••••••••••

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



Primary

IP: 192.168.110.40

Attempting to pair..

Secondary

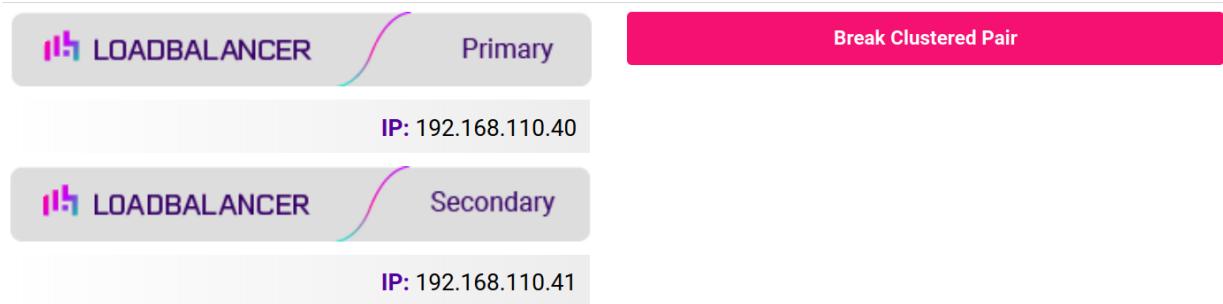
IP: 192.168.110.41

configuring

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



High Availability Configuration - primary



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

Note

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

Note

For more details on configuring HA with 2 appliances, please refer to [Appliance Clustering for HA](#).

Note

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



16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	8 November 2019	Initial version		IG
1.0.1	11 December 2019	Changed some instructions based on feedback from Panzura	Required updates	IG
1.0.2	2 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.1.0	1 November 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	AH
1.1.3	2 February 2023	Updated screenshots	Branding update	AH
1.1.4	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.2.0	24 March 2023	New document theme Modified diagram colours	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://www.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.

