LOADBALANCER

DEPLOYMENT BRIEF

Load Balancing Cerberus FTP Server

Version 1.0.0

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. Cerberus FTP Server . . . 3
4. Cerberus FTP Server . . . 3
5. Load Balancing Cerberus FTP Server . . . 3
    5.1. Load Balancing & HA Requirements . . . 4
    5.2. Virtual Service (VIP) Requirements . . . 4
    5.3. SSL/TLS Termination . . . 5
6. Deployment Concept . . . 5
7. Load Balancer Deployment Methods . . . 5
    7.1. Layer 7 SNAT Mode . . . 5
    7.2. A Note on Transparency . . . 6
8. Loadbalancer.org Appliance – the Basics . . . 7
    8.1. Virtual Appliance . . . 7
    8.2. Initial Network Configuration . . . 7
    8.3. Accessing the Appliance WebUI . . . 7
        8.3.1. Main Menu Options . . . 9
    8.4. Appliance Software Update . . . 10
        8.4.1. Online Update . . . 10
        8.4.2. Offline Update . . . 10
    8.5. Ports Used by the Appliance . . . 11
    8.6. HA Clustered Pair Configuration . . . 12
9. Configuring Cerberus FTP Server for Load Balancing . . . 12
    9.1. Layer 7 SNAT Mode . . . 12
10. Appliance Configuration for Cerberus FTP Server . . . 13
    10.1. VIP 1 - Cerberus-FTP-FTPS . . . 13
        10.1.1. Virtual Service (VIP) Configuration . . . 13
        10.1.2. Define the Associated Real Servers (RIPs) . . . 14
    10.2. VIP 2 - Cerberus-SFTP . . . 14
        10.2.1. Virtual Service (VIP) Configuration . . . 14
        10.2.2. Define the Associated Real Servers (RIPs) . . . 15
    10.3. VIP 3 - Cerberus-HTTPS . . . 16
        10.3.1. Virtual Service (VIP) Configuration . . . 16
        10.3.2. Define the Associated Real Servers (RIPs) . . . 17
    10.4. Finalizing the Configuration . . . 17
11. Testing & Verification . . . 18
    11.1. Using System Overview . . . 18
    11.2. Accessing Cerberus FTP Server via the Load Balancer . . . 18
12. Additional Considerations . . . 19
13. Technical Support . . . 19
14. Further Documentation . . . 20
15. Appendix . . . 21
    15.1. Configuring HA - Adding a Secondary Appliance . . . 21
        15.1.1. Non-Replicated Settings . . . 21
        15.1.2. Configuring the HA Clustered Pair . . . 22
16. Document Revision History . . . 24

1. About this Brief

This brief outlines the steps required to configure a load balanced Cerberus FTP Server environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Cerberus FTP Server 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 Cerberus FTP Server. 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

- V8.13.2 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. Cerberus FTP Server

- Any currently supported release

> **Note** The Professional or Enterprise edition is required if the built-in Sync Manager is to be used to replicate the configuration between nodes.

4. Cerberus FTP Server

Cerberus FTP Server (also referred to as Cerberus MFT) from Redwood Software is a secure managed file transfer (MFT) server for Windows Server. It supports FTP, FTPS, SFTP, SCP and HTTP/S, and provides a browser based HTTPS web client for end users, a web based admin console, Active Directory/LDAP integration, MFT event automation, and auditing and reporting aimed at regulatory compliance.

5. Load Balancing Cerberus FTP Server

> **Note** It's highly recommended that you have a working Cerberus FTP Server environment first before

implementing the load balancer.

5.1. Load Balancing & HA Requirements

Cerberus FTP Server does not natively support clustering. High availability is instead achieved by deploying two or more identically configured Cerberus nodes behind a load balancer. For more information, please refer to the vendor's guide to failover, high availability and load balancing.

The following points should be considered when designing the deployment:

- **Keeping the nodes identical** - user accounts and settings are kept in step across the nodes using the built-in Sync Manager (Professional and Enterprise editions), which replicates the configuration from a primary server to one or more secondary servers, and/or by pointing all nodes at the same Active Directory/LDAP directory.
- **Shared storage** - Sync Manager replicates configuration only, it does not replicate file data. All nodes must therefore use the same shared storage (e.g. an SMB/NAS share) so that users see the same files whichever node serves their connection.
- **Persistence** - each node maintains its own local session database, so sticky (persistent) load balancing is required for the HTTP/S web client. ***Source IP*** persistence is used on all VIPs in this design, which also keeps the FTP data channel on the same node as the FTP control channel.
- **Licensing** - each Cerberus node requires its own license.

The FTP, FTPS, SFTP and HTTPS services are load balanced. The web based admin console (default HTTPS port 8443) is ***not*** load balanced - see Additional Considerations.

5.2. Virtual Service (VIP) Requirements

To provide load balancing and HA for Cerberus FTP Server, the following VIPs are required:

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	Cerberus-FTP-FTPS	L7 SNAT (TCP)	21, 990, 50000-50100	Source IP	Connect to Port (21)
VIP 2	Cerberus-SFTP	L7 SNAT (TCP)	22	Source IP	Connect to Port (22)
VIP 3	Cerberus-HTTPS	L7 SNAT (TCP)	443	Source IP	Connect to Port (443)

Note

Port **990** is only required if implicit FTPS is enabled on the Cerberus servers - explicit FTPS runs over port 21. The range **50000-50100** is the passive FTP data range configured on the Cerberus servers and must match the range set there. The FTP/FTPS control ports and the passive data range are deliberately kept in a single Virtual Service so that, combined with Source IP persistence, data connections always reach the same node as the associated control connection.

Note

If the plaintext HTTP web listener is enabled on the Cerberus servers, add port **80** to the *Ports*

field of VIP 3.

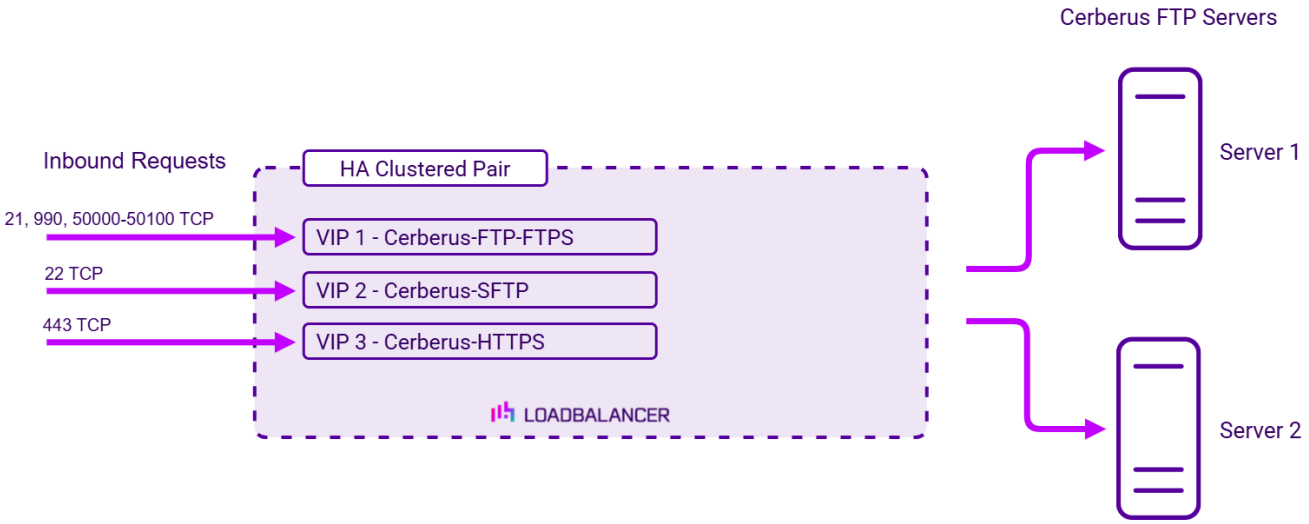
5.3. SSL/TLS Termination

No SSL termination is configured on the load balancer. FTPS and HTTPS are passed through at TCP level and TLS is terminated on the Cerberus servers themselves.

The same server certificate must be installed on every node so that clients are presented with a consistent certificate whichever node answers the connection.

6. Deployment Concept

Once the load balancer is deployed, clients connect to the Virtual Services (VIPs) rather than connecting directly to one of the Cerberus servers. These connections are then load balanced across the Cerberus servers to distribute the load according to the load balancing algorithm selected.



> **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

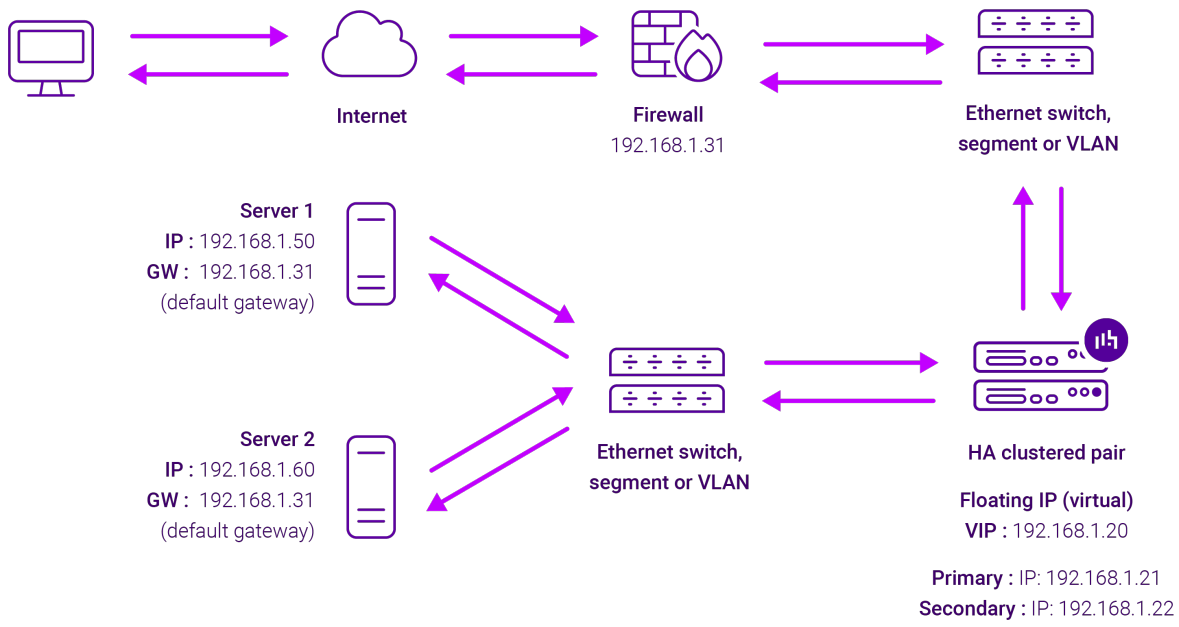
The load balancer can be deployed in 4 fundamental ways: ***Layer 4 DR mode***, ***Layer 4 NAT mode***, ***Layer 4 SNAT mode***, and ***Layer 7 SNAT mode***.

For Cerberus FTP Server, layer 7 SNAT mode is recommended. This mode is described below and is used for the configuration presented in this guide.

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.

7.2. A Note on Transparency

Layer 7 SNAT mode is not transparent by default, which means that the Cerberus servers see connections as originating from the load balancer rather than from the real client. Per-client IP allow/deny rules and per-IP logging on Cerberus will therefore not see the real client addresses.

If client IP transparency is required, either:

- Enable TProxy on the layer 7 Virtual Services. Certain topology requirements must be met - for more details please refer to Transparency at Layer 7 in the Administration Manual, or
- Load balance the FTP/FTPS service at layer 4 instead. Layer 4 DR mode and layer 4 NAT mode support both active and passive FTP, layer 4 SNAT mode supports passive FTP only. For more details please refer to Example 4 - Load Balancing FTP in the Administration Manual.

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.

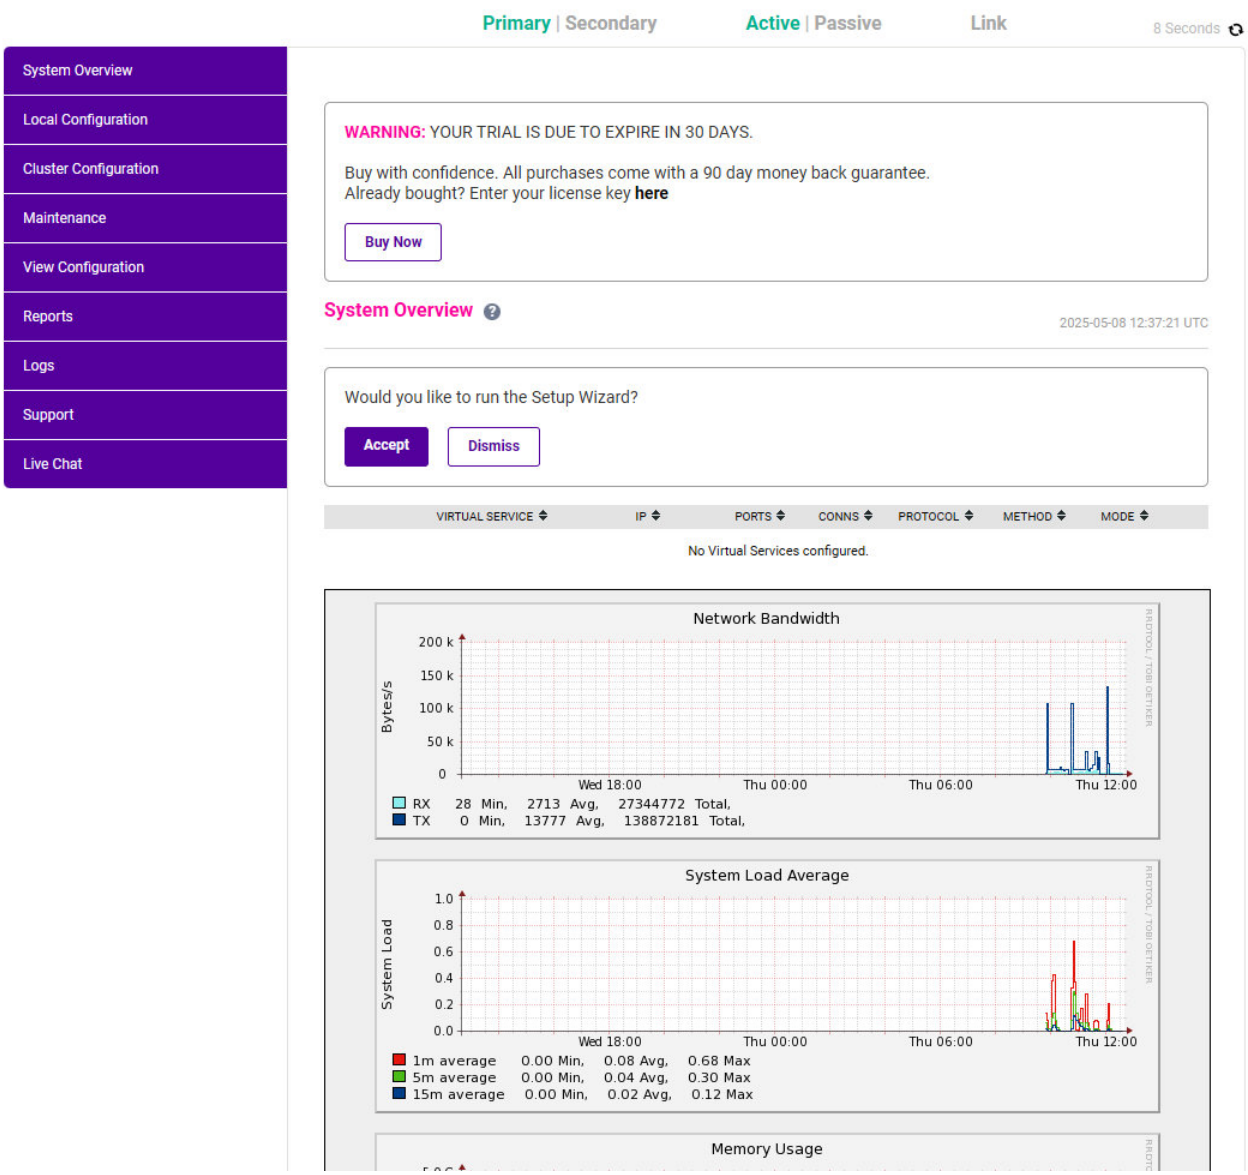
> **Note** If you need to change the port, IP address or protocol that the WebUI listens on, please refer to Service Socket Addresses.

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

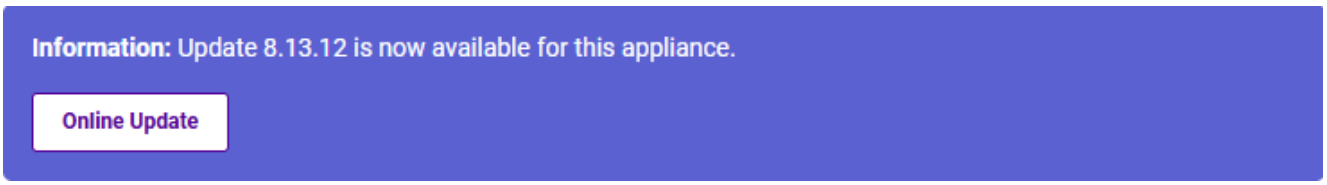
We recommend that the appliance is kept up to date to ensure that you benefit from the latest bug fixes, security updates and feature improvements. Both online and offline update are supported.

> **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

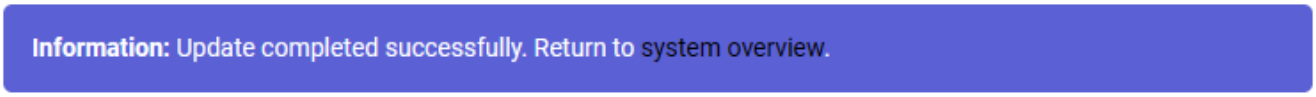
The appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. This is the default behavior and can be disabled if preferred. If an update is found, a notification similar to the example below will be displayed at the top of the WebUI:



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:



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

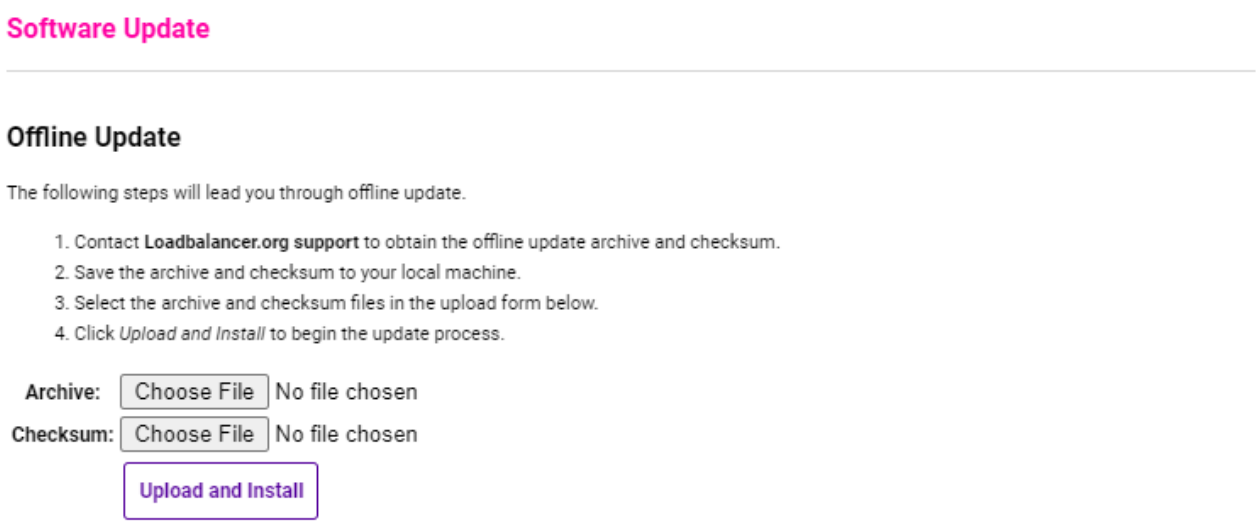
8.4.2. Offline Update

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:



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
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. Configuring Cerberus FTP Server for Load Balancing

Configure one node fully - this becomes the primary - and then bring the remaining nodes into line with it using the steps below.

1. **Replication** - use Sync Manager (Professional and Enterprise editions) to replicate the primary's configuration to the secondary nodes, and/or configure all nodes to authenticate against the same Active Directory/LDAP directory. All nodes must run the same version of Cerberus FTP Server, so plan to upgrade them together. For details, please refer to the vendor's article on configuring server replication.
2. **Passive FTP port range** - set the same passive port range on every node, e.g. **50000-50100**. The passive mode options are located in the ***Advanced*** section of the Server Manager. This range must match the range configured on VIP 1.
3. **Passive mode external IP** - on every node, set the passive/WAN IP to the VIP 1 address and disable WAN IP autodetection. These options are located in the ***General*** section of the Server Manager. This ensures that the PASV reply hands the client the VIP address rather than the node's own address. For details, please refer to the vendor's article on configuring external access for the FTP protocol.
4. **SFTP host key** - every node must present the same SSH host key. Export the host key from the primary node and import it on each secondary node, otherwise clients will raise host-key-changed warnings when a subsequent connection lands on a different node.
5. **TLS certificate** - install the same server certificate on all nodes. This certificate is used for both FTPS and the HTTPS web client.
6. **Shared storage** - point the user home directories and virtual directories on every node at the same UNC path, and run the Cerberus service under an account that has rights to that share.
7. **Windows Firewall** - on each node, allow inbound TCP 21, 990 (if implicit FTPS is used), 22 and 443 (and 80 if the plaintext HTTP listener is enabled) plus the passive data range 50000-50100. If Sync Manager is used, each secondary node must also allow inbound TCP 8443 (the HTTPS Admin port, by default) from the primary node, as this is the port Sync Manager connects to. For a full list of the ports used, please refer to the vendor's article on ports used by Cerberus FTP Server.

9.1. Layer 7 SNAT Mode

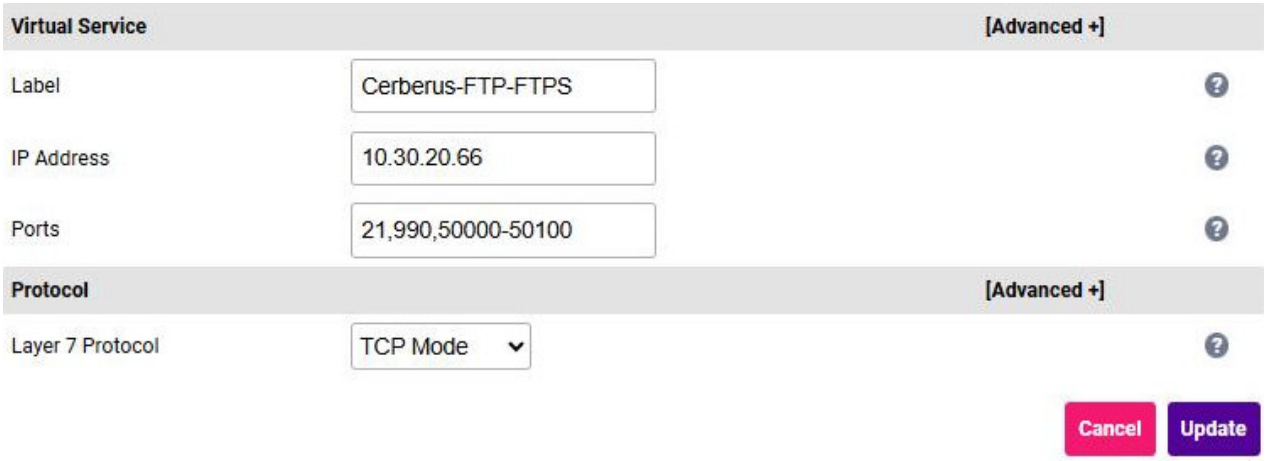
Layer 7 SNAT mode VIPs do not require any mode specific configuration changes to the load balanced Real Servers (the Cerberus servers).

10. Appliance Configuration for Cerberus FTP Server

10.1. VIP 1 - Cerberus-FTP-FTPS

10.1.1. Virtual Service (VIP) Configuration

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



3. Enter an appropriate *Label* for the VIP, e.g. **Cerberus-FTP-FTPS**.
4. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.30.20.66**.
5. Set the *Virtual Service Ports* field to **21,990,50000-50100**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Click **Modify** next to the newly created VIP.
9. Scroll to the *Connection Distribution Method* section.
   - Set the *Balance Mode* to **Weighted Round Robin**.
10. Scroll to the *Persistence* section.
    - Ensure that the *Persistence Mode* is set to **Source IP**.
11. Scroll to the *Health Checks* section and click **[Advanced]**.
    - Set the *Health Checks* to **Connect to port**.
    - Set the *Check Port* to **21**.
12. Scroll to the *Other* section and click **[Advanced]**.
    - Enable (check) the *Timeout* checkbox.
    - Set the *Client Timeout* to **5m**.
    - Set the *Real Server Timeout* to **5m**.

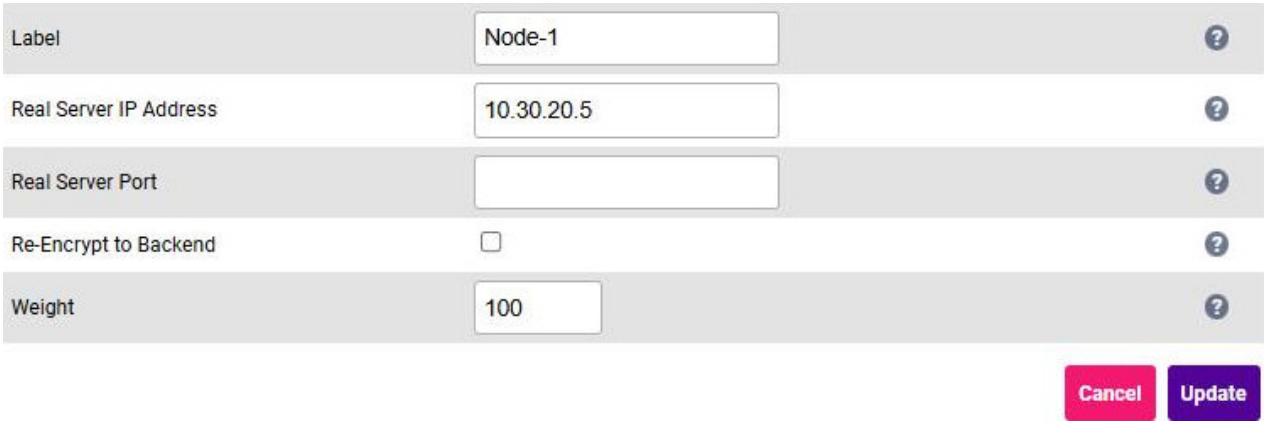
> **Note** 5 minutes is the minimum recommended value. Increase it further if large file transfers are expected to run for longer than this.

13. Leave all other settings at their default value.
14. Click **Update**.

> **Note** Port 990 only needs to be included in the *Virtual Service Ports* field if implicit FTPS is enabled on the Cerberus servers.

10.1.2. Define the Associated Real Servers (RIPs)

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



3. Define the *Label* for the Real Server as required, e.g. **Node-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **10.30.20.5**.
5. Leave the *Real Server Port* field ***blank***.

> **Note** The port field must be left blank because this is a multi-port Virtual Service - each connection is then forwarded to the Real Server on the same port it arrived on.

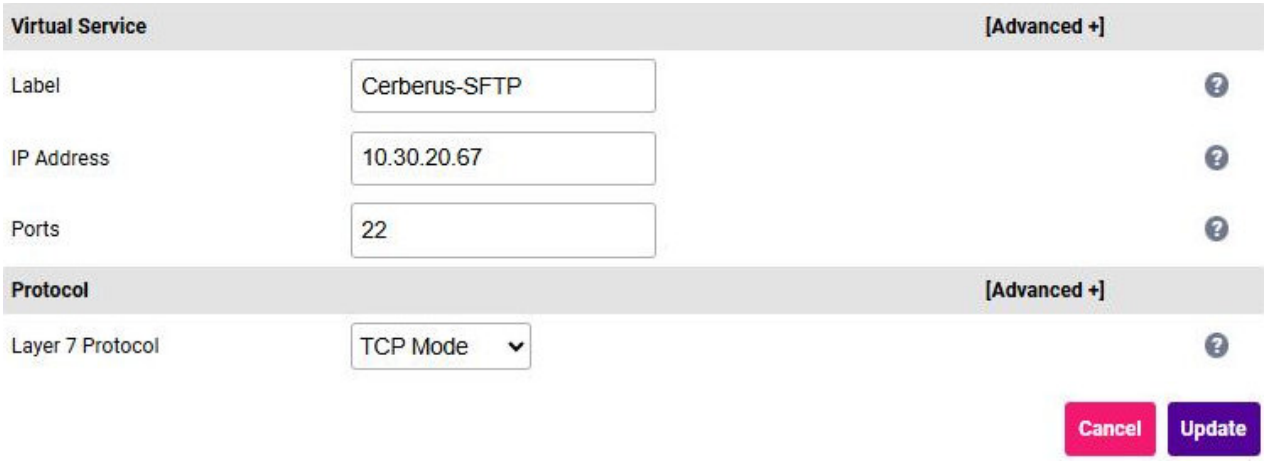
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining Cerberus server(s), e.g. **Node-2** at **10.30.20.6**.

10.2. VIP 2 - Cerberus-SFTP

10.2.1. Virtual Service (VIP) Configuration

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

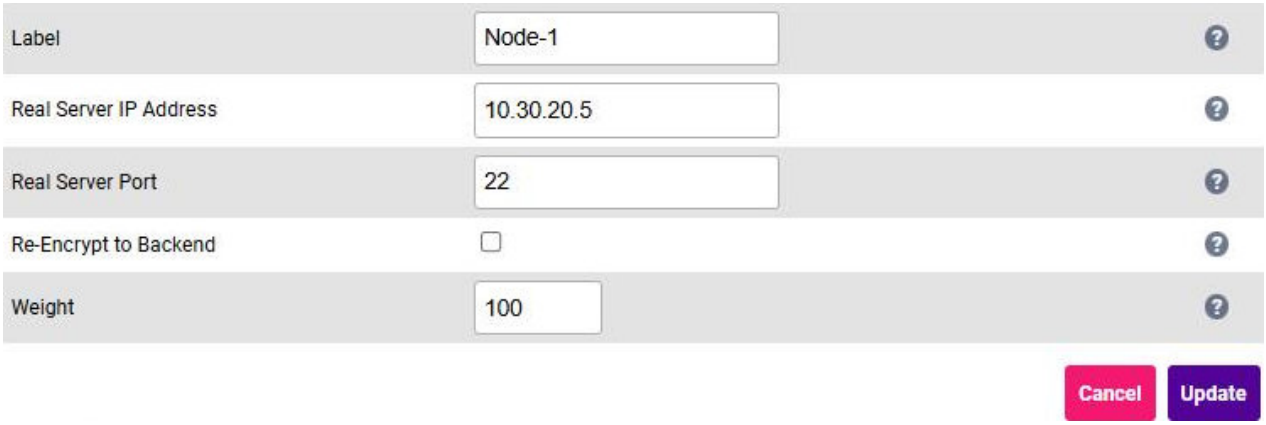
2. Enter the following details:



3. Enter an appropriate *Label* for the VIP, e.g. **Cerberus-SFTP**.
4. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.30.20.67**.
5. Set the *Virtual Service Ports* field to **22**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Click **Modify** next to the newly created VIP.
9. Scroll to the *Connection Distribution Method* section.
    - Set the *Balance Mode* to **Weighted Round Robin**.
10. Scroll to the *Persistence* section.
    - Ensure that the *Persistence Mode* is set to **Source IP**.
11. Scroll to the *Health Checks* section.
    - Set the *Health Checks* to **Connect to port**.
12. Scroll to the *Other* section and click **[Advanced]**.
    - Enable (check) the *Timeout* checkbox.
    - Set the *Client Timeout* to **5m**.
    - Set the *Real Server Timeout* to **5m**.
13. Leave all other settings at their default value.
14. Click **Update**.

10.2.2. Define the Associated Real Servers (RIPs)

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

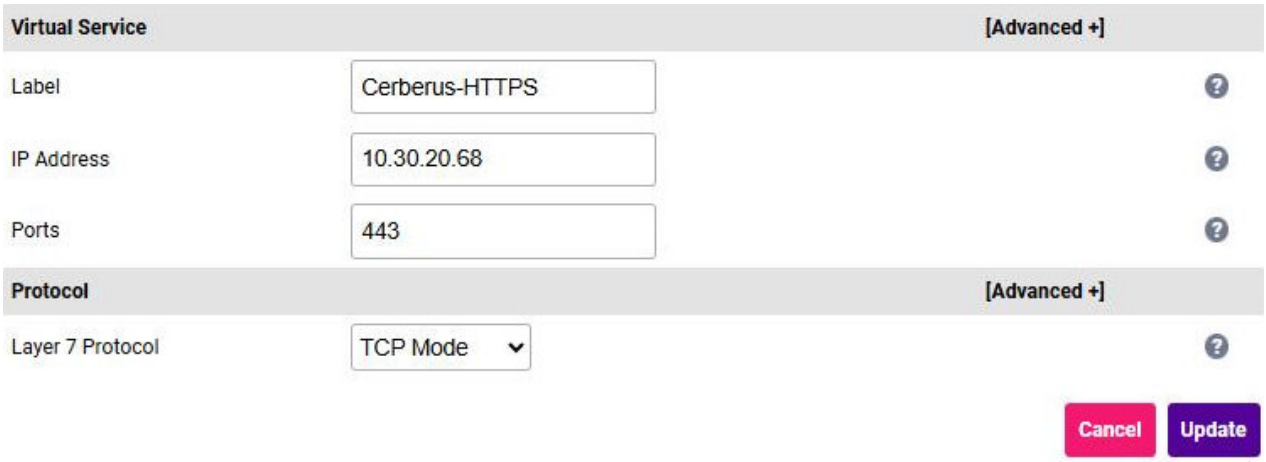


3. Define the *Label* for the Real Server as required, e.g. **Node-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **10.30.20.5**.
5. Set the *Real Server Port* field to **22**.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining Cerberus server(s), e.g. **Node-2** at **10.30.20.6**.

10.3. VIP 3 - Cerberus-HTTPS

10.3.1. Virtual Service (VIP) Configuration

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



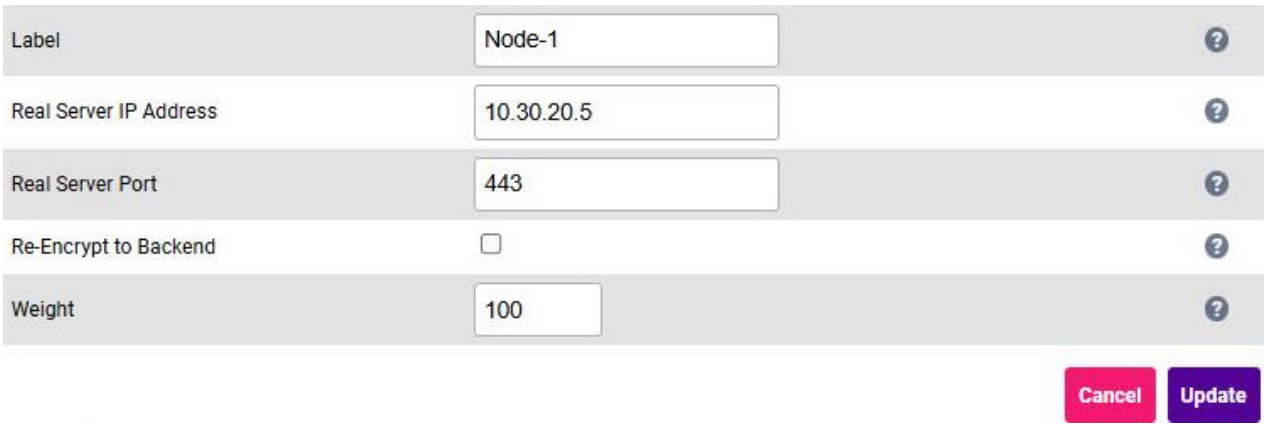
3. Enter an appropriate *Label* for the VIP, e.g. **Cerberus-HTTPS**.
4. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.30.20.68**.
5. Set the *Virtual Service Ports* field to **443**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.

7. Click **Update**.
8. Click **Modify** next to the newly created VIP.
9. Scroll to the *Connection Distribution Method* section.
    - Set the *Balance Mode* to **Weighted Round Robin**.
10. Scroll to the *Persistence* section.
    - Ensure that the *Persistence Mode* is set to **Source IP**.
11. Scroll to the *Health Checks* section.
    - Set the *Health Checks* to **Connect to port**.
12. Leave all other settings at their default value.
13. Click **Update**.

> **Note** If the plaintext HTTP web listener is also enabled on the Cerberus servers, set the *Virtual Service Ports* field to **80,443**, leave the *Real Server Port* field blank when defining the Real Servers, and set the *Check Port* to **443** in the *Health Checks* section.

10.3.2. Define the Associated Real Servers (RIPs)

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



3. Define the *Label* for the Real Server as required, e.g. **Node-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **10.30.20.5**.
5. Set the *Real Server Port* field to **443**.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining Cerberus server(s), e.g. **Node-2** at **10.30.20.6**.

10.4. 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:

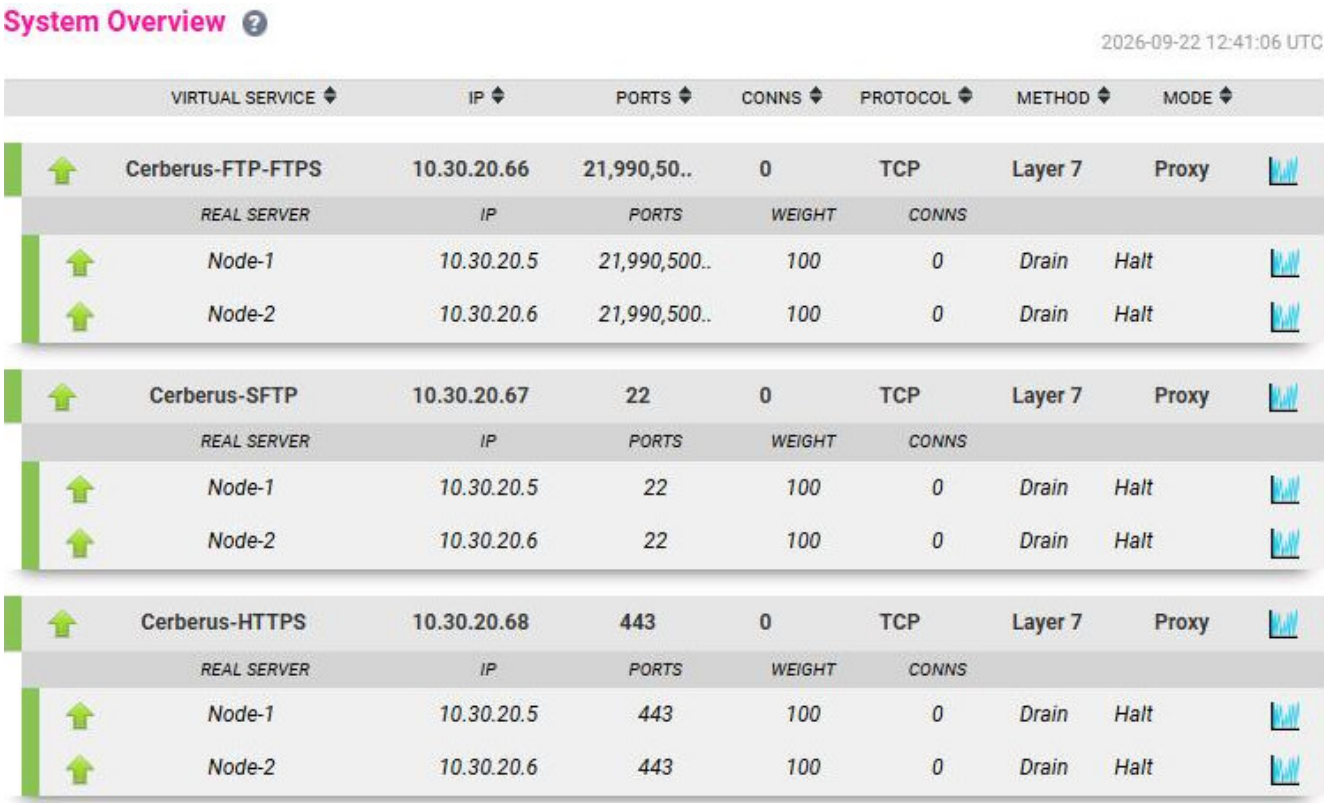
1. Using the WebUI, navigate to: ***Maintenance > Restart Services***.
2. Click **Reload HAProxy**.

11. Testing & Verification

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

11.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all Virtual Services & the associated Real Servers (i.e. the Cerberus servers) and shows the state/health of each server as well as the overall state of each cluster. The example below shows that all servers are healthy (green) and available to accept connections:



11.2. Accessing Cerberus FTP Server via the Load Balancer

Verify that you're able to successfully access all load balanced applications and services via the Virtual Services on the load balancer.

FTP & FTPS - using an FTP client such as FileZilla or WinSCP, connect to the VIP 1 address using plain FTP and explicit FTPS on port 21, and implicit FTPS on port 990 if it is enabled. Upload a test file and download it again, then confirm that the checksums match. Also confirm that the PASV reply hands back the VIP address rather

than the address of an individual node.

SFTP - connect to the VIP 2 address, e.g. `sftp testuser@10.30.20.67`, then upload and download a test file. Repeat the test after a failover and confirm that no host-key-changed warning is raised.

HTTPS web client - browse to the VIP 3 address, e.g. **https://10.30.20.68**, log in, and upload and download a file through the browser.

Replication - create a test user on the primary node, allow Sync Manager to run (or push the configuration manually), then verify that the new user can log in through the VIPs when the connection lands on a secondary node.

Failover - stop the Cerberus service on one node. The System Overview should mark that node as down and all of the tests above should continue to succeed via the remaining node(s). Restart the service and confirm that the node returns to the up state.

> **Note** Make sure that DNS is updated so that any FQDNs used point to the VIPs rather than individual servers.

12. Additional Considerations

- **The web admin console should not be load balanced** - the Cerberus web admin console (default HTTPS port 8443) is deliberately excluded from this design. All configuration changes must be made on the primary node only: Sync Manager pushes the primary's configuration to the secondary nodes, so changes made directly on a secondary node will be overwritten.
- **Automatic Threat Blocking** - Cerberus blocks by client IP address, and in this design every connection reaches the Cerberus servers from the load balancer rather than from the client. A handful of failed logins from a single user can therefore block the load balancer itself, and with it every other client. If *DoS Protection* is enabled, the load balancer's health checks also count towards the limit. To avoid this, add the load balancer's IP addresses - both appliances in a clustered pair - to the Cerberus IP Manager as **Never Block This IP**.
- **Automatic Threat Blocking cannot act on individual clients** - because every connection arriving through the Virtual Services shares the load balancer's address, Cerberus cannot block one client without blocking them all. If per-client blocking is required, the Virtual Services must be made transparent using TProxy - see A Note on Transparency.
- **HTTP/S public file sharing** - a share link created on a secondary server is not propagated to the other nodes. The vendor currently documents no workaround for this.
- **Licensing** - each Cerberus node requires its own license, and Sync Manager requires the Professional or Enterprise edition.

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

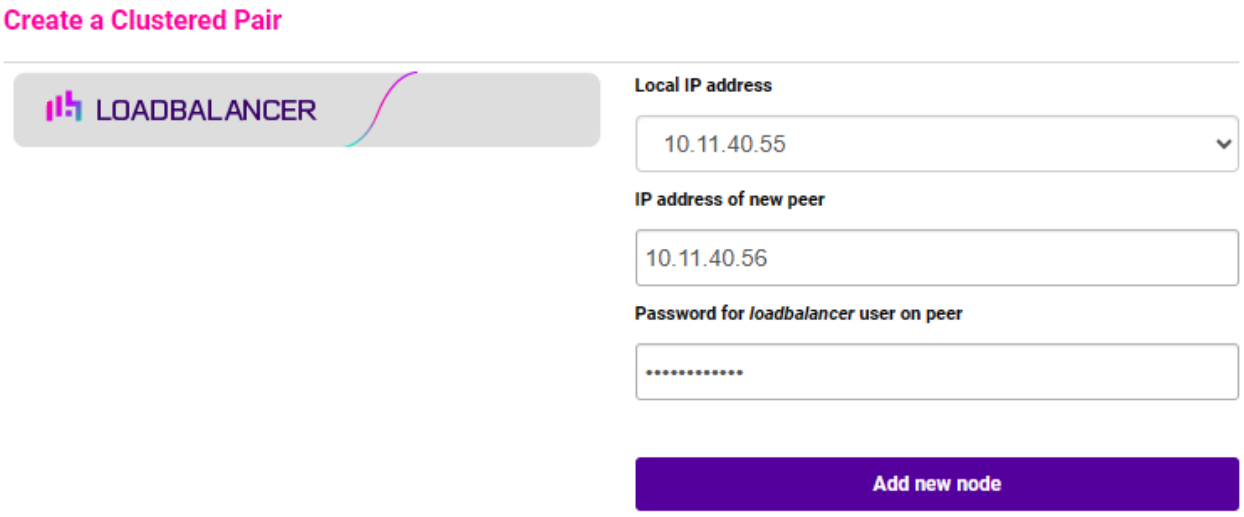
> **Important** Make sure that where any of the above have been configured on the Primary appliance, they're also configured on the Secondary.

15.1.2. Configuring the HA Clustered Pair

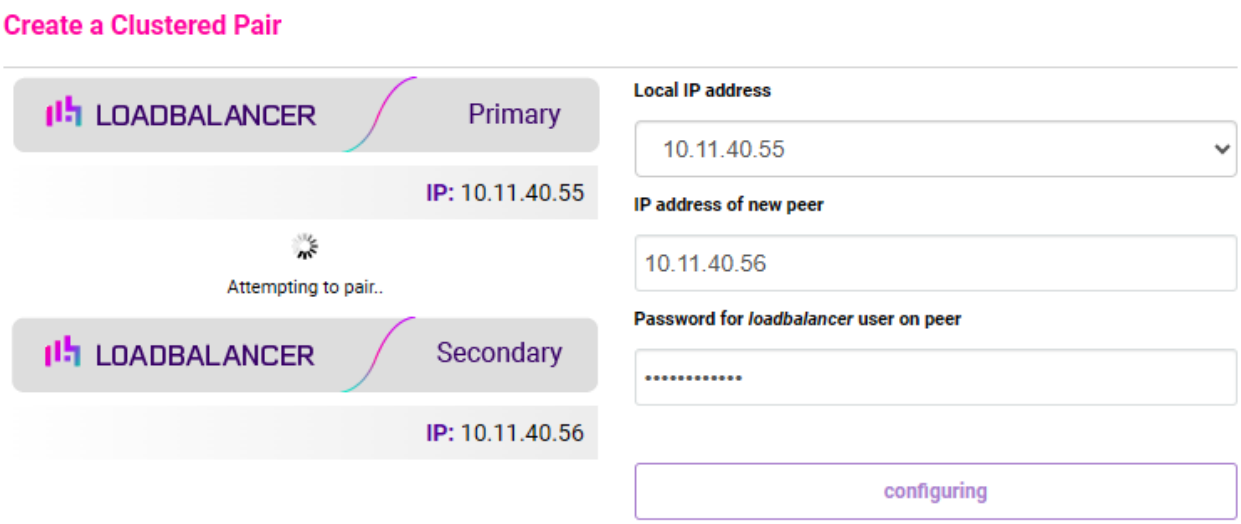
> **Important** During HA pairing, all WebUI users and passwords are synchronized from the Primary to the Secondary. After clustering completes (you will be logged out of the Secondary when this occurs), the Primary's credentials should be used to login to both nodes.

> **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.

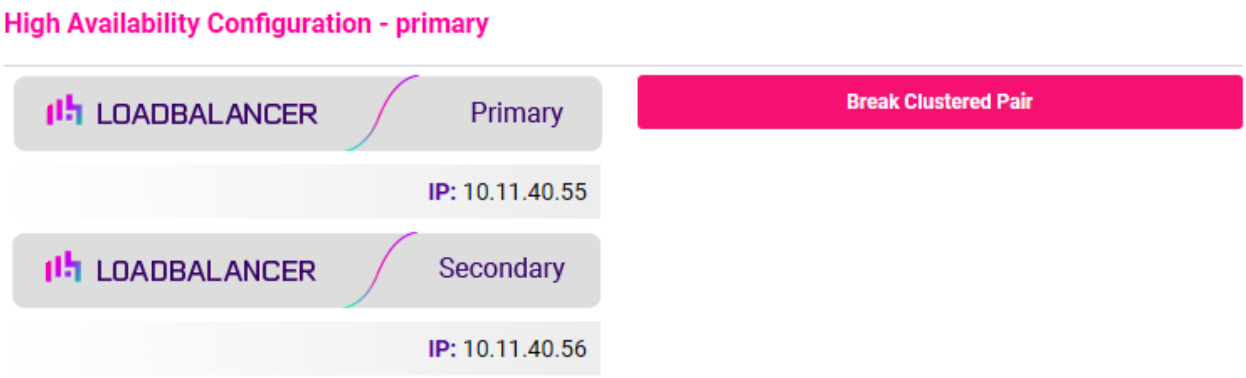
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: ***Cluster Configuration > High-Availability Configuration***.



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:



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



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.

16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	22 October 2026	Initial version		RJC

LOADBALANCER

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.

