

# Load Balancing CTERA Portal

Version 1.0.0



# Table of Contents

1. About this Brief	4
2. Loadbalancer.org Appliances Supported	4
3. Software Versions Supported	4
3.1. Loadbalancer.org Appliance	4
3.2. CTERA Portal	4
4. CTERA Portal	4
5. Load Balancing CTERA Portal	4
5.1. Load Balancing & HA Requirements	5
5.2. Virtual Service (VIP) Requirements	5
SSL Termination on the CTERA Portal Servers	5
SSL Termination on the Load Balancer	5
6. Deployment Concept	5
6.1. SSL Termination on the CTERAL Portal Servers	5
6.2. SSL Termination on the Load Balancer	6
7. Load Balancer Deployment Methods	6
7.1. Layer 7 SNAT Mode	6
8. Configuring CTERA Portal for Load Balancing	7
9. Loadbalancer.org Appliance – the Basics	8
9.1. Virtual Appliance	8
9.2. Initial Network Configuration	9
9.3. Accessing the Appliance WebUI	9
Main Menu Options	10
9.4. Appliance Software Update	11
Determining the Current Software Version	11
Checking for Updates using Online Update	11
Using Offline Update	11
9.5. Ports Used by the Appliance	12
9.6. HA Clustered Pair Configuration	13
10. Appliance Configuration for CTERA Portal - SSL Termination on the CTERA Portal Servers	13
10.1. VIP 1 - CTERA_Portal	13
Virtual Service (VIP) Configuration	13
Define the Associated Real Servers (RIPs)	14
11. Appliance Configuration for CTERA Portal - SSL Termination on the Load Balancer	15
11.1. VIP 1 - CTERA_Portal_Web_1	15
Virtual Service (VIP) Configuration	15
Define the Associated Real Servers (RIPs)	16
Upload the SSL Certificate	17
Configure SSL Termination	17
11.2. VIP 2 - CTERA_Portal_Web_2	18
Virtual Service (VIP) Configuration	18
Define the Associated Real Servers (RIPs)	20
Configure SSL Termination	20
11.3. VIP 3 - CTERA_Portal_CTTTP	21
Virtual Service (VIP) Configuration	21
Define the Associated Real Servers (RIPs)	22
11.4. Finalizing the Configuration	23
12. Testing & Verification	23
12.1. Accessing CTERA Portal via the Load Balancer	23

12.2. Using System Overview .....	23
13. Technical Support .....	24
14. Further Documentation .....	24
15. Appendix .....	25
15.1. Configuring HA - Adding a Secondary Appliance .....	25
Non-Replicated Settings .....	25
Configuring the HA Clustered Pair .....	26
16. Document Revision History .....	28

# 1. About this Brief

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

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. CTERA Portal

- v8.1.x and later

## 4. CTERA Portal

CTERA Portal is a scalable cloud service delivery platform installed in a user's own data center or in a cloud environment. It allows users to create, deliver and manage cloud storage applications, including a Global File System, file access via stubbing/caching, backup, and mobile collaboration.

CTERA Portal is compatible with cloud storage infrastructure from multiple vendors, including EMC, Wasabi, Scality, IBM COS, Hitachi HCP, and cloud storage providers such as AWS, Azure, Google Cloud Platform, and IBM Cloud.

## 5. Load Balancing CTERA Portal

#### Note

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



## 5.1. Load Balancing & HA Requirements

An installation of CTERA Portal consists of a Primary Server and optionally, one or more additional servers that mirror its configuration. This allows a high availability (HA) configuration of multiple servers that host one or more Virtual Portals. A load balancer is required to distribute traffic across multiple Virtual Portals and to facilitate HA behavior of the solution.

## 5.2. Virtual Service (VIP) Requirements

The VIPs required depend on where SSL termination is performed.

### SSL Termination on the CTERA Portal Servers

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	CTERA_Portal	L7 SNAT (TCP)	443,995,8443	Source IP	HTTPS (GET)

### SSL Termination on the Load Balancer

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	CTERA_Portal_Web_1	L7 SNAT (HTTP)	80	Source IP	HTTPS (GET)
VIP 2	CTERA_Portal_Web_2	L7 SNAT (HTTP)	8080	Source IP	HTTPS (GET)
VIP 3	CTERA_Portal_CCTP	L7 SNAT (TCP)	995	Source IP	Connect to Port

SSL Termination is configured on the load balancer for the following VIPs:

- VIP 1 - **CTERA\_Portal\_Web\_1**
- VIP 2 - **CTERA\_Portal\_Web\_2**

This provides a corresponding HTTPS Virtual Service on port 443 for VIP 1 and a corresponding HTTPS Virtual Service on port 8443 for VIP 2. Certificates in PEM or PFX format can be uploaded to the load balancer.


The connection from the load balancer to the CTERA Portal servers for these VIPs is also encrypted, providing full end-to-end encryption from client to server.

## 6. Deployment Concept

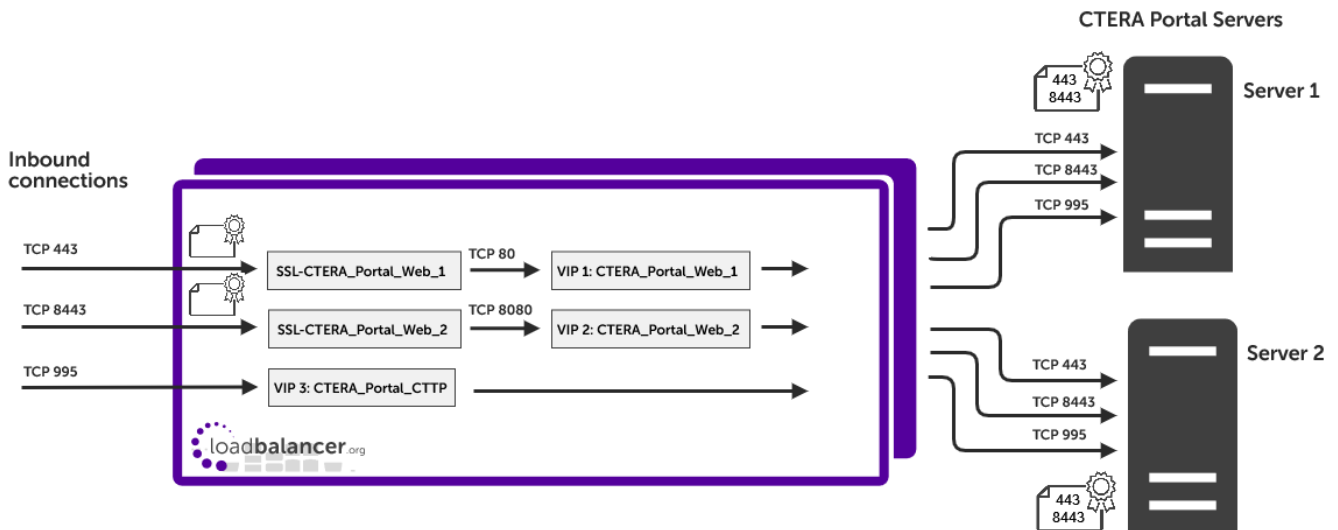
Once the load balancer is deployed, clients connect to the Virtual Service(s) (VIP(s)) rather than connecting directly to one of the CTERA Portal servers.

### 6.1. SSL Termination on the CTERAL Portal Servers





## 6.2. SSL Termination on the Load Balancer



### 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*.

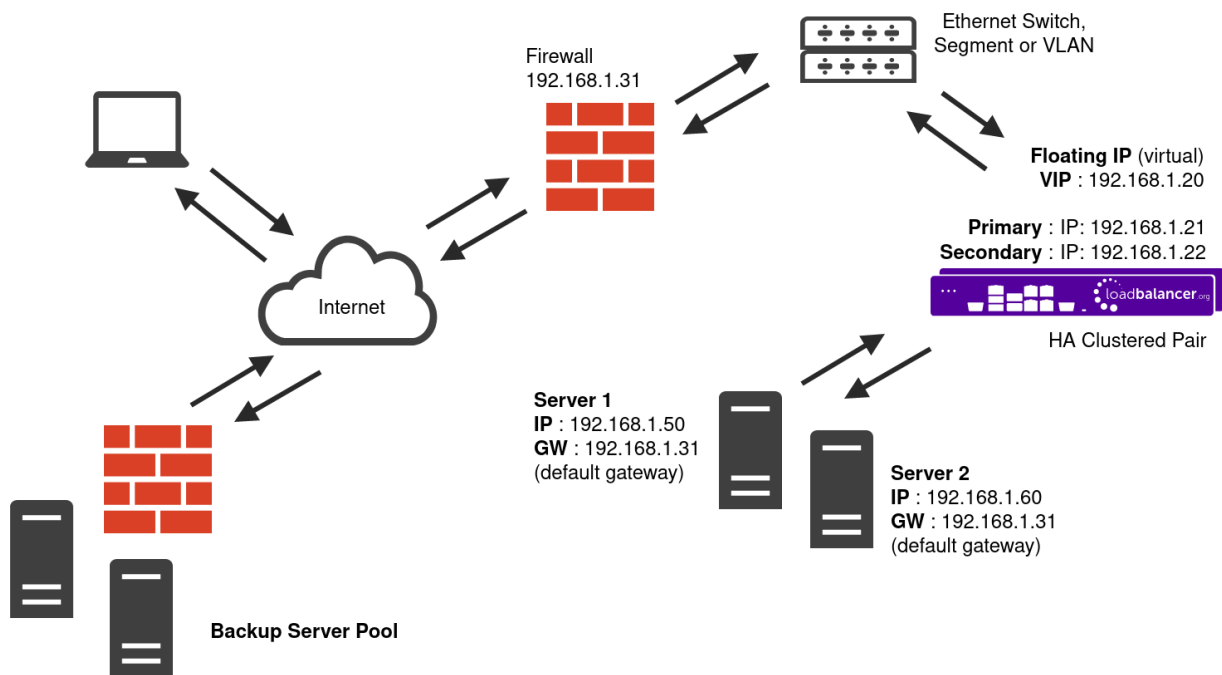
For CTERA Portal, layer 7 SNAT mode is recommended. This mode is described below and is used for the configuration presented in this guide.

### 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.



- 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.

## 8. Configuring CTERA Portal for Load Balancing

The CTPP keepalive interval must be configured using the Virtual Portal Settings window.

To configure this setting:

1. In the CTERA Portal Platform, navigate to *Virtual Portal Settings*.



2. Under the *Advanced* section, enable (check) the *Send CTTTP keepalive messages every... seconds* checkbox and set the value to **840**.
3. Click **Save** at the bottom of the dialog.

## 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.

**(!) 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    15 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** 2024-03-15 16:27:21 UTC

Would you like to run the Setup Wizard?

[Accept](#)    [Dismiss](#)

VIRTUAL SERVICE    IP    PORTS    CONNS    PROTOCOL    METHOD    MODE

No Virtual Services configured.

**Network Bandwidth**

Bytes/s

20 k  
15 k  
10 k  
5 k  
0

Thu 18:00    Fri 00:00    Fri 06:00    Fri 12:00

■ RX    3k Min, 4k Avg, 32675k Total  
■ TX    6k Min, 7k Avg, 56693k Total

**System Load Average**

System Load

1.0  
0.8  
0.6  
0.4  
0.2  
0.0

Thu 18:00    Fri 00:00    Fri 06:00    Fri 12:00

■ 1m average    0.00 Min, 0.12 Avg, 0.60 Max  
■ 5m average    0.00 Min, 0.06 Avg, 0.21 Max  
■ 15m average    0.00 Min, 0.02 Avg, 0.08 Max

Memory Usage

- You'll be asked if you want to run the Setup Wizard which can be used to configure layer 7 services. Click **Dismiss** if you're following a guide or want to configure the appliance manually or click **Accept** to start the wizard.

## 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

**Live Chat** - Start a live chat session with one of our Support Engineers

## 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 – 2024

ENTERPRISE VA Max - v8.11.4

English ▼

### Checking for Updates using Online Update

#### Note

By default, the appliance periodically contacts the Loadbalancer.org update server and checks for updates. An update check can also be manually triggered as detailed below.

1. Using the WebUI, navigate to: *Maintenance > Software Update*.
2. Select **Online Update**.
3. If the latest version is already installed, a message similar to the following will be displayed:

**Information:** Version v8.11.4 is the current release. No updates are available

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

#### Note

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

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

**Information:** Update completed successfully.

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

### Using Offline Update

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



## Note

Please contact [support@loadbalancer.org](mailto:support@loadbalancer.org) to check if an update is available and obtain the latest offline update files.

To perform an offline update:

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

### Software Update

#### Offline Update

The following steps will lead you through offline update.

1. Contact **Loadbalancer.org support** to obtain the offline update archive and checksum.
2. Save the archive and checksum to your local machine.
3. Select the archive and checksum files in the upload form below.
4. Click *Upload and Install* to begin the update process.

Archive:  No file chosen

Checksum:  No file chosen

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

## 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



Protocol	Port	Purpose
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 Addresses](#).

## 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 the section [Configuring HA - Adding a Secondary Appliance](#) of the appendix.

# 10. Appliance Configuration for CTERA Portal - SSL Termination on the CTERA Portal Servers

## 10.1. VIP 1 - CTERA\_Portal

### Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="CTERA_Portal"/>	?
IP Address	<input type="text" value="192.168.95.190"/>	?
Ports	<input type="text" value="443,995,8443"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

- Define the *Label* for the Virtual Service as required, e.g. **CTERA\_Portal**.
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.95.190**.
- Set the *Ports* field to **443,995,8443**.
- Set the *Layer 7 Protocol* to **TCP Mode**.
- Click **Update** to create the Virtual Service.



8. Now click **Modify** next to the newly created VIP.
9. Scroll to the *Protocol* section and click **[Advanced]**.
  - Enable (check) the *TCP Keep alive* checkbox.
10. Scroll to the *Health Checks* section and click **[Advanced]**.
  - Set the *Health Checks* to **Negotiate HTTPS (GET)**.
  - Set the *Request to Send* to **/admin/startup**.
  - Set the *Check Port* to **443**.
11. Scroll to the *Other* section and click **[Advanced]**.
  - Enable (check) the *Timeout* checkbox.
  - Set both *Client Timeout* and *Real Server Timeout* to **500m** (500 minutes).
12. Leave all other settings at their default value.
13. Click **Update**.

### Define the Associated Real Servers (RIPs)

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

Label	<input type="text" value="CTERA_Portal-1"/>	?
Real Server IP Address	<input type="text" value="192.168.95.180"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Define the *Label* for the Real Server as required, e.g. **CTERA\_Portal-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.95.180**.
5. Leave the *Real Server Port* field blank.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining CTERA Portal server(s).

# 11. Appliance Configuration for CTERA Portal - SSL Termination on the Load Balancer

## 11.1. VIP 1 - CTERA\_Portal\_Web\_1

### Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="CTERA_Portal_Web_1"/>	?
IP Address	<input type="text" value="192.168.95.190"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?

**Cancel** **Update**

3. Define the *Label* for the Virtual Service as required, e.g. **CTERA\_Portal\_Web\_1**.
4. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.95.190**.
5. Set the *Ports* field to **80**.
6. Set the *Layer 7 Protocol* to **HTTP Mode**.
7. Click **Update** to create the Virtual Service.
8. Now click **Modify** next to the newly created VIP.
9. Scroll to the *Persistence* section and click **[Advanced]**.
  - Set the *Persistence Mode* to **Application Cookie**.
  - Set the *Application Cookie Name* to **JSESSIONID**.
  - Set the *Persistence Timeout* to **180** (i.e. 180 seconds).
10. Scroll to the *Health Checks* section.
  - Set the *Health Checks* to **Negotiate HTTPS (GET)**.
  - Set the *Request to Send* to **/admin/startup**.
11. Scroll to the *ACL Rules* section, and click the **Add Rule** button.
12. Set the **Type** to **Freetype** and copy and paste the following ACL rules into the *Freetype* field ensuring that the default "#" in the *Freetype* field is first removed.

```
option tcpka
http-response replace-header Set-Cookie '^((?!(?i)httponly).)*$' "\1; HttpOnly"
http-response replace-header Set-Cookie '^((?!(?i)secure).)*$' "\1; Secure" if { ssl_fc }
```

### Note

These ACLs are used to ensure that the application cookie is correctly secured when it is transmitted back to the client, and that TCP keepalive mode is enabled.

- Click **OK** to save and close the ACL.
- Scroll to the *Fallback Server* section.
  - Enable (check) the *Disable Fallback* checkbox.
- Scroll to the *SSL* section.
  - Enable (check) *Backend Encryption*.
- Scroll to the *Other* section and click **[Advanced]**.
  - Set *Force to HTTPS* to **Yes**.
  - Set *Enable HSTS* to **Yes**.
- Leave all other settings at their default value.
- Click **Update**.

## Define the Associated Real Servers (RIPs)

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

Label	<input type="text" value="CTERA_Portal-1"/>	?
Real Server IP Address	<input type="text" value="92.168.95.180"/>	?
Real Server Port	<input type="text" value="443"/>	?
Re-Encrypt to Backend	<input checked="" type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

**Cancel** **Update**

- Define the *Label* for the Real Server as required, e.g. **CTERA\_Portal-1**.
- Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.95.180**.
- Set the *Real Server Port* field to **443**.



6. Ensure that *Re-Encrypt to Backend* is enabled (checked).
7. Leave all other settings at their default value.
8. Click **Update**.
9. Repeat these steps to add the remaining CTERA Portal server(s).

## Upload the SSL Certificate

Certificates in either PEM or PFX format can be uploaded.

1. Using the WebUI, navigate to *Cluster Configuration > SSL Certificate* and click **Add a new SSL Certificate**.
2. Select the option **Upload prepared PEM/PFX file**.
3. Enter the following details:

The screenshot shows a form for uploading an SSL certificate. It features three radio button options under the heading "I would like to:". The first option, "Upload prepared PEM/PFX file", is selected. Below this are two unselected options: "Create a new SSL Certificate Signing Request (CSR)" and "Create a new Self-Signed SSL Certificate.". A "Label" field contains the text "Cert-CTERA-Portal". A "File to upload" field contains a "Choose File" button and the filename "ctera-portal.pem". A purple "Upload Certificate" button is located at the bottom right of the form.

4. Specify an appropriate *Label*, e.g. **Cert-CTERA-Portal**.
5. Click **Choose File**.
6. Browse to and select the relevant PEM or PFX file.
7. For PFX files specify the password if required.
8. Click **Upload Certificate**.

## Configure SSL Termination

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

Label	SSL-CTERA_Portal_Web_1	?
Associated Virtual Service	CTERA_Portal_Web_1	?
Virtual Service Port	443	?
SSL Operation Mode	High Security	
SSL Certificate	cert-ctera_portal	?
Source IP Address		?
Enable Proxy Protocol	<input checked="" type="checkbox"/>	?
Bind Proxy Protocol to L7 VIP	CTERA_Portal_Web_1	?

Cancel
Update

- Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **CTERA\_Portal\_Web\_1**.

 **Note**

Once the VIP is selected, the *Label* field will be auto-populated with **SSL-CTERA\_Portal\_Web\_1**. This can be changed if preferred.

- Ensure that the *Virtual Service Port* is set to **443**.
- Leave *SSL Operation Mode* set to **High Security**.
- Select the *SSL Certificate* uploaded previously.
- Leave all other settings at their default value.
- Click **Update**.

## 11.2. VIP 2 - CTERA\_Portal\_Web\_2

### Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="CTERA_Portal_Web_2"/>	?
IP Address	<input type="text" value="192.168.95.190"/>	?
Ports	<input type="text" value="8080"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?

3. Define the *Label* for the virtual service as required, e.g. **CTERA\_Portal\_Web\_2**.
4. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.95.190**.
5. Set the *Ports* field to **8080**.
6. Set the *Layer 7 Protocol* to **HTTP Mode**.
7. Click **Update** to create the Virtual Service.
8. Now click **Modify** next to the newly created VIP.
9. Scroll to the *Persistence* section and click **[Advanced]**.
  - Set the *Persistence Mode* to **Application Cookie**.
  - Set the *Application Cookie Name* to **JSESSIONID**.
  - Set the *Persistence Timeout* to **180** (i.e. 180 seconds).
10. Scroll to the *Health Checks* section.
  - Set the *Health Checks* to **Negotiate HTTPS (GET)**.
  - Set the *Request to Send* to **/admin/startup**.
11. Scroll to the *ACL Rules* section, and click the **Add Rule** button.
12. Set the **Type** to *Freetype* and copy and paste the following ACL rules into the *Freetype* field ensuring that the default "#" in the *Freetype* field is first removed.

```
option tcpka
http-response replace-header Set-Cookie '(^((?!(?i)httponly).)*$)' "\1; HttpOnly"
http-response replace-header Set-Cookie '(^((?!(?i)secure).)*$)' "\1; Secure" if { ssl_fc }
```

 **Note**

These ACLs are used to ensure that the application cookie is correctly secured when it is transmitted back to the client, and that TCP keepalive mode is enabled.

13. Click **OK** to save and close the ACL.
14. Scroll to the *Fallback Server* section.

- Enable (check) the *Disable Fallback* checkbox.
15. Scroll to the *SSL* section.
    - Enable (check) *Backend Encryption*.
  16. Scroll to the *Other* section and click **[Advanced]**.
    - Set *Force to HTTPS* to **Yes**.
    - Set *Enable HSTS* to **Yes**.
  17. Leave all other settings at their default value.
  18. Click **Update**.

## Define the Associated Real Servers (RIPs)

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

Label	<input type="text" value="CTERA_Portal-1"/>	?
Real Server IP Address	<input type="text" value="192.168.95.180"/>	?
Real Server Port	<input type="text" value="8443"/>	?
Re-Encrypt to Backend	<input checked="" type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Define the *Label* for the Real Server as required, e.g. **CTERA\_Portal-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.95.180**.
5. Set the *Real Server Port* field to **8443**.
6. Ensure that *Re-Encrypt to Backend* is enabled (checked).
7. Leave all other settings at their default value.
8. Click **Update**.
9. Repeat these steps to add the remaining CTERA Portal server(s).

## Configure SSL Termination

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

Label	SSL-CTERA_Portal_Web_2	?
Associated Virtual Service	CTERA_Portal_Web_2 ▼	?
Virtual Service Port	8443	?
SSL Operation Mode	High Security ▼	
SSL Certificate	Default Self Signed Certificate ▼	?
Source IP Address		?
Enable Proxy Protocol	<input checked="" type="checkbox"/>	?
Bind Proxy Protocol to L7 VIP	CTERA_Portal_Web_2 ▼	?

Cancel
Update

- Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **CTERA\_Portal\_Web\_2**.

 **Note**

Once the VIP is selected, the *Label* field will be auto-populated with **SSL-CTERA\_Portal\_Web\_2**. This can be changed if preferred.

- Ensure that the *Virtual Service Port* is set to **8443**.
- Leave *SSL Operation Mode* set to **High Security**.
- Select the *SSL Certificate* uploaded previously.
- Leave all other settings at their default value.
- Click **Update**.

## 11.3. VIP 3 - CTERA\_Portal\_CTPP

### Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="CTERA_Portal_CTTTP"/>	?
IP Address	<input type="text" value="192.168.95.191"/>	?
Ports	<input type="text" value="995"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

3. Define the *Label* for the virtual service as required, e.g. **CTERA\_Portal\_CTTTP**.
4. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.95.191**.
5. Set the *Ports* field to **995**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update** to create the Virtual Service.
8. Now click **Modify** next to the newly created VIP.
9. Scroll to the *Protocol* section and click **[Advanced]**.
  - Enable (check) the *TCP Keep alive* checkbox.
10. Scroll to the *Other* section and click **[Advanced]**.
  - Enable (check) the *Timeout* checkbox.
  - Set both *Client Timeout* and *Real Server Timeout* to **500m** (500 minutes).
11. Leave all other settings at their default value.
12. Click **Update**.

## Define the Associated Real Servers (RIPs)

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

Label	<input type="text" value="CTERA_Portal-1"/>	?
Real Server IP Address	<input type="text" value="192.168.95.180"/>	?
Real Server Port	<input type="text" value="995"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Define the *Label* for the Real Server as required, e.g. **CTERA\_Portal-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.95.180**.
5. Set the *Real Server Port* field to **995**.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining CTERA Portal server(s).

## 11.4. Finalizing the Configuration

To apply the new settings, HAProxy and STunnel (if SSL is terminated on the load balancer) 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**.
3. Click **Reload STunnel** (if applicable).

## 12. Testing & Verification

### Note

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

### 12.1. Accessing CTERA Portal via the Load Balancer

Verify that users can connect to CTERA Portal via the load balancer.

### Note

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





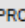
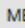
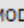






### 12.2. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the CTERA

Portal servers) and shows the state/health of each server as well as the state of the cluster as a whole. The example below shows a deployment where both CTERA Portal servers are healthy (green) and available to accept connections on both VIPs:

### System Overview

2024-09-03 14:03:46 UTC

	VIRTUAL SERVICE 	IP 	PORTS 	CONNS 	PROTOCOL 	METHOD 	MODE 	
	CTERA_Portal	192.168.95.190	443,995,8..	0	TCP	Layer 7	Proxy	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	CTERA_Portal-1	192.168.95.180	443,995,8443	100	0	Drain	Halt	
	CTERA_Portal-2	192.168.95.181	443,995,8443	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](mailto: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.

### 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	Fallback Page	Fallback page configuration
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.

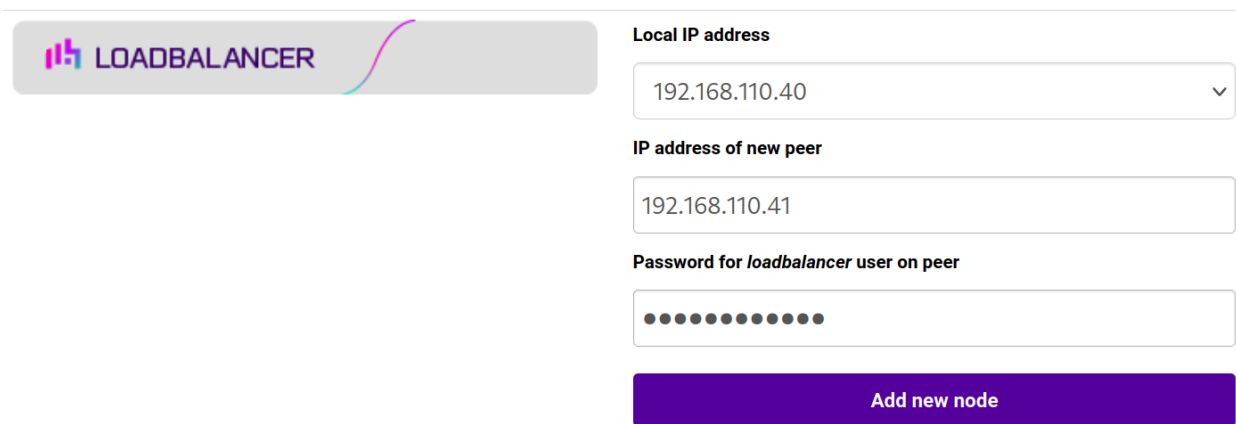
## Configuring the HA Clustered Pair

### 📌 Note

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

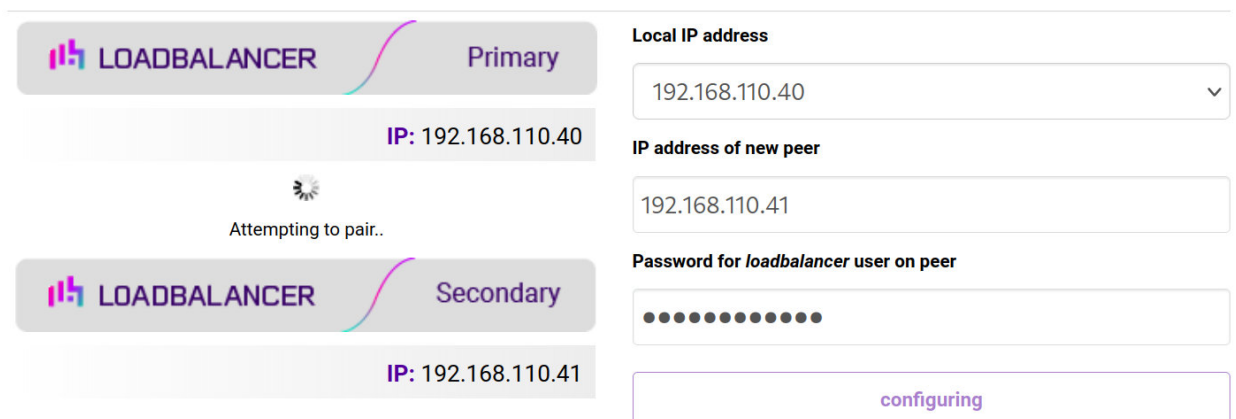
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: **Cluster Configuration > High-Availability Configuration**.

### Create a Clustered Pair



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

The screenshot displays a configuration interface for a High Availability (HA) setup. It features two load balancer appliances in a clustered pair. The top appliance is the Primary, with IP 192.168.110.40. The bottom appliance is the Secondary, with IP 192.168.110.41. A red button labeled "Break Clustered Pair" is located to the right of the appliances.

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	4 September 2024	Initial version		RJC





**Visit us:** [www.loadbalancer.org](http://www.loadbalancer.org)

**Phone us:** +44 (0)330 380 1064

**Phone us:** +1 833 274 2566

**Email us:** [info@loadbalancer.org](mailto:info@loadbalancer.org)

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

## About Loadbalancer.org

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

