

# Load Balancing Zadara VPSA Object Storage

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. Zadara VPSA Object Storage	4
4. Load Balancing Zadara VPSA Object Storage	
4.1. Port Requirements	
4.2. Deployment Concept	
4.3. Virtual Service (VIP) Requirements	
4.4. Deployment Mode	
5. Loadbalancer.org Appliance – the Basics	
5.1. Virtual Appliance	
5.2. Initial Network Configuration	
5.3. Accessing the Appliance WebUI	
Main Menu Options	
5.4. Appliance Software Update	
Determining the Current Software Version	
Checking for Updates using Online Update.	
Using Offline Update	
5.5. Ports Used by the Appliance	
5.6. HA Clustered Pair Configuration	
6. Appliance & VPSA Node Configuration	
6.1. Appliance Configuration	
Configuring VIP1 – OBS Data	
Configuring VIP2 – VPSA GUI	
Configuring VIP 3 – VPSA Authentication	
Finalizing the Configuration	
7. Additional Configuration Options & Settings	
7.1. SSL Termination	
7.2. SSL Termination on the load balancer - SSL Offloading	
Certificates	
Uploading Certificates	
7.3. Configuring SSL Termination on the Load Balancer	
Configure SSL Termination	
Finalizing the Configuration	
8. Testing & Verification	
8.1. Using System Overview	
9. Technical Support	
10. Further Documentation	
11. Appendix	
11.1. Alternative Load Balancing Method for Read-Intensive Deployments (Direct Routing)	
11.2. Caveats	
11.3. Appliance Configuration for Zadara VPSA Nodes – Using Layer 4 DR Mode (Direct Routing)	
Configuring VIP 1 – OBS Data	
Configuring VIP 2 – VPSA GUI	
Configuring VIP 3 – VPSA Authentication	
11.4. Configuring HA - Adding a Secondary Appliance	
Non-Replicated Settings	24

Adding a Secondary Appliance - Create an HA Clustered Pair	25
12. Document Revision History	27

## 1. About this Guide

This guide details the steps required to configure a highly available Zadara VPSA cluster environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Zadara VPSA configuration changes that are required.

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 Zadara VPSA Object Storage. The complete list of models is shown below:

All our products can be used for load balancing Zadara VPSA Object Storage. For full specifications of available models please refer to https://www.loadbalancer.org/products. Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

# 3. Software Versions Supported

## 3.1. Loadbalancer.org Appliance

V8.4.1 and later

8 Note

The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If using an older software version, note that the screenshots presented here may not match the WebUI exactly.

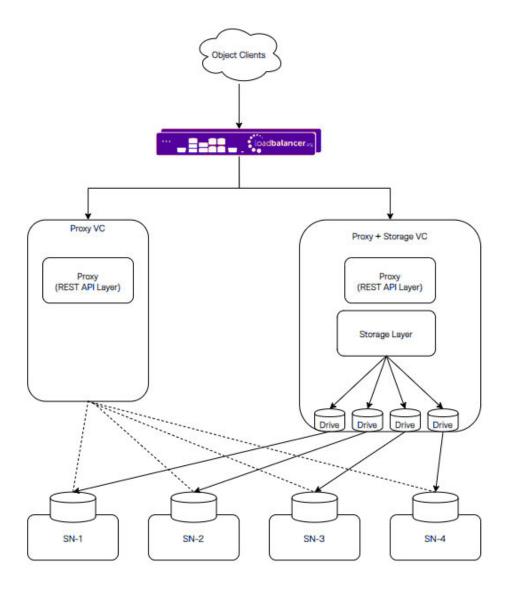
# 3.2. Zadara VPSA Object Storage

All versions

# 4. Load Balancing Zadara VPSA Object Storage

VPSA Object Storage is Zadara's object storage service. It is provided on Zadara clouds, side by side with the VPSA that provides block and file storage services.

VPSA Object Storage architecture is a scale out cluster of Virtual Controllers that together provide the service. The number Of VC's is automatically determined as needed to serve the capacity and performance of the system.



This figure shows high level logical view of VPSA Object Storage. It is a Virtual Object Store cluster, with two distinct layers:

- "Storage Layer" that manages individual disks
- "Proxy REST API Layer" that provides a REST API front-end to the Object Storage.

The typical VC runs both functions and is referred to as "Proxy+Storage" VC. It is possible to add VCs with the Proxy layer only. There are referred to as "Proxy" VC.

Each VPSA Object Storage is typically composed of several Proxy+Storage VCs and optionally one or more Proxy VCs with each VC having dedicated CPU / RAM / networking. Proxy+Storage VCs consume raw Physical drives (like SAS / SATA / SSD) exposed from Storage Nodes (SNs). Proxy+Storage and Proxy VCs run the 'Object Storage Stack' which provides Amazon S3 and Swift REST API interfaces.

Capacity & Performance can be independently scaled up/down by adding/removing disks and proxy-VCs respectively. VPSA Object Storage typically has a set of load balancers to distribute REST API traffic across the Proxy REST API Layers. Each VPSA Object Storage natively being multi-tenant allows for the creation of multiple accounts within it, with each account having multiple users who can work with the object interface (GET/PUT objects).

A single Zadara Storage Cloud can host several virtual object stores, making it truly disruptive and unique. Each VPSA Object Store has entirely provisioned resources of CPU / RAM / networking / disks and runs the object stack in isolated Virtual Machines (i.e. there is no sharing of resources anywhere across VPSAs) thereby providing complete performance and fault isolation.

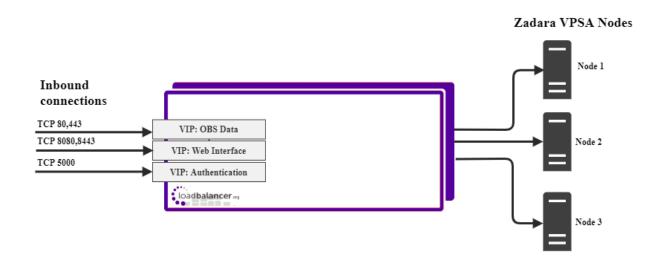
### 4.1. Port Requirements

The following table shows the ports used by the Zadara VPSA nodes. The load balancer must be configured to listen on the same ports.

ort Protocols		Use	
80, 443	TCP / HTTP, HTTPS	Object storage data	
8080, 8443	TCP / HTTP, HTTPS	Web interface	
5000	TCP	Authentication	

# 4.2. Deployment Concept

When the Zadara VPSA nodes are deployed with the load balancer, clients connect to the Virtual Service (VIP) on the load balancer rather than connecting directly to one of the VPSA nodes.



8 Note

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

# 4.3. Virtual Service (VIP) Requirements

To provide load balancing for Zadara VPSA nodes the following VIPs are required

- VIP 1: OBS Data
- VIP 2: VPSA GUI
- VIP 3: VPSA Authentication



### 4.4. Deployment Mode

We recommend using Layer 7 as no network changes are required and SSL termination can be implemented. This mode offers high performance and implementation flexibility, however as Layer 7 is a reverse proxy the client source IP address is not visible at the real server. Instead, the IP address of the load balancer is visible at the real server. In order to retain the client source IP address, the load balancer inserts an *X-Forwarded-For* header into the load balanced traffic, which the VPSA nodes can log for troubleshooting issues while seeing the true source IP address of connecting clients.

# 5. Loadbalancer.org Appliance – the Basics

## 5.1. 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.

å Note	The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.
8 Note	Please refer to Virtual Appliance Installation and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.
	The VA has 4 network adapters. For VM were only the first adapter (ath0) is connected by
å Note	The VA has 4 network adapters. For VMware only the first adapter ( <b>eth0</b> ) 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.

# 5.2. Initial Network Configuration

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



## 5.3. 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.	es. For details,
---	------------------

8 Note

A number of compatibility issues have been found with various versions of Microsoft Internet Explorer and Edge. The WebUI has been tested and verified using both Chrome & Firefox.

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

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

8 Note

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

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

Username: loadbalancer

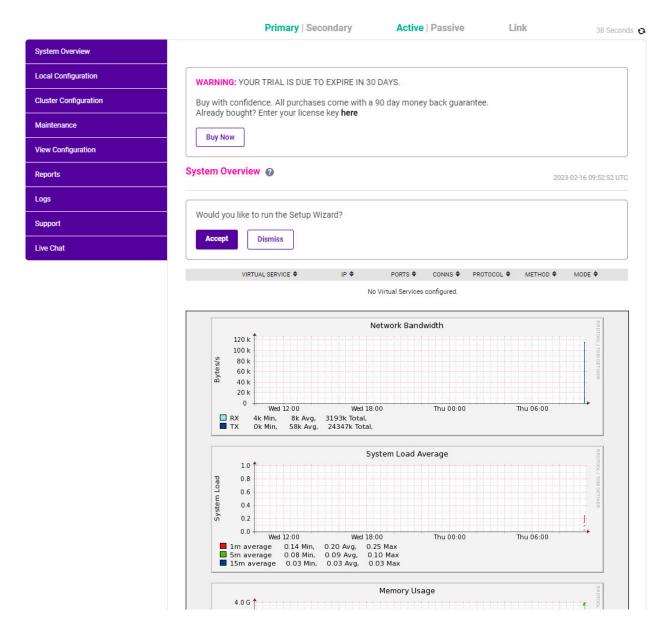
Password: <configured-during-network-setup-wizard>

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

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

IL LOADBALANCER





3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

#### Main Menu Options

System Overview - Displays a graphical summary of all VIPs, RIPs and key appliance statistics

Local Configuration - Configure local host settings such as IP address, DNS, system time etc.

Cluster Configuration - Configure load balanced services such as VIPs & RIPs

Maintenance - Perform maintenance tasks such as service restarts and taking backups

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

Support - Create a support download, contact the support team & access useful links



## 5.4. Appliance Software Update

To ensure that the appliance(s) are running the latest software version, we recommend a software update check is performed.

#### **Determining the Current Software Version**

The software version is displayed at the bottom of the WebUI as shown in the example below:

Copyright © Loadbalancer.org Inc. 2002 – 2023 ENTERPRISE VA Max - v8.9.0



#### Checking for Updates using Online Update

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

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

Information: Version v8.9.0 is the current release. No updates are available

- 4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
- 5. Click **Online Update** to start the update process.
  - Note Do not navigate away whilst the update is ongoing, this may cause the update to fail.
- 6. Once complete (the update can take several minutes depending on download speed and upgrade version) the following message will be displayed:

Information: Update completed successfully.

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

#### **Using Offline Update**

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



8 Note

Please contact support@loadbalancer.org to check if an update is available and obtain the latest offline update files.

#### To perform an offline update:

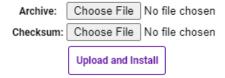
- 1. Using the WebUI, navigate to: *Maintenance > Software Update*.
- 2. Select Offline Update.
- 3. The following screen will be displayed:

#### **Software Update**

#### Offline Update

The following steps will lead you through offline update.

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



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

# 5.5. Ports Used by the Appliance

By default, the appliance uses the following TCP & UDP ports:

Protocol	Port	Purpose	
TCP	22	SSH	
TCP & UDP	53	DNS	
TCP & UDP	123	NTP	
TCP & UDP	161	SNMP	
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode	
TCP	7778	HAProxy persistence table replication	
TCP	9080	WebUI - HTTP (disabled by default)	
TCP	9081	Nginx fallback page	
TCP	9443	WebUI - HTTPS	

## 5.6. 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.

# 6. Appliance & VPSA Node Configuration

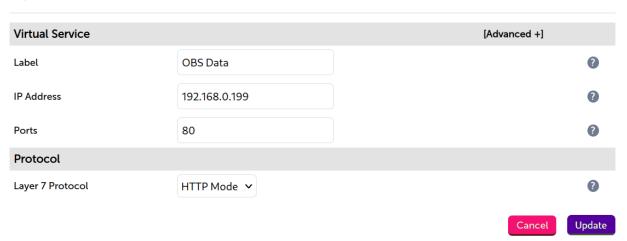
## 6.1. Appliance Configuration

#### Configuring VIP1 - OBS Data

#### a) Setting up the Virtual Service (VIP)

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

Layer 7 - Add a new Virtual Service



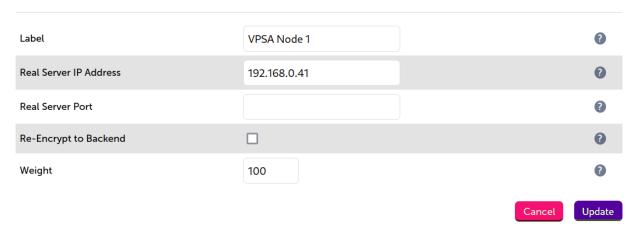
- 3. Enter an appropriate label (name) for the VIP, e.g. OBS Data.
- 4. Set the Virtual Service IP address field to the required IP address, e.g. 192.168.0.199.
- 5. Set the Virtual Service Ports field to 80.
- 6. Set *Protocol* to **HTTP mode**.
- 7. Click Update.
- 8. Click Modify next to the newly created VIP.
- 9. Set Persistence Mode to None.
- 10. Set Health Checks to Negotiate HTTP (GET).
- 11. Set the *Request to send* to /healthcheck.
- 12. Click [Advanced] and set the Check Port to 80.
- 13. Under the Other section click [Advanced].

- 14. Under *Timeout* check the box.
- 15. Set the Client Timeout and Real Server Timeout to 50000.
- 16. Click Update.

#### b) Setting up the 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 OBS Data VIP.
- 2. Enter the following details:

#### Layer 7 Add a new Real Server - OBS\_Data



- 3. Enter an appropriate label (name) for the RIP, e.g. VPSA Node 1.
- 4. Set the Real Server IP Address field to the IP address of the VPSA node 1.
- 5. Click **Update**.
- 6. Repeat these steps to add additional VPSA nodes as real servers as required.

#### Configuring VIP2 - VPSA GUI

#### a) Setting up the Virtual Service (VIP)

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

#### Layer 7 - Add a new Virtual Service

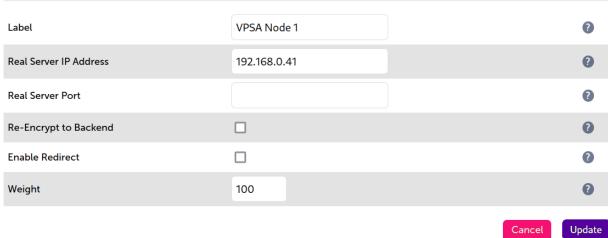


- 3. Enter an appropriate label (name) for the VIP, e.g. VPSA GUI.
- 4. Set the Virtual Service IP address field to the required IP address, e.g. 192.168.0.199.
- 5. Set the Virtual Service Ports field to 8080.
- 6. Set *Protocol* to HTTP Mode.
- 7. Click Update.
- 8. Click Modify next to the newly created VIP.
- 9. Set Persistence Mode to None.
- 10. Set Health Checks to Negotiate HTTP (HEAD).
- 11. Set the *Request to send* to /healthcheck.
- 12. Click [Advanced] and set the Check Port to 8080.
- 13. Under the *Other* section click [Advanced].
- 14. Under *Timeout* check the box.
- 15. Set the Client Timeout and Real Server Timeout to 50000.
- 16. Click Update.

#### b) Setting up the 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 VPSA Cluster VIP.
- 2. Enter the following details:

#### Layer 7 Add a new Real Server - VPSA\_GUI

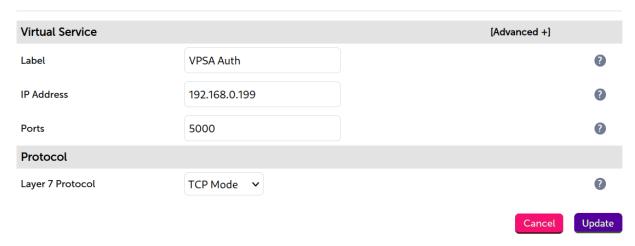


- 3. Enter an appropriate label (name) for the RIP, e.g. VPSA Node 1.
- 4. Set the Real Server IP Address field to the IP address of the VPSA node 1.
- 5. Click Update.
- 6. Repeat these steps to add additional VPSA nodes as real servers as required.

#### Configuring VIP 3 – VPSA Authentication

- a) Setting up the Virtual Service (VIP)
  - 1. Using the WebUI, navigate to Cluster Configuration > Layer 7 Virtual Services and click Add a new Virtual Service.
  - 2. Enter the following details:

Layer 7 - Add a new Virtual Service

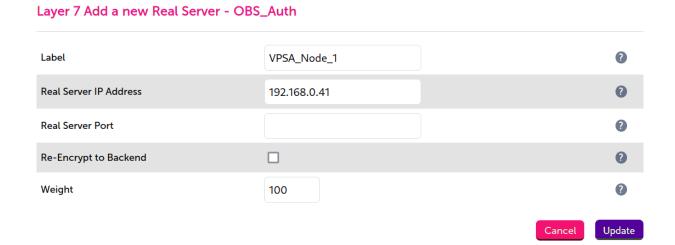


- 3. Enter an appropriate label (name) for the VIP, e.g. **VPSA Auth**.
- 4. Set the Virtual Service IP address field to the required IP address, e.g. 192.168.0.199.
- 5. Set the Virtual Service Ports field to 5000.

- 6. Leave Protocol set to TCP.
- 7. Click Update.
- 8. Click Modify next to the newly created VIP.
- 9. Set Persistence Mode to None.
- 10. Set Health Checks to Connect to Port.
- 11. Click Update.

#### b) Setting up the 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 VPSA Cluster VIP.
- 2. Enter the following details:



- 3. Enter an appropriate label (name) for the RIP, e.g. VPSA Node 1.
- 4. Set the Real Server IP Address field to the IP address of the VPSA node 1.
- 5. Click Update.
- 6. Repeat these steps to add additional VPSA nodes as real servers as required.

#### 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.

# 7. Additional Configuration Options & Settings

#### 7.1. SSL Termination

SSL termination can be handled in the following ways:

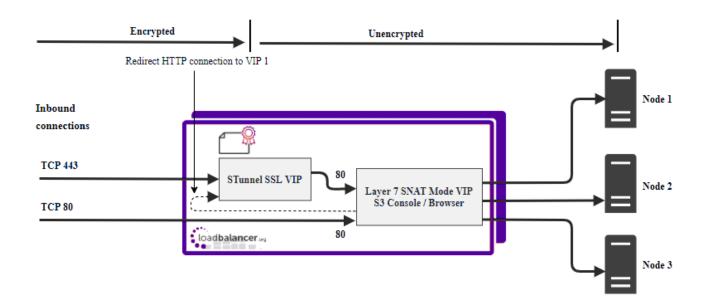


- 1. On the Real Servers aka SSL Pass-through.
- 2. On the load balancer aka SSL Offloading.
- 3. On the load balancer with re-encryption to the backend servers aka SSL Bridging.

By default, a self-signed certificate is used for the new SSL VIP. Certificates can be requested on the load balancer or uploaded as described in the section below. The default self-signed certificate can be regenerated if needed using the WebUI menu option: SSL Certificate and clicking the Regenerate Default Self Signed Certificate button.

The backend for the SSL VIP can be either a Layer 7 SNAT mode VIP or a Layer 4 NAT or SNAT mode VIP. Layer 4 DR mode cannot be used since stunnel acts as a proxy, and the VPSA node servers see requests with a source IP address of the VIP. However, since the VPSA node servers believe that they own the VIP (due to the loopback adapter configured to handle to ARP problem) they are unable to reply to stunnel.

## 7.2. SSL Termination on the load balancer - SSL Offloading



In this case, an SSL VIP utilizing stunnel is configured on the appliance and an SSL certificate is uploaded and associated to the Virtual Service. Data is encrypted from the client to the load balancer, but is un-encrypted from the load balancer to the backend servers as shown above.

#### Certificates

If you already have an SSL certificate in either PFX or PEM file format, this can be uploaded to the Load balancer using the certificate upload option as explained in Uploading Certificates. Alternatively, you can create a Certificate Signing Request (CSR) on the load balancer and send this to your CA to create a new certificate. For more information please refer to Generating a CSR on the Load Balancer.

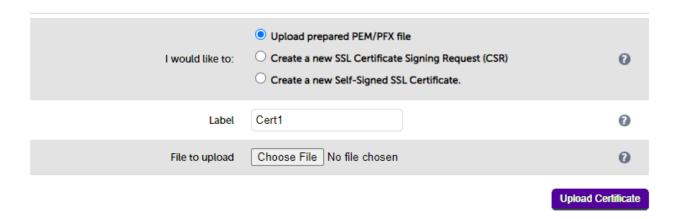
## **Uploading Certificates**



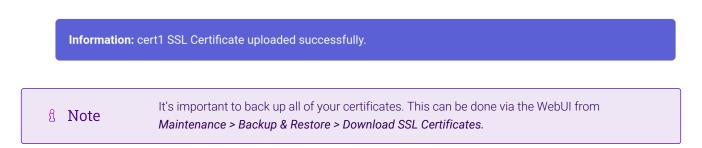
If you already have a certificate in either PEM or PFX format, this can be uploaded to the load balancer.

To upload a Certificate:

- 1. Using the WebUI, navigate to: Cluster Configuration > SSL Certificates.
- 2. Click Add a new SSL Certificate & select Upload prepared PEM/PFX file.



- 3. Enter a suitable Label (name) for the certificate, e.g. Cert1.
- 4. Browse to and select the certificate file to upload (PEM or PFX format).
- 5. Enter the password, if applicable.
- 6. Click **Upload Certificate**. If successful, a message similar to the following will be displayed:

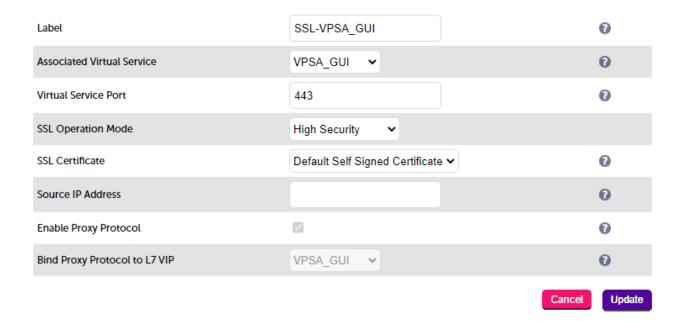


# 7.3. Configuring SSL Termination on the Load Balancer

Configure SSL termination for the VPSA GUI and OBS Data VIPs.

#### **Configure SSL Termination**

1. Using the WebUI, navigate to Cluster Configuration > SSL Termination and click Add a new Virtual Service.



2. Set *Associated Virtual Service* to the appropriate VIP, e.g. **VPSA GUI**. This will automatically fill in the label as the VIP name with SSL inserted in front of the VIP name e.g. **SSL-VPSA GUI**.

The Associated Virtual Service drop-down is populated with all single port, standard (i.e. non-manual) Layer 7 VIPs available on the load balancer. Using a Layer 7 VIP for the backend is the recommended method although as mentioned earlier, Layer 4 NAT mode and layer 4 SNAT mode VIPs can also be used if required. To forward traffic from the SSL VIP to these type of VIPs, you'll need to set Associated Virtual Service to **Custom**, then configure the IP address & port of the required VIP.

- 3. Leave Virtual Service Port set to 443.
- 4. Leave SSL operation Mode set to High Security.
- 5. Select the required certificate from the *SSL Certificate* drop-down.
- 6. Click Update.

Now repeat the above steps for the **OBS Data** VIP. In this case set *Associated Virtual Service* to **OBS Data**.

Once configured, HTTP traffic will be load balanced by the Layer 7 SNAT mode VIP and HTTPS traffic will be terminated by the SSL VIP, then passed on to the Layer 7 SNAT mode VIP as unencrypted HTTP for load balancing.

#### Finalizing the Configuration

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

- 1. Using the WebUl, navigate to: *Maintenance > Restart Services*.
- 2. Click Reload HAProxy.
- 3. Click Reload STunnel.

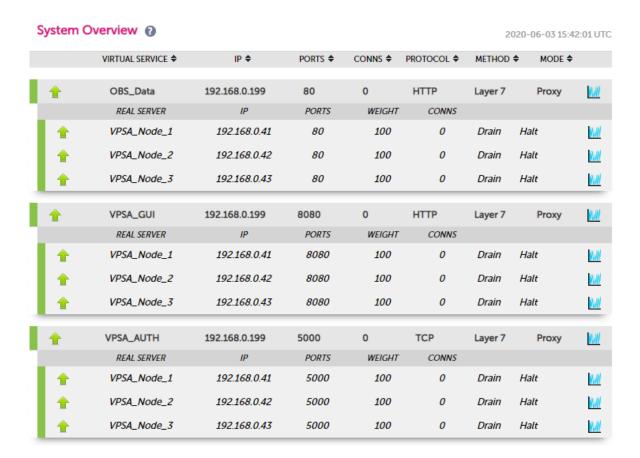
# 8. Testing & Verification

8 Note

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

### 8.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. VPSA\_GUI and VPSA\_Auth) and shows the state/health of each server as well as the state of the each cluster as a whole. The example below shows that all VPSA nodes are healthy and available to accept connections.



# 9. 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.

# 10. Further Documentation

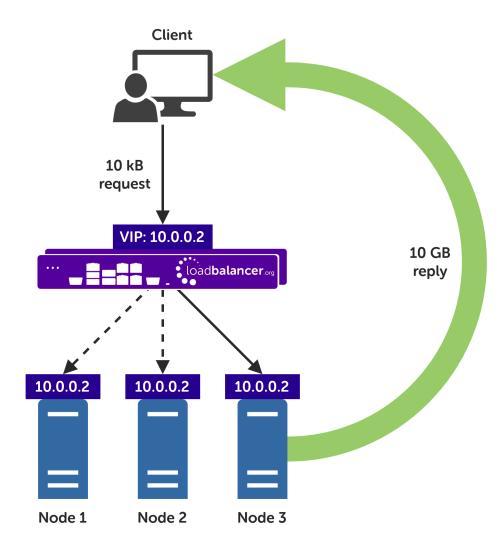
For additional information, please refer to the Administration Manual.

# 11. Appendix

# 11.1. Alternative Load Balancing Method for Read-Intensive Deployments (Direct Routing)

Direct routing, also known as direct server return or DSR, is a method of load balancing. With direct routing, reply traffic flows directly from the back end servers to the clients. In this way, the load balancer is completely bypassed on the return journey for a given connection, thus removing the load balancer as a potential bottleneck for traffic on the return path.

This alternative method of load balancing can benefit read-intensive deployments which feature a large reply traffic to request traffic ratio. For example, consider the scenario where a typical client request is 10 kB in size while a typical reply is 10 GB in size (perhaps file retrieval or video streaming). Direct routing benefits such scenarios: the much larger volume of reply traffic bypasses the load balancer and is *not* limited by the load balancer's network throughput. The reply traffic is instead limited by the total available network bandwidth between the servers and the clients, which is limited only by the underlying infrastructure.



#### 11.2. Caveats

There are caveats for using the direct routing load balancing method which should be considered:

- The load balancers must be on the same network segment / switching fabric as the VPSA nodes (due to the
  fact that this load balancing method works by rewriting MAC addresses, i.e. operates at layer 2 of the OSI
  model).
- Each VPSA node must own the VIP address so that they can all accept and reply to the load balanced traffic. This address should be assigned to a loopback network adaptor.
- Each VPSA node must be configured to not reply to ARP requests for the VIP address or advertise that they own the address.

For guidance on configuring the VPSA nodes for direct routing, in the context of the caveats described above, please consult with the Zadara team or Support.

# 11.3. Appliance Configuration for Zadara VPSA Nodes – Using Layer 4 DR Mode (Direct Routing)

#### Configuring VIP 1 - OBS Data

#### Configuring the Virtual Service (VIP)

- 1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 Virtual Services* and click on **Add** a new Virtual Service.
- 2. Define the Label for the virtual service as required, e.g. OBS\_Data.
- 3. Set the Virtual Service IP Address field to the required IP address, e.g. 192.168.0.167.
- 4. Set the Ports field to 80.
- 5. Leave the *Protocol* set to **TCP**.
- 6. Leave the Forwarding Method set to Direct Routing.
- 7. Click **Update** to create the virtual service.
- 8. Click Modify next to the newly created VIP.
- 9. Ensure that the Persistence Enable checkbox is unchecked.
- 10. Set the *Health Checks Check Type* to **Negotiate**.
- 11. Set the Check Port to 80.
- 12. Set the Protocol to HTTP.
- 13. Set the *Request to send* to /healthcheck.
- 14. Click **Update**.

#### 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. **VPSA-node1**.
- 3. Set the Real Server IP Address field to the required IP address, e.g. 192.168.0.41.
- 4. Click Update.



5. Repeat these steps to add additional VPSA nodes as real servers as required.

#### Configuring VIP 2 - VPSA GUI

#### Configuring the Virtual Service (VIP)

- 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. VPSA\_GUI.
- 3. Set the Virtual Service IP Address field to the required IP address, e.g. 192.168.0.167.
- 4. Set the Ports field to 8080.
- 5. Leave the *Protocol* set to **TCP**.
- 6. Leave the Forwarding Method set to Direct Routing.
- 7. Click **Update** to create the virtual service.
- 8. Click Modify next to the newly created VIP.
- 9. Ensure that the *Persistence Enable* checkbox is unchecked.
- 10. Set the *Health Checks Check Type* to **Negotiate**.
- 11. Set the Check Port to 8080.
- 12. Set the Protocol to HTTP.
- 13. Set the *Request to send* to /healthcheck.
- 14. Click Update.

#### **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. **VPSA-node1**.
- 3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.0.41**.
- 4. Click Update.
- 5. Repeat these steps to add additional VPSA nodes as real servers as required.

#### Configuring VIP 3 - VPSA Authentication

#### Configuring the Virtual Service (VIP)

- 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. **VPSA\_Auth**.
- 3. Set the Virtual Service IP Address field to the required IP address, e.g. 192.168.0.167.
- 4. Set the Ports field to 5000.
- 5. Leave the *Protocol* set to **TCP**.



- 6. Leave the Forwarding Method set to Direct Routing.
- 7. Click **Update** to create the virtual service.
- 8. Click Modify next to the newly created VIP.
- 9. Ensure that the *Persistence Enable* checkbox is unchecked.
- 10. Click **Update**.

#### **Defining the Real Servers (RIPs)**

- Using the web user interface, navigate to Cluster Configuration > Layer 4 Real Servers and click on Add a new Real Server next to the newly created VIP.
- 2. Define the *Label* for the real server as required, e.g. **VPSA-node1**.
- 3. Set the Real Server IP Address field to the required IP address, e.g. 192.168.0.41.
- 4. Click Update.
- 5. Repeat these steps to add additional VPSA nodes as real servers as required.

## 11.4. Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution.

We recommend that the Primary appliance is configured first and then the Secondary should be added. Once the Primary and Secondary are paired, all load balanced services configured on the Primary are automatically replicated to the Secondary over the network using SSH/SCP.

8 Note

For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

The clustered HA pair uses Heartbeat to determine the state of the other appliance. Should the active device (normally the Primary) suffer a failure, the passive device (normally the Secondary) will take over.

#### **Non-Replicated Settings**

A number of settings are not replicated as part of the Primary/Secondary pairing process and therefore must be manually configured on the Secondary appliance. These are listed by WebUI menu option in the table below:

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	All network settings including IP address(es), bonding configuration and VLANs
Local Configuration	Routing	Routing configuration including default gateways and static routes

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	System Date & time	All time and date related settings
Local Configuration	Physical – Advanced Configuration	Various settings including Internet Proxy, Management Gateway, Firewall connection tracking table size, NIC offloading, SMTP relay, logging and Syslog Server
Local Configuration	Security	Appliance security settings
Local Configuration	SNMP Configuration	Appliance SNMP settings
Local Configuration	Graphing	Appliance graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Software Updates	Appliance software update management
Maintenance	Firewall Script	Appliance firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

(!) Important

Make sure that if these settings/updates have been configured on the Primary appliance, they're also configured on the Secondary appliance.

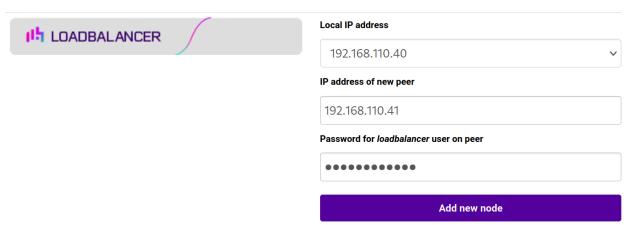
#### Adding a Secondary Appliance - Create an HA Clustered Pair

Note

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

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

#### **Create a Clustered Pair**

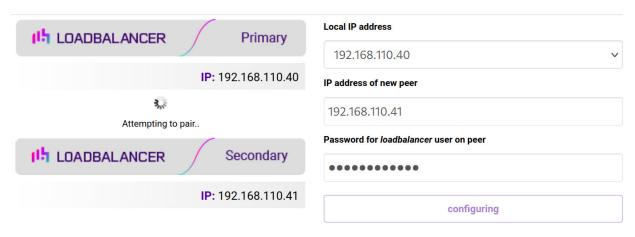


3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.



- 4. Click Add new node.
- 5. The pairing process now commences as shown below:

#### **Create a Clustered Pair**

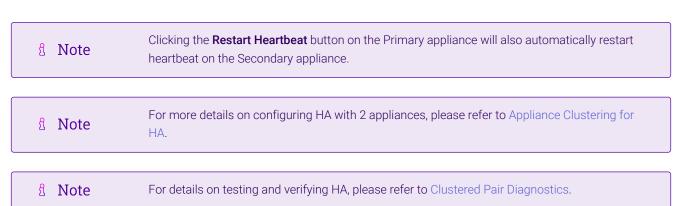


6. Once complete, the following will be displayed on the Primary appliance:

#### **High Availability Configuration - primary**



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



# 12. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	1 April 2020	Initial version		IBG
1.0.1	3 June 2020	VIP Configuration  New title page	Added new OBS data vip and SSL termination	IBG, AH
		Updated Canadian contact details	Branding update	
			Change to Canadian contact details	
1.1.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH,RJC,ZAC
1.1.1	28 January 2022	Remove references to retired ZIOS branding	Request from Zadara	АН
1.1.2	26 April 2022	Updated SSL related content to reflect latest software version	New software release	RJC
1.1.3	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	АН
1.1.4	5 January 2023	Combined software version information into one section  Added one level of section numbering  Added software update instructions  Added table of ports used by the appliance  Reworded 'Further Documentation' section  Removed references to the colour of certain UI elements	Housekeeping across all documentation	АН
1.1.5	2 February 2023	Updated screenshots	Branding update	АН
1.1.6	7 March 2023	Removed conclusion section	Updates across all documentation	АН
1.2.0	24 March 2023	New document theme  Modified diagram colours	Branding update	АН



Visit us: www.loadbalancer.org

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

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

Email us: info@loadbalancer.org

Follow us: @loadbalancer.org

## **About Loadbalancer.org**

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

