LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing Metaswitch Virtual EAS SSS

Version 1.3.0

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced Metaswitch Virtual EAS SSS environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Metaswitch Virtual EAS SSS configuration changes that are required to enable load balancing.

For more information about initial appliance deployment, network configuration and using the Web User Interface (WebUI), please also refer to the Administration Manual.

2. Loadbalancer.org Appliances Supported

All our products can be used for load balancing Metaswitch Virtual EAS SSS. For full specifications of available models please refer to https://www.loadbalancer.org/products/enterprise.

Some features may not be available or fully supported in all cloud platforms due to platform specific limitations. For more details, please refer to the "Main Differences to our Standard (Non-Cloud) Product" section in the appropriate cloud platform Quick Start Guide or check with Loadbalancer.org support.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.9.1 and later

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

3.2. Metaswitch Virtual EAS SSS

- Version 9.x and above

4. Metaswitch Virtual EAS SSS

The following Metaswitch related acronyms are used throughout this document. They are presented here in full for clairty.

- EAS – Enhanced Application Server
- SSS – Stackable Server Solution

This guide specifically describes configuring a virtual load balancer to be used with the deployment of Metaswitch Virtual EAS SSS. The principles and instructions presented here would also apply to a hardware EAS SSS implementation. Loadbalancer.org and Metaswitch have a long-standing partnership for the resale of a fully supported hardware EAS SSS implementation.

In a Metaswitch Virtual EAS SSS deployment, the EAS servers may also be connected to separate signalling and media networks. These additional networks carry SIP and RTP traffic respectively. It is not necessary to connect

the load balancer to them as **the load balancer is not responsible for load balancing SIP or RTP traffic**.

5. Sizing, Capacity, and Performance

For deployments up to 250,000 subscribers, your virtual host should be allocated a minimum of 8 vCPUs, 16 GB of RAM, and 8 GB of disk storage.

This specification will support the following bandwidth and connection thresholds:

- Internet → EAS bandwidth: 100 Mbit/s
- Internet → EAS packets/s: 70,000 pkts/s
- EAS → Internet bandwidth: 700 Mbit/s
- EAS → Internet packets/s: 55,000 pkts/s
- Concurrent connections: 380,000 connections

For larger deployments, your Metaswitch support representative will give you details of the expected load on your load balancers based on your predicted usage profile.

6. Using WAF Gateways

A service provided by a Metaswitch EAS SSS deployment, for example CommPortal access, can be protected with a WAF gateway service on a Loadbalancer.org appliance. This can be done for both virtual and hardware deployments.

A set of five custom WAF rules have been developed to protect a Metaswitch EAS deployment. These rules protect from a range of different attacks, including:

- denial-of-service attacks on login pages
- brute-force attacks to guess passwords
- attempts to gain access to accounts by trying the same common passwords many consecutive times

Instructions on how to deploy WAF gateways, as well as explanations of the custom WAF rules, form a separate deployment guide titled 'Loadbalancer.org WAF Gateway with Metaswitch EAS'.

7. Load Balancing Metaswitch Virtual EAS SSS

Note	It's highly recommended that you have a working Metaswitch Virtual EAS SSS environment first before implementing the load balancer.

7.1. Persistence (aka Server Affinity)

Some of the virtual services needed to load balance Metaswitch Virtual EAS SSS require source IP persistence.

The EAS servers include a file named pools.txt which describes the virtual services that need to be set up,

whether or not they require persistence, and if so what persistence timeout value should be used.

Full instructions on correctly configuring persistence settings can be found in sections Appliance Configuration for Metaswitch Virtual EAS SSS – Two Internal Networks (Scenario 1) and Appliance Configuration for Metaswitch Virtual EAS SSS – One Internal Network (Scenario 2).

7.2. Port Requirements

The ports that are load balanced vary from one EAS deployment to another.

The EAS servers include a file named pools.txt which describes the virtual services that need to be set up and which ports they should be listening on.

Full instructions on correctly setting up virtual services listening on the correct ports can be found in sections Appliance Configuration for Metaswitch Virtual EAS SSS – Two Internal Networks (Scenario 1) and Appliance Configuration for Metaswitch Virtual EAS SSS – One Internal Network (Scenario 2).

8. Network Configuration Options

There are two ways that Metaswitch Virtual EAS SSS can be deployed. The deployment type used determines how the load balancer must be configured.

It is essential to know which type of deployment is in place or being planned **before** attempting to set up a load balancer. If the deployment type is not clear, please contact Metaswitch support for further information.

8.1. Scenario 1 – Two Separate Internal Networks (Recommended)

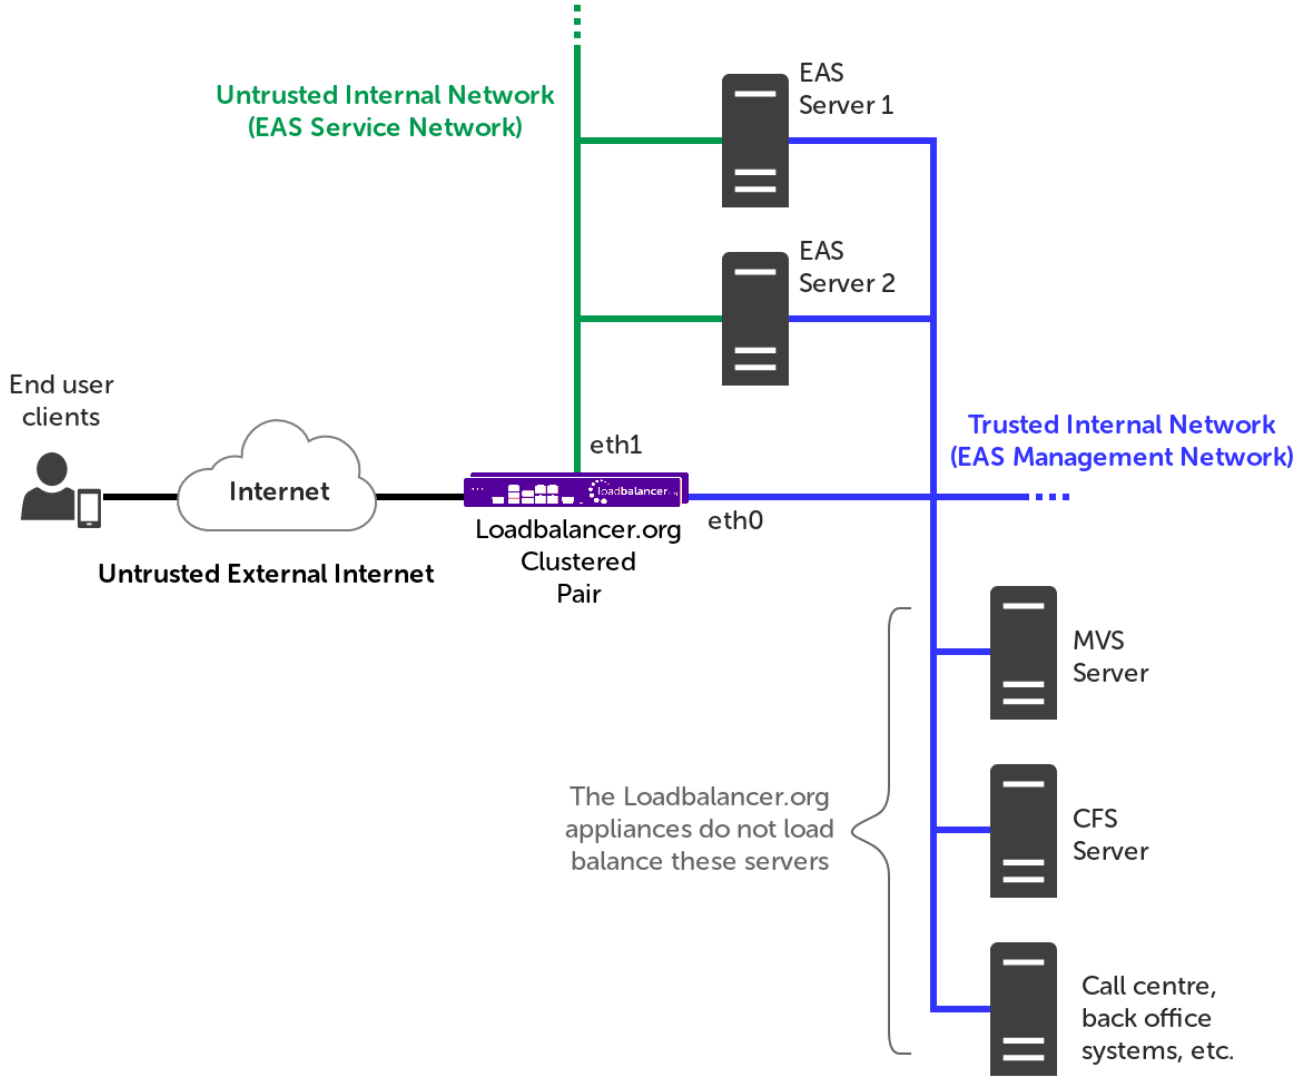

This is the recommended network configuration by Metaswitch. Two separate internal networks are used in this configuration:

- **Service network**: this handles traffic from untrusted external clients, which is load balanced using layer 4 NAT mode virtual services
- **Management network**: this handles traffic from trusted internal clients, which is load balanced using layer 7 SNAT mode virtual services

It can be the case that the load balancer needs to be connected to a third separate network, which is the **external network** where traffic from external clients goes to and from. It could also be the case that external traffic flows through a router sitting in either the management or service networks, which removes the need for this third network connection. If it is not clear how external traffic is routed in your deployment, please contact Metaswitch support for further information.

In this type of deployment, traffic flow through the load balancer looks like the following:

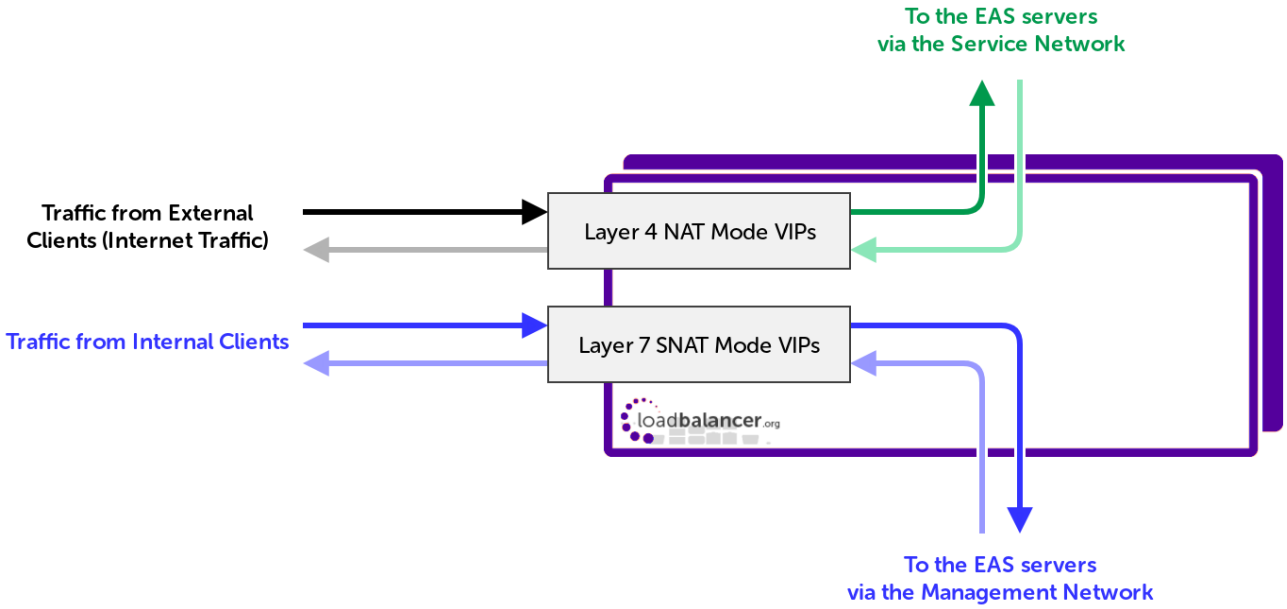

There are three main benefits to using two separate internal networks:

- Traffic from external clients, a significant proportion of incoming traffic, can be load balanced using layer 4 NAT mode. This is much faster and less intensive for the load balancer compared to load balancing at layer 7
- Untrusted external traffic is isolated from the trusted internal management network and traffic
- Source IP addresses are preserved for incoming requests from external clients, i.e. clients out on the public Internet. This makes identifying and blocking malicious clients easier

Note

It is possible to use two separate HA pairs of load balancers in this scenario.

In this case, the first pair of load balancers would connect the external clients to the service network, and would host layer 4 NAT mode VIPs. The second pair of load balancers would sit in the internal trusted management network, and would only load balance internal traffic using layer 7 VIPs.

In practice, the functionality of these two HA pairs of load balancers can be combined into a single pair, provided that they have 'arms' in each of the necessary networks (the service network, management network, and external network too if applicable).

This guide describes setting up a single HA pair of load balancers.

8.2. Scenario 2 – One Internal Network

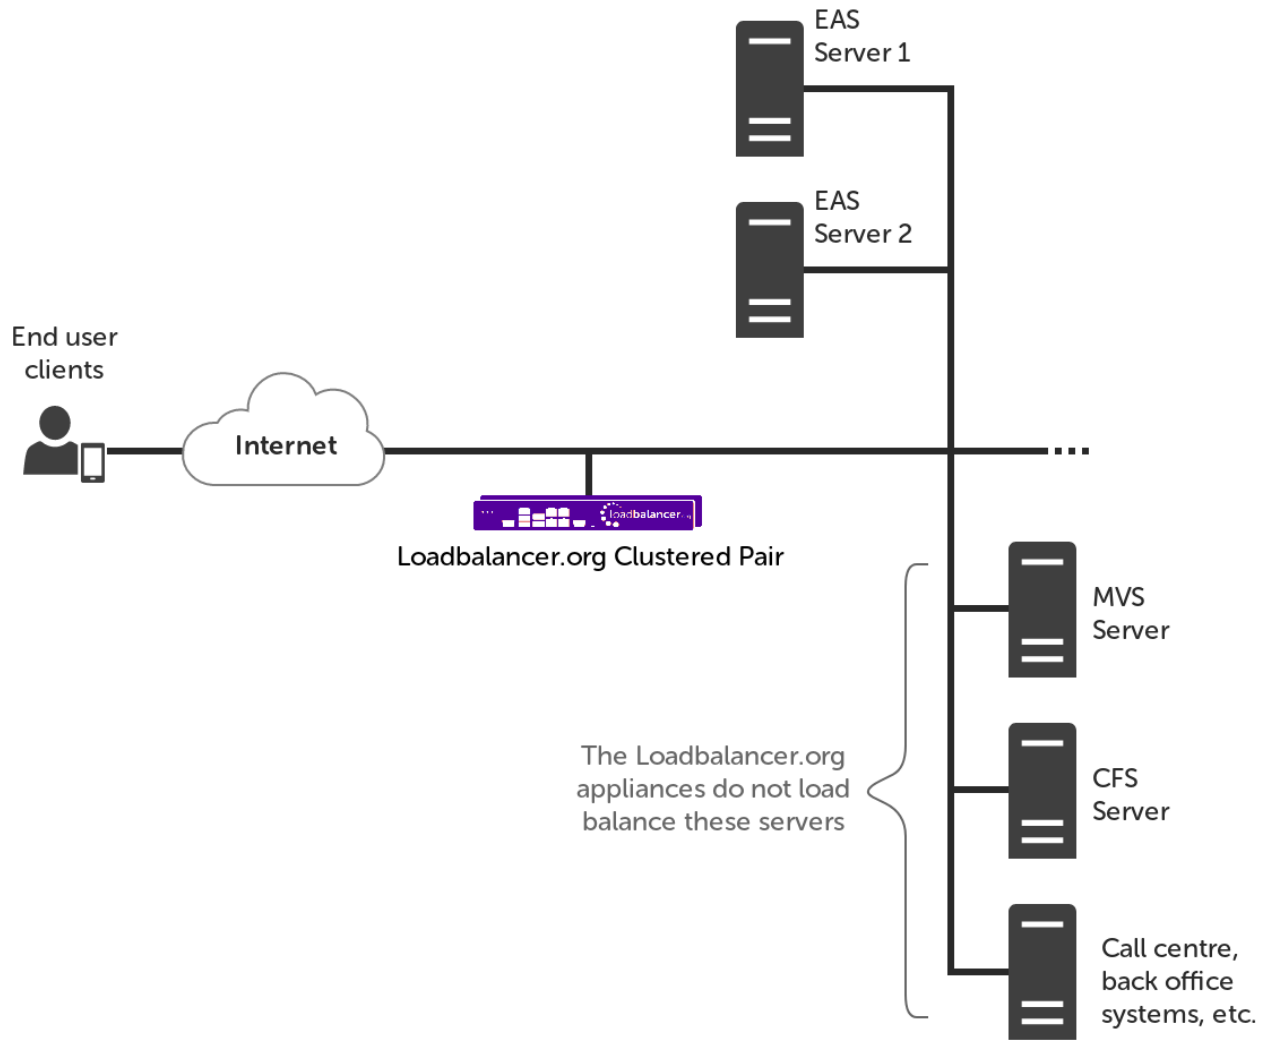

An EAS deployment can be built without a separate service network. In this case, external and internal traffic is all handled on a single internal network. All traffic must be load balanced using layer 7 SNAT mode virtual services.

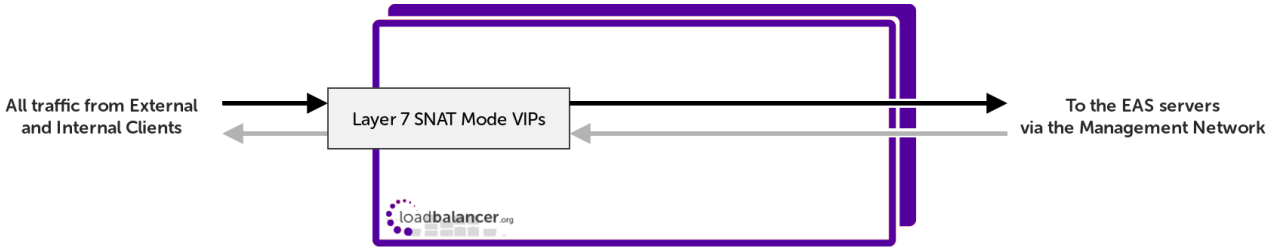

When load balancing at layer 7, performance is not as fast as the layer 4 option available in scenario 1. Layer 7 load balancing is also more intensive for the load balancer.

The load balancer acts as a full proxy in this setup, and as such load balancing is not source IP transparent. This means the EAS servers see all inbound traffic as originating from one of the load balancer's IP addresses. This makes it difficult to identify and block malicious requests from external clients on the public Internet, as their source IP addresses are obscured.

8.3. Virtual Service (VIP) Requirements

The number of virtual services required on the load balancer varies between Metaswitch Virtual EAS SSS

deployments.

The EAS writes a file named pools.txt which details every virtual service that needs to be configured. This file can be found on any EAS server, in the directory /home/defcraft/files.

Instructions on how to configure these virtual services can be found in the following sections, depending on which network configuration / scenario you are using:

Appliance Configuration for Metaswitch Virtual EAS SSS – Two Internal Networks (Scenario 1) Appliance Configuration for Metaswitch Virtual EAS SSS – One Internal Network (Scenario 2)

8.4. Load Balancing Methods

The load balancer can be deployed in one of 4 fundamental ways: ***Layer 4 DR mode***, ***Layer 4 NAT mode***, ***Layer 4 SNAT mode***, or ***Layer 7 SNAT mode***. For Metaswitch Virtual EAS SSS, layer 4 NAT mode and layer 7 SNAT mode virtual services are supported. Both of these supported load balancing methods are described below.

8.5. Layer 4 NAT Mode

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

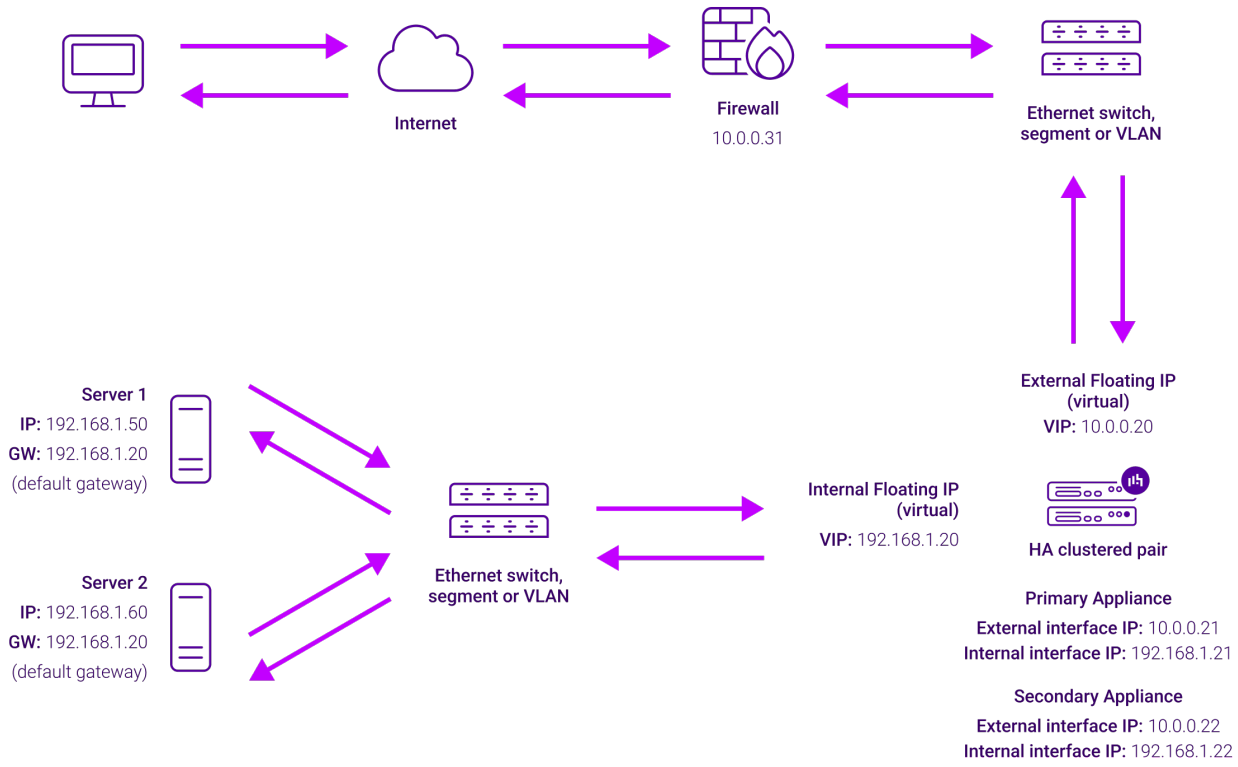

- The load balancer translates all requests from the Virtual Service to the Real Servers.
- NAT mode can be deployed in the following ways:
  - **Two-arm (using 2 Interfaces)** (as shown above) - Here, 2 subnets are used. The VIP is located in one subnet and the load balanced Real Servers are located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

interfaces, one in each subnet.

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

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

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

- Clients can be located in the same subnet as the VIP or any remote subnet provided they can route to the VIP.

- **One-arm (using 1 Interface)** - Here, the VIP is brought up in the same subnet as the Real Servers.

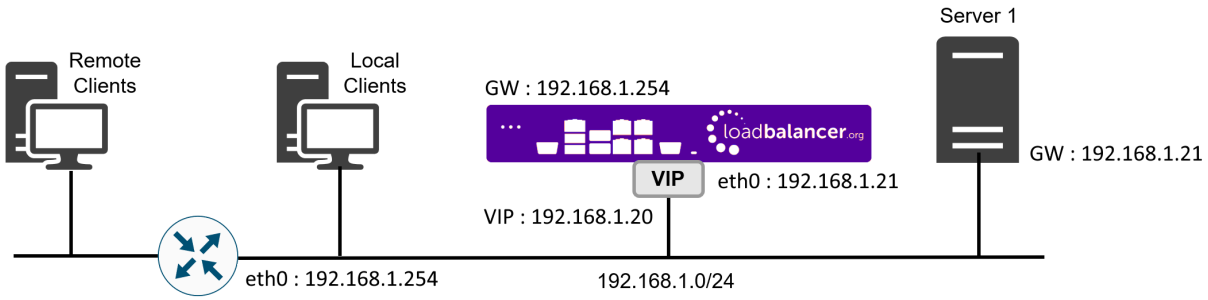

- To support remote clients, the default gateway on the Real Servers must be an IP address on the load balancer and routing on the load balancer must be configured so that return traffic is routed back via the router.

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

- To support local clients, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer. For more information please refer to One-Arm (Single Subnet) NAT Mode.

- If you want Real Servers to be accessible on their own IP address for non-load balanced services, e.g. RDP, you will need to setup individual SNAT and DNAT firewall script rules for each Real Server or add additional VIPs for this.

- Port translation is possible with Layer 4 NAT mode, e.g. VIP:80 → RIP:8080 is supported.
- NAT mode is transparent, i.e. the Real Server will see the source IP address of the client.

NAT Mode Packet re-Writing

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

The following table shows an example NAT mode setup:

Protocol	VIP	Port	RIP	Port
TCP	10.0.0.20	80	192.168.1.50	80

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

Packet rewriting works as follows:

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

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

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

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

3) Replies return to the load balancer as:

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

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

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

8.6. 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.

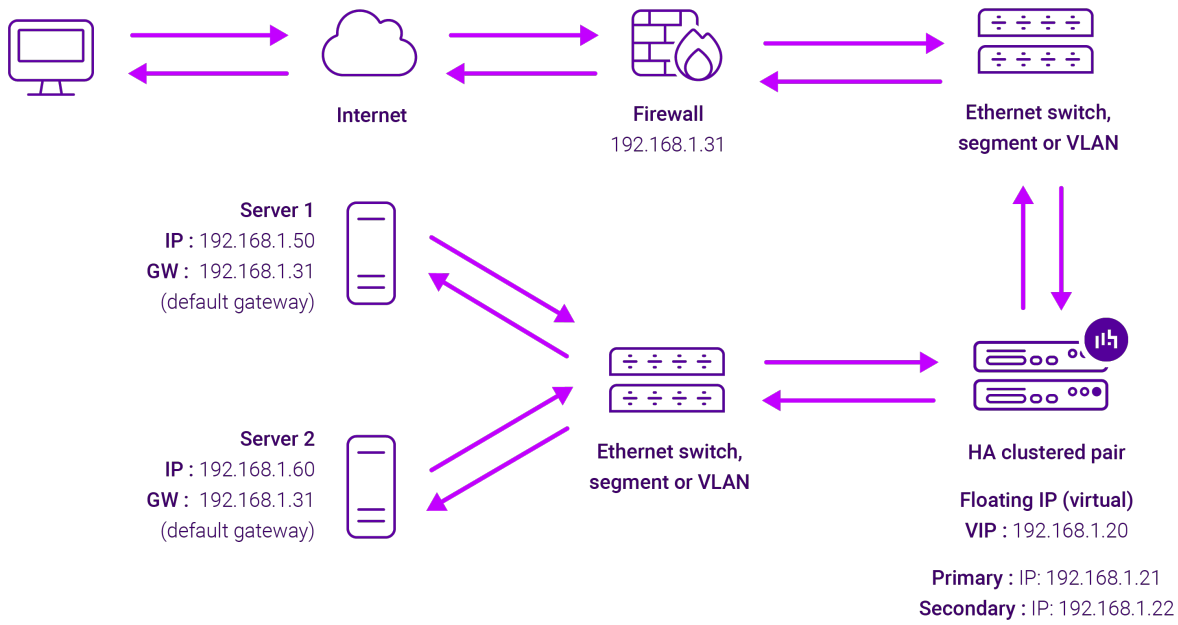

- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's 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.

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:

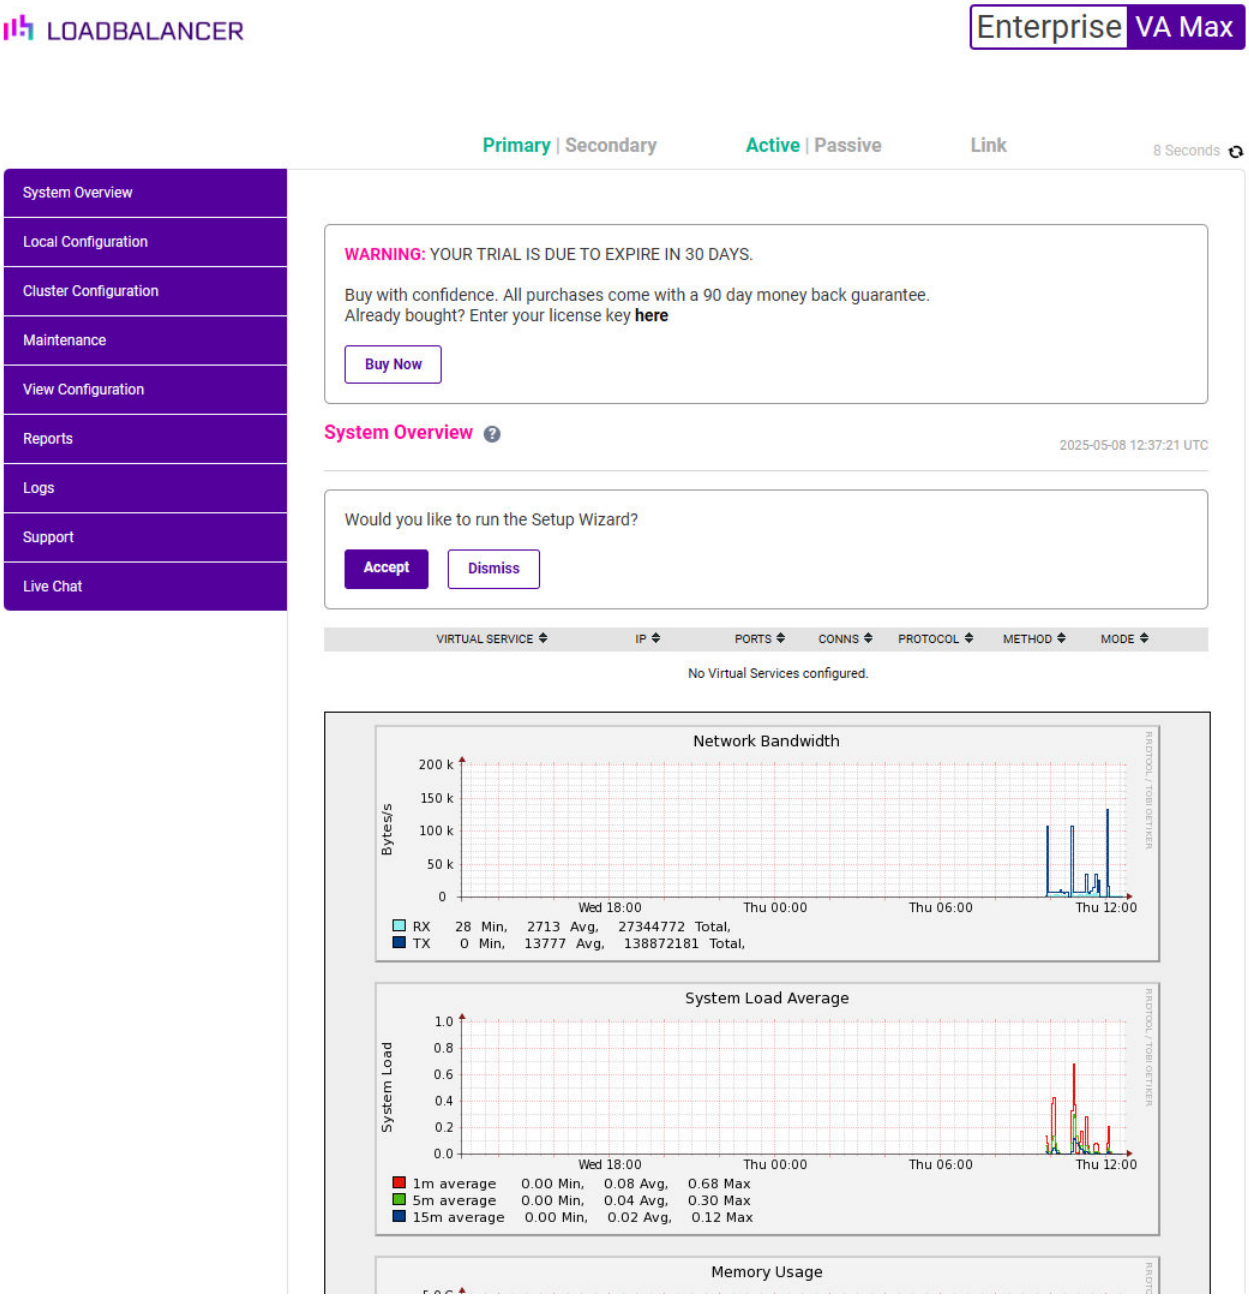

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

9.3.1. 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

9.4. Appliance Software Update

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

> **Note** For full details, please refer to Appliance Software Update in the Administration Manual.

> **Note** Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

9.4.1. Online Update

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

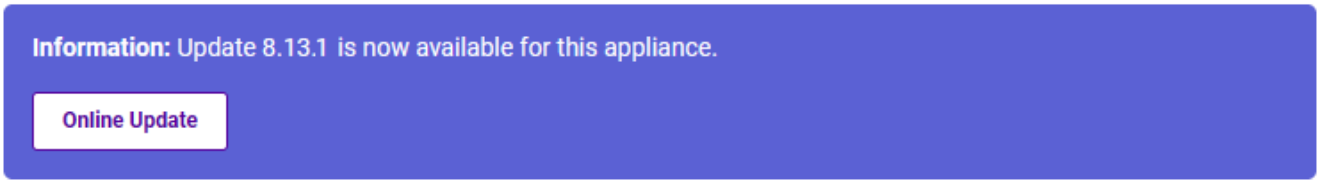

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

> **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:

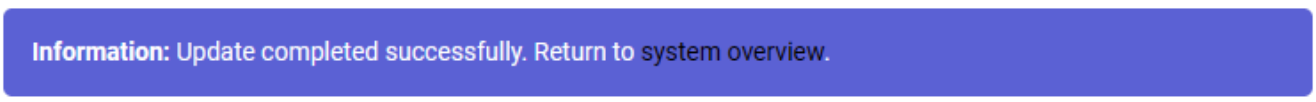

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

9.4.2. Offline Update

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

To check for the latest version, please refer to our product roadmap page available here. To obtain the latest offline update files contact support@loadbalancer.org.

To perform an offline update:

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

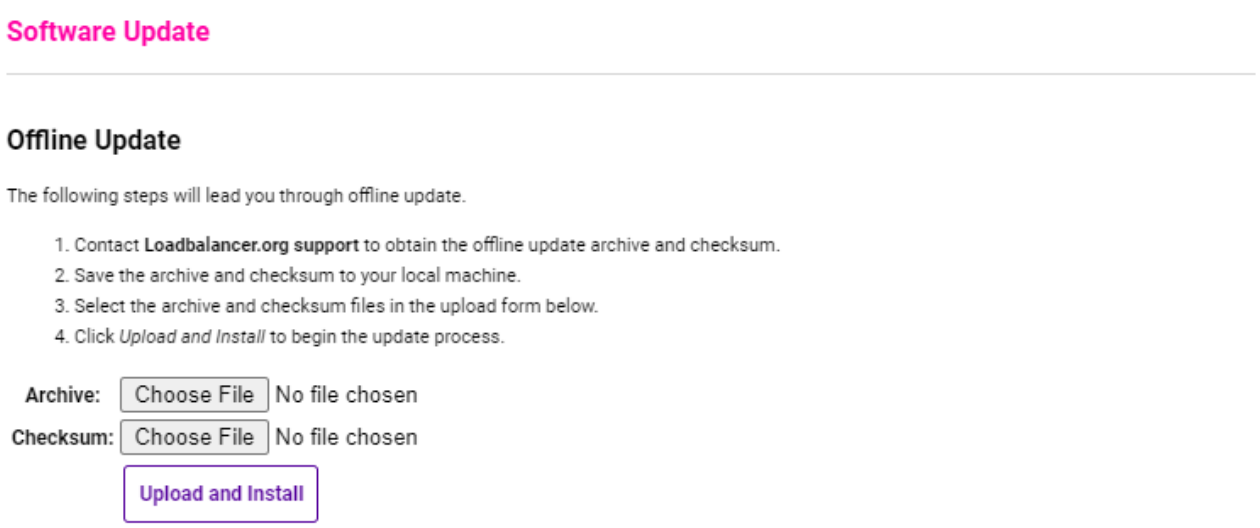

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

9.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22 *	SSH
TCP & UDP	53 *	DNS / GSLB
TCP & UDP	123	NTP
TCP & UDP	161 *	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000 *	Gateway service (Centralized/Portal Management)
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

> **Note** The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to Service Socket 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 Configuring HA - Adding a Secondary Appliance.

10. Appliance Configuration for Metaswitch Virtual EAS SSS – Two Internal Networks (Scenario 1)

Two separate internal networks are used in this scenario.

During initial setup, the load balancer should have been assigned an IP address on the **trusted internal management network**. This IP address is assigned to the **eth0** network interface. With this IP address, the load balancer has an 'arm'/connection set up to the trusted internal management network.

The load balancer must now be given an arm in the **internal service network**, as well as an arm in the external network if that is a separate network.

10.1. Connecting the Load Balancer to the Service Network

The load balancer needs to be assigned an IP address in the service network to use that network. This IP address should be assigned to a separate network interface.

To assign an IP address from the WebUI:

1. Go to *Local Configuration > Network Interface Configuration*.
2. Add the required IP address and subnet mask next to the appropriate NIC under *IP Address Assignment*.
3. Press the **Configure Interfaces** button.

In the example presented here, the IP address 192.168.85.240/18, which sits in the service network, is being added to interface **eth1**:

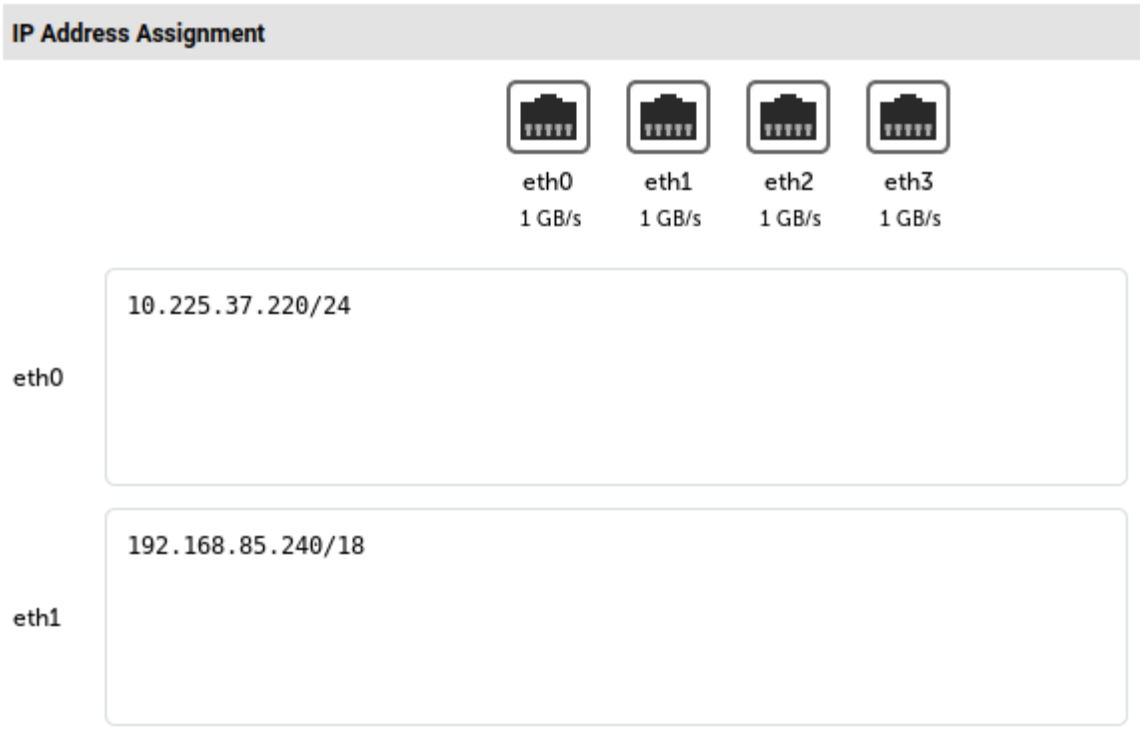

> **Note** If the load balancer is being set up as part of a highly available pair, the Secondary load balancer must also be assigned its own IP address in the service network by following the same steps.

10.2. Connecting the Load Balancer to the External Network

The load balancer must have an 'arm' in the external network through which client traffic enters the EAS deployment. How this is accomplished will depend on how the network is set up.

For some deployments, external client traffic enters through a firewall/router that sits in the service network. In this situation, **if the external facing firewall can be reached** from the load balancer's IP address in either the management network or the service network then **no additional configuration is necessary**.

If external client traffic enters the deployment from a third distinct network then the load balancer will need an arm / IP address in that network. To do this, follow the same steps used in the previous section Connecting the Load Balancer to the Service Network. A third NIC can be used for this, e.g. eth2.

10.3. Configuring the Virtual Services

A list of every virtual service that needs to be configured on the load balancer can be found on any of the EAS servers (the list files are identical). The EAS writes this list to the file pools.txt. This file can be found on any EAS server in the directory /home/defcraft/files.

10.3.1. Configuring the Layer 7 Management VIPs

The first section of the pools.txt file is headed "Management". It describes the VIPs that need to be set up on the internal management network. These are all layer 7 SNAT mode VIPs. This section of the file looks like this:

```
POOL POOL IP PROTO PORT STICKY SERVER IPS
# MANAGEMENT
```

```
HTTP-80 10.225.37.227 tcp 80 2100 10.225.37.231 10.225.37.116
HTTPS-443 10.225.37.227 tcp 443 2100 10.225.37.231 10.225.37.116
HTTPS-10000 10.225.37.227 tcp 10000 2100 10.225.37.231 10.225.37.116
IMAP-143 10.225.37.223 tcp 143 0 10.225.37.231 10.225.37.116
MVWEBPROV-8087 10.225.37.224 tcp 8087 86400 10.225.37.231 10.225.37.116
```

Each line after the "Management" heading describes an individual VIP to be set up. These lines are read as follows:

- **Pool**: a descriptive name for the VIP, which could be used as the VIP label on the load balancer
- **Pool IP**: the IP address that must be used for the virtual service
- **Proto**: the protocol for the virtual service, which should always be listed as TCP
- **Port**: the port that must be used for the virtual service
- Sticky: the source IP persistence setting to be used. A value of 0 means persistence must be disabled. A non-zero value means source IP persistence must be enabled and have a timeout equal to that number of seconds
- Server IPs: the RIPs (IP addresses of each real server) that must be added to the VIP

Every layer 7 virtual service described in the file needs to be set up. Instructions on how to set up a layer 7 VIP are presented below.

Configuring a Layer 7 VIP

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as desired, e.g. **MGMT_HTTPS-443**.
3. Set the *Virtual Service IP Address* field to the "Pool IP" value, e.g. **10.225.37.227**.
4. Set the *Ports* field to the "Port" value, e.g. **443**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

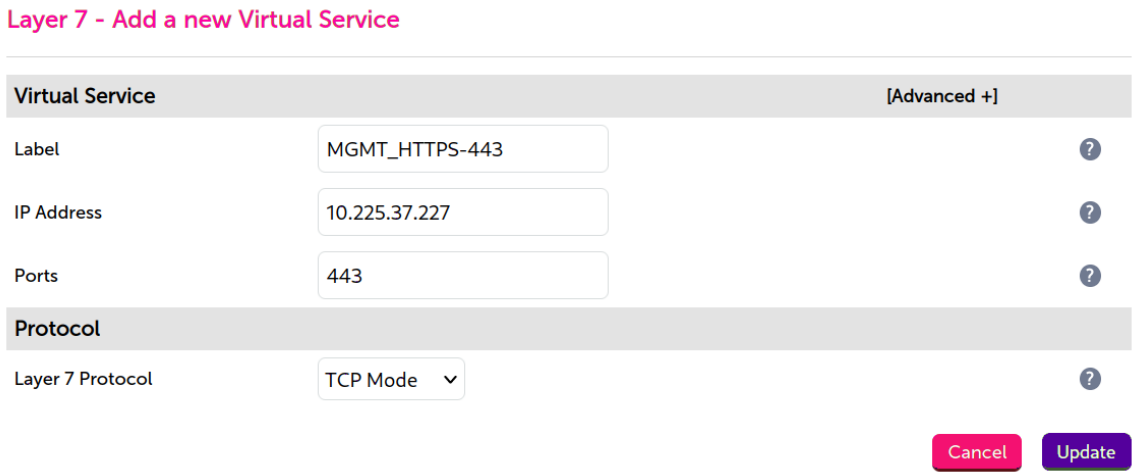

7. Click **Modify** next to the newly created VIP.
8. Set *Balance Mode* to **Weighted Round Robin**.
9. Set *Persistence Mode* as required:
    - If the "Sticky" column has a value of 0 then *Persistence Mode* must be set to **None**
    - If the "Sticky" column has a non-zero value then under *Persistence* click **Advanced** to show more options. *Persistence Mode* must be set to **Source IP** and the *Persistence Timeout* must be set equal to the "Sticky" value divided by 60. The division is necessary because the layer 7 persistence timeout units are *minutes* as opposed to the 'Sticky' units which are *seconds*
10. Set *Health Checks* to **Connect to port**.
11. Set *Check Port* to the "Port" value, e.g. **443**.
12. Under the *Other* section, click the **Advanced** button.
13. Check the **Timeout** checkbox.
14. Set the *Client Timeout* value to **900000**.
15. Set the *Real Server Timeout* value to **901000**.
16. Click **Update**.

Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Define the *Label* for the real server as required, e.g. **EAS_SSS-1**.
3. Set the *Real Server IP Address* field to the first "Server IPs" value, e.g. **10.225.37.231**.
4. Set the *Real Server Port* field to the "Port" value, e.g. **443**.
5. Click **Update**.
6. Repeat these steps until all the listed EAS servers have been added.

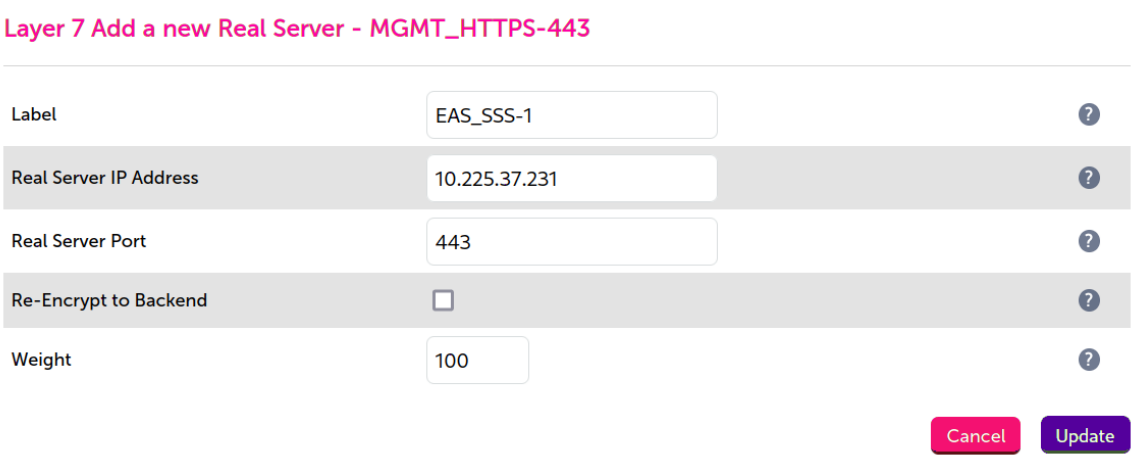

Changing Layer 7 Global Settings

Once all of the individual VIPs have been set up, one of the layer 7 global settings needs to be changed.

1. From the WebUI, go to *Cluster Configuration > Layer 7 – Advanced Configuration*.
2. Set *Interval* to **3000**.
3. Click **Update**.

Finalizing the Layer 7 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**.

10.3.2. Configuring the Layer 4 Service VIPs

The second section of the pools.txt file is headed "Service". It describes the VIPs that need to be set up on the internal service network. These are all layer 4 NAT mode VIPs. This section of the file looks like this:

```
# SERVICE
HTTP-80 - tcp 80 2100 192.168.85.61 192.168.85.5
HTTPS-443 - tcp 443 2100 192.168.85.61 192.168.85.5
HTTPS-10000 - tcp 10000 2100 192.168.85.61 192.168.85.5
IMAP-143 - tcp 143 0 192.168.85.61 192.168.85.5
```

Each line after the "Service" heading describes an individual VIP to be set up. These lines are read in the same way as the "Management" VIPs, and a list of descriptions of each column can be found in Configuring the Layer 7 Management VIPs.

Every layer 4 virtual service described in the file needs to be set up. Instructions on how to set up a layer 4 VIP are presented below.

Configuring a Layer 4 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 desired, e.g. **SERV_IMAP-143**.
3. Set the *Virtual Service IP Address* field to an IP address of your choice, bearing in mind the following two conditions for NAT mode to work:
    - The VIP address must be accessible from the external network
    - The VIP address must be in a different subnet to the real server IP addresses

In this example the VIP address 192.168.85.150 is used, as this IP address is not in the same subnet as the example real servers which are sitting in the 192.168.85.0/25 subnet

1. Set the *Ports* field to the "Port" value, e.g. **143**.
2. Set the *Protocol* to **TCP**.

3. Set the *Forwarding Method* to **NAT**.
4. Click **Update** to create the virtual service.

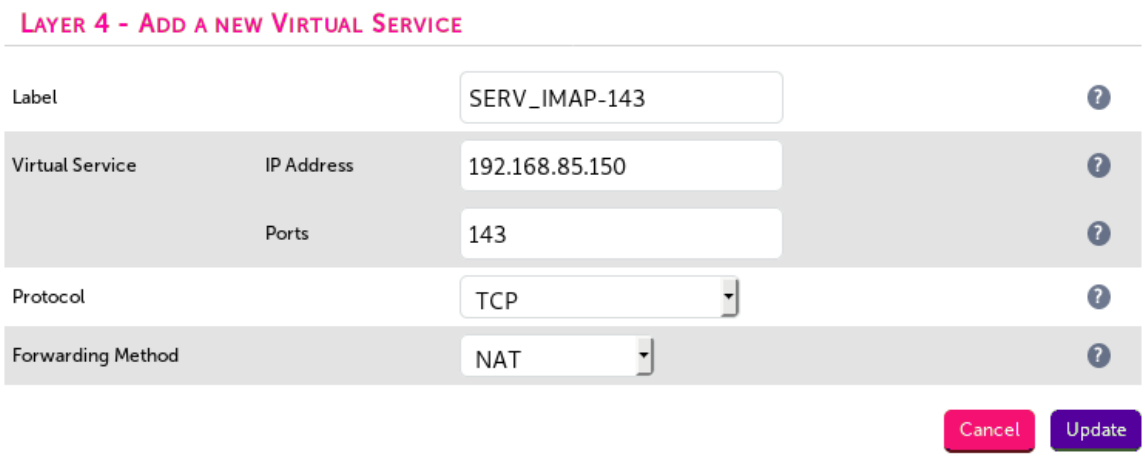

5. Click **Modify** next to the newly created VIP.
6. Set *Balance Mode* to **Weighted Round Robin**.
7. Set *Persistent* as required:
   - If the "Sticky" column has a value of 0 then the *Persistent* checkbox must not be checked
   - If the "Sticky" column has a non-zero value then the *Persistent* checkbox must be checked and the *Persistence Timeout* must be set to the "Sticky" value (the layer 4 persistence timeout units are seconds, so no division is necessary as it is layer 7 services)

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. **EAS_SSS-1**.
3. Set the *Real Server IP Address* field to the first "Server IPs" value, e.g. **192.168.85.61**.
4. Click **Update**.
5. Repeat these steps until all the listed EAS servers have been added.

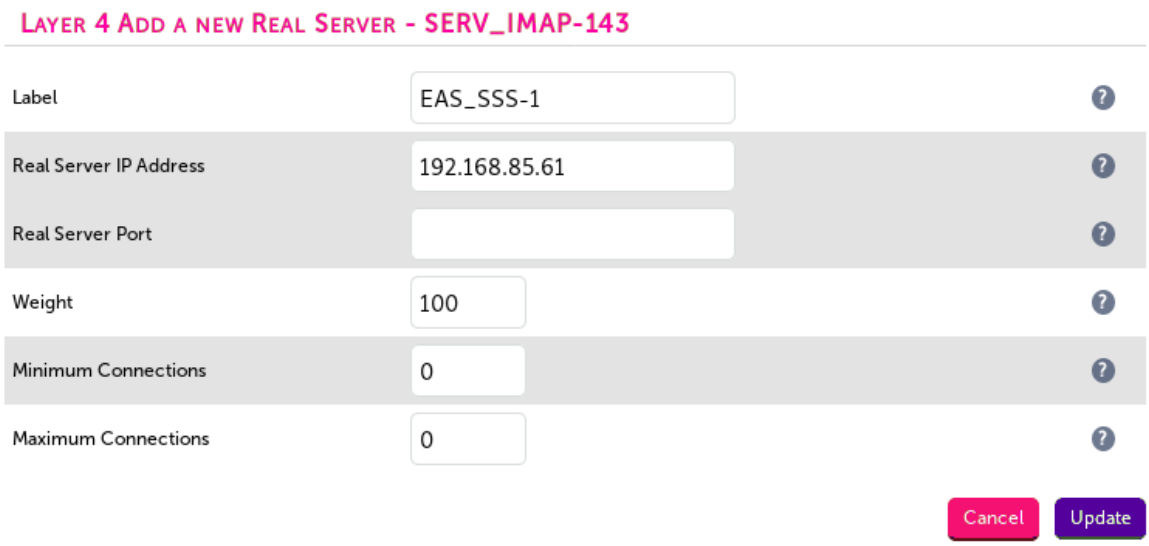

Changing Layer 4 Global Settings

Once all of the individual VIPs have been set up, some of the layer 4 global settings regarding health checking need to be changed.

1. From the WebUI, go to *Cluster Configuration > Layer 4 – Advanced Configuration*.
2. Set *Check Interval* to **3**.
3. Set *Check Timeout* to **3**.
4. Set *Failure Count* to **2**.
5. Click **Update**.

Creating a Floating IP Address on the Service Network

If deploying an HA pair of load balancers, a floating IP address sitting on the service network **must** be created on the load balancers.

Because of how layer 4 NAT mode works, the EAS servers must have their gateway for the service network set to be an IP address that is owned by the load balancer. This causes the EAS servers to reply to incoming service network traffic via the load balancer. The gateway is configured on the EAS servers, and the IP address to be used for this must be known prior to EAS commissioning. Instructions for confirming the gateways configured on an EAS server are presented in Confirming the Gateway Settings on the EAS Servers.

For a pair of load balancers, the gateway IP address must be a floating IP address. This means that during a fail over from one load balancer to the other this IP address is picked up by the other appliance, enabling traffic to continue to route from the real servers on toward the external clients.

To add the floating IP address from the WebUI on the Primary load balancer:

1. Navigate to *Cluster Configuration > Floating IPs*.
2. Specify the new floating IP addresses.
3. Click the **Add Floating IP** button.

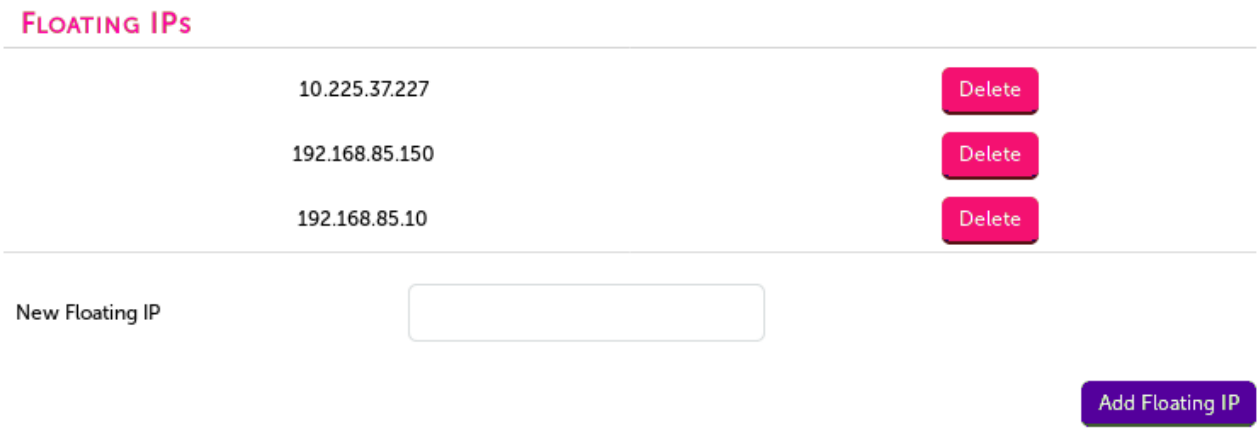

11. Appliance Configuration for Metaswitch Virtual EAS SSS – One Internal Network (Scenario 2)

A single internal network is used in this scenario. All traffic is handled by a single set of layer 7 virtual services, traffic from both from internal clients and external clients out on the Internet.

11.1. Configuring the Virtual Services

A list of every virtual service that needs to be configured on the load balancer can be found on any of the EAS servers (the list files are identical). The EAS writes this list to the file pools.txt. This file can be found on any EAS server in the directory /home/defcraft/files.

11.1.1. Configuring the Layer 7 Management VIPs

If the EAS servers are configured for use in a single internal network environment, the pools.txt file should contain a single section headed "Management". It describes the layer 7 VIPs that need to be set up. These are all layer 7 SNAT mode VIPs. The file looks like this:

```
POOL POOL IP PROTO PORT STICKY SERVER IPS
# MANAGEMENT
HTTP-80 10.225.37.227 tcp 80 2100 10.225.37.231 10.225.37.116
HTTPS-443 10.225.37.227 tcp 443 2100 10.225.37.231 10.225.37.116
HTTPS-10000 10.225.37.227 tcp 10000 2100 10.225.37.231 10.225.37.116
IMAP-143 10.225.37.223 tcp 143 0 10.225.37.231 10.225.37.116
MVWEBPROV-8087 10.225.37.224 tcp 8087 86400 10.225.37.231 10.225.37.116
```

Each line after the "Management" heading describes an individual VIP to be set up. These lines are read as follows:

- **Pool**: a descriptive name for the VIP, which could be used as the VIP label on the load balancer
- **Pool IP**: the IP address that must be used for the virtual service
- **Proto**: the protocol for the virtual service, which should always be listed as TCP
- **Port**: the port that must be used for the virtual service

- **Sticky**: the source IP persistence setting to be used. A value of 0 means persistence must be disabled. A non-zero value means source IP persistence must be enabled and have a timeout equal to that number of seconds
- **Server IPs**: the RIPs (IP addresses of each real server) that must be added to the VIP

Every layer 7 virtual service described in the file needs to be set up. Instructions on how to set up a layer 7 VIP are presented below.

Configuring a Layer 7 VIP

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as desired, e.g. **MGMT_HTTPS-443**.
3. Set the *Virtual Service IP Address* field to the "Pool IP" value, e.g. **10.225.37.227**.
4. Set the *Ports* field to the "Port" value, e.g. **443**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

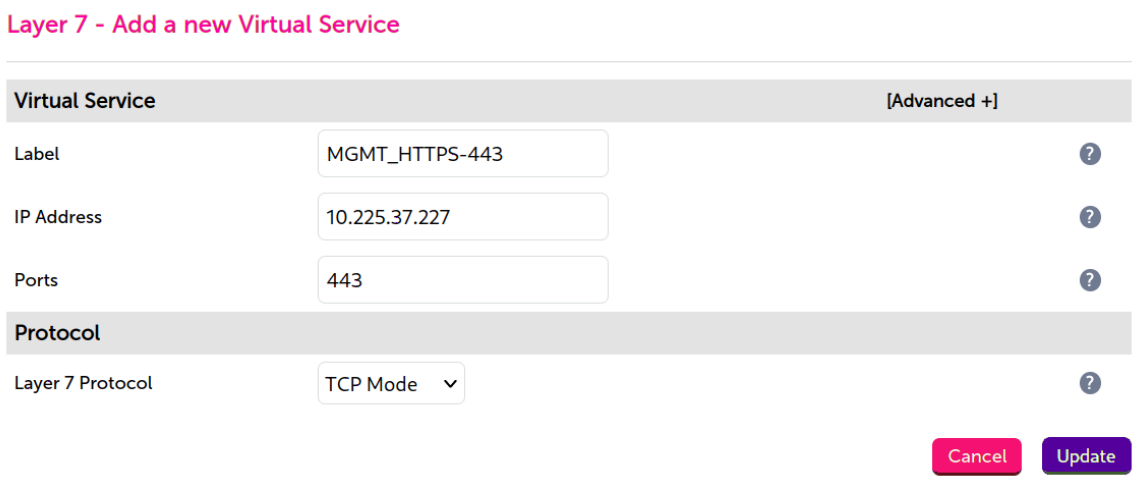

7. Click **Modify** next to the newly created VIP.
8. Set *Balance Mode* to **Weighted Round Robin**.
9. Set *Persistence Mode* as required:
   - If the "Sticky" column has a value of 0 then *Persistence Mode* must be set to **None**
   - If the "Sticky" column has a non-zero value then under *Persistence* click **Advanced** to show more options. *Persistence Mode* must be set to **Source IP** and the *Persistence Timeout* must be set equal to the "Sticky" value divided by 60. The division is necessary because the layer 7 persistence timeout units are *minutes* as opposed to the 'Sticky' units which are *seconds*
10. Set *Health Checks* to **Connect to port**.
11. Set *Check Port* to the "Port" value, e.g. **443**.
12. Under the *Other* section, click the **Advanced** button.

13. Check the **Timeout** checkbox.
14. Set the *Client Timeout* value to **900000**.
15. Set the *Real Server Timeout* value to **901000**.
16. Click **Update**.

Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Define the *Label* for the real server as required, e.g. **EAS_SSS-1**.
3. Set the *Real Server IP Address* field to the first "Server IPs" value, e.g. **10.225.37.231**.
4. Set the *Real Server Port* field to the "Port" value, e.g. **443**.
5. Click **Update**.
6. Repeat these steps until all the listed EAS servers have been added.

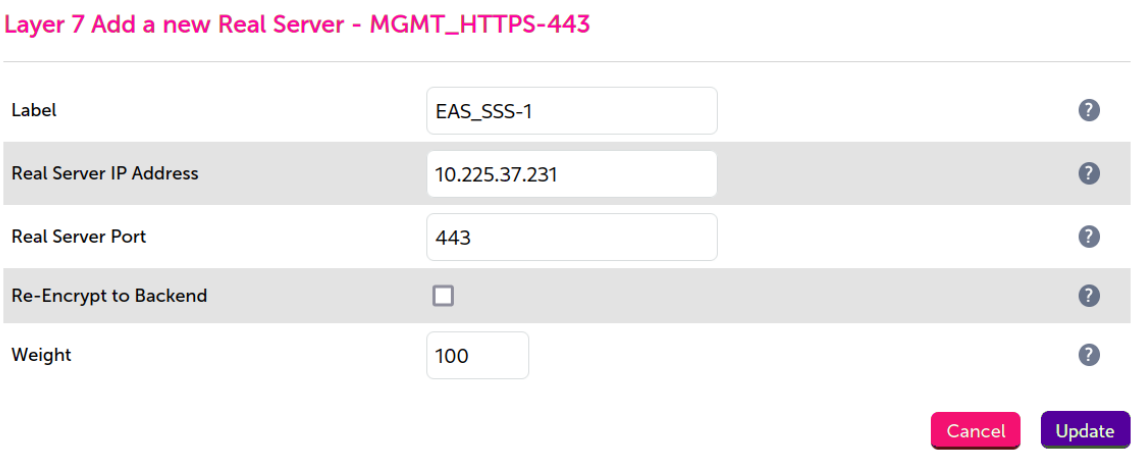

Changing Layer 7 Global Settings

Once all of the individual VIPs have been set up, one of the layer 7 global settings needs to be changed.

1. From the WebUI, go to *Cluster Configuration > Layer 7 – Advanced Configuration*.
2. Set *Interval* to **3000**.
3. Click **Update**.

Finalizing the Layer 7 Configuration

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

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

12. Testing & Verification

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

12.1. Metaswitch Specific Tests

Metaswitch recommend using some specific tests which make use of the EAS Craft menu and the CommPortal web portal.

12.1.1. Pool Configuration Test

Metaswitch provide a Python test script which reads the pools.txt file and checks that each virtual service has been created correctly and is accepting connections. This script can be run on any Linux or Windows machine that has Python 2.7 installed and can access the appropriate networks for the deployment.

The following points should be considered before running the test:

- At least one EAS server must be running for a successful test
- Depending on the network configuration, the script may need to be run separately from two different machines to test both the internal 'management network' services and the external facing 'service network' services
- If 'service network' services have been set up, the script will ask for the IP address and port that has been assigned for each virtual service, as these are set by the user and are not pre-defined
- The test script must not be run from an EAS server, as this will return false positives for all tests

Running the Test

1. Copy the Python script to an appropriate test machine.
2. Put the script in the same directory as the pools.txt file.
3. Run the test by executing the following, either from a Linux shell or the Windows command prompt:

```
python testpools.py
```

4. Follow the script's prompts and make sure that both the management and the service network virtual services are tested if testing a deployment that uses two internal networks (otherwise only the management services will have been set up and these should be tested).
5. Address any reported errors by checking the configuration of any problematic virtual services on the load balancer.

12.1.2. Health Checking Test

This test is disruptive to end users and should not be run on a live production system.

This test verifies that the load balancer's real server health checking is correctly configured. When an EAS server goes offline the load balancer should detect this change.

1. Ensure that all EAS servers are online.
2. Open the ***System Overview*** page of the load balancer's WebUI.
3. Check that all virtual services and real servers are displayed as online/green.
4. Use the EAS Craft interface to stop the software services on one or more real servers, by navigating to ***Main > Manage servers and passwords > Stop a server***.
5. Check the load balancer's ***System Overview*** page again, and verify that after a few seconds the appropriate EAS servers are now displayed as offline/red.
6. Use the EAS Craft interface to start all servers, by navigating to ***Main > Manage servers and passwords > Start ALL servers***.
7. Check the load balancer's ***System Overview*** page again, and verify that after a few seconds all virtual services and real servers are now displayed as online/green.

12.1.3. Persistence Test

This test is disruptive to end users and should not be run on a live production system.

This test verifies that persistence is configured correctly. For virtual services with persistence enabled, an end user should be load balanced to the same EAS server until their persistence session/entry times out.

Metaswitch note that this test doesn't have to be conducted for every virtual service with persistence enabled (i.e. services with non-zero "STICKY" values in pools.txt), and that provided the test works for one virtual service, and that other 'sticky' services are configured in the same way, this is sufficient.

1. Use the EAS Craft interface to stop approximately half of the EAS servers, **making sure that at least one is left running**. Do this by navigating to ***Main > Manage servers and passwords > Stop a server***.
2. Use a browser to access the CommPortal web interface via one of the configured virtual services (one of HTTP-80, HTTPS-443, HTTPS-100XX).
3. Log in using a test account.
4. Once successfully logged in, start up all servers by navigating to ***Main > Manage servers and passwords > Start ALL servers***.
5. Return to the CommPortal session, click on some of the tabs, and press CtrlF5 several times to force reload the page. The log in session should persist without issue.
6. Use the EAS Craft interface to stop all of the EAS servers that were left running at the start of this test, by navigating to ***Main > Manage servers and passwords > Stop a server***.
7. Return to the CommPortal session and press CtrlF5 to force reload the page. The test user should now be logged out, and the login screen should be displayed again.
8. Log back in using the test account, and then press CtrlF5 several times to force reload the page. The log in session should persist without issue.
9. Start up all servers again by navigating to ***Main > Manage servers and passwords > Start ALL servers***.
10. Return to the CommPortal session and press CtrlF5 several times to force reload the page. The log in session should persist without issue. The test is now complete.

12.1.4. Fail Over / High Availability Test

This test is disruptive to end users and should not be run on a live production system.

If using a highly available pair of load balancers, the fail over functionality between them can be tested.

1. Log into the CommPortal web interface.
2. Trigger a fail over from your active load balancer to your passive load balancer. You could force this by powering off the active load balancer.
3. Press CtrlF5 in browser to force refresh the CommPortal page.
4. Once a successful fail over has taken place, the passive load balancer will become active and will start serving traffic. The browser should show the CommPortal again. Note that a new log in may need to be performed following a fail over.

12.2. Useful Load Balancer Based Checks

12.2.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the EAS servers) 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 EAS servers are healthy and available to accept connections.

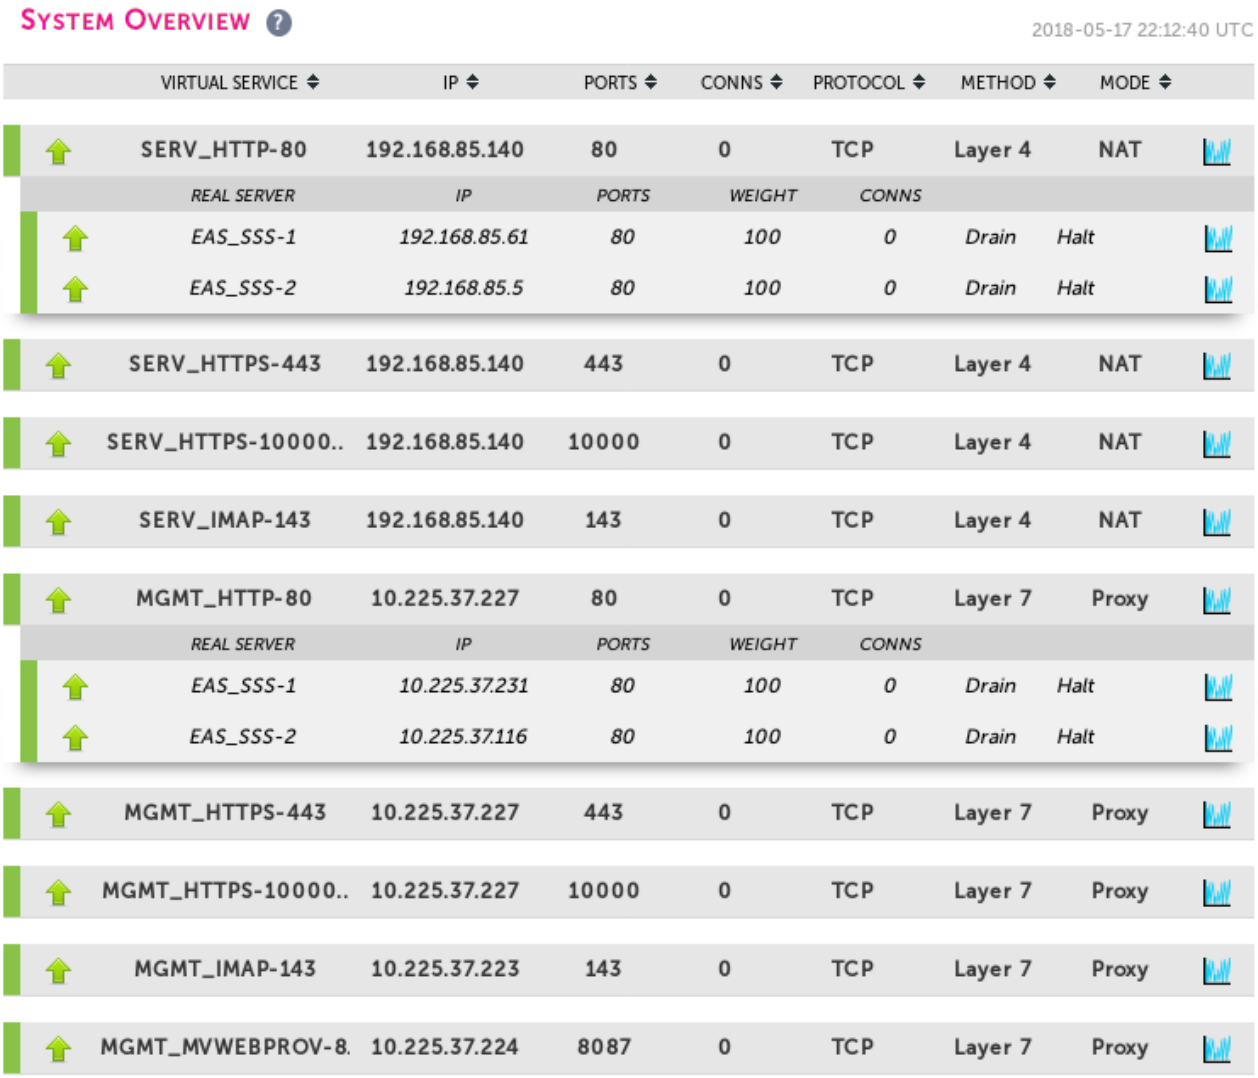

13. Technical Support

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. Confirming the Gateway Settings on the EAS Servers

It is possible to confirm what network gateways are defined on an EAS server. **If a gateway needs to be changed** to enable load balancing to work correctly, please contact Metaswitch support for assistance with this.

Confirming the gateway settings is done using the EAS Craft interface like so:

1. Sign into the EAS Craft interface.
2. Select option 1, *Configuration*.

```
Craft on svr01 (lbsdc)
Version V9.4.20-02
WARNING: svr01 is unlicensed
[Main]
Main menu

ENTER  Refresh
0)   < Log off the craft menu
1)   > Configuration
2)   > Manage software version
3)   > Manage servers and passwords
4)   > Diagnostics and recovery
5)   > Web configuration
6)   > Administration
8)   > TUI Customizations




Select an item: 1
_
```

3. Select option 2, *Network parameters*.

```
Craft on svr01 (lbsdc)
Version V9.4.20-02
WARNING: svr01 is unlicensed
[Main->Config]
Configuration

ENTER  Refresh
0)   < Back to previous menu
1)   > EAS parameters
2)   > Network parameters
3)   > Configure EAS services
4)   > Configure EAS servers
5)   > Configure 3rd-party hardware
6)   > Service Federation configuration
7)   > Manage TUI language packs
8)   > Licensing




Select an item: 2
```

_

4. Select option 6, *Manage subnets*.

```
Craft on svr01 (lbsdc)
Version V9.4.20-02
WARNING: svr01 is unlicensed
[Main->Config->Network]
Network parameters

ENTER  Refresh
0)   < Back to previous menu
1)   > SIP Settings
2)   > NTP, DNS, and rsyslog servers
3)   > Trusted Networks
4)   > MetaView Network Config
5)   > Server IP addresses and UDP multicast configuration
6)   > Manage subnets
7)   > External Servers




Select an item: 6

_
```

5. Select option 1, *List all subnets*.

```
Craft on svr01 (lbsdc)
Version V9.4.20-02
WARNING: svr01 is unlicensed
[Main->Config->Network->Subnet]
Manage subnets

ENTER  Refresh
0)   < Back to previous menu
1)     List all subnets









Select an item: 1

_
```

6. The different gateways configured on the server are listed. In the example presented here, the IP address 172.60.5.60 (highlighted) is used for the gateway on the service network. In the example deployment in question, a floating IP address of 172.60.5.60 is configured on the load balancer pair to enable traffic to flow correctly

```
Subnet Name                Protocol Subnet Gateway/Subnet Mask
-------------------------------------------------------------------------------
management                 IPv4     10.60.5.1/255.255.255.0
                           IPv6     not configured
-------------------------------------------------------------------------------
service                    IPv4     172.60.5.60/255.255.255.0
                           IPv6     not configured

Command succeeded.  Press a key to continue
```

15.2. Full Configuration Backup

When using a third party load balancer in a Metaswitch Virtual EAS SSS deployment, taking backups of the load balancer is an exercise that is left to the user.

Having a full configuration backup is very useful. In the unlikely event of a catastrophic failure, a fresh load balancer could be quickly deployed and put into production. In another scenario, if the load balancer's configuration were altered and broken in the future then the backed up working configuration could be imported and restored.

15.3. Taking a Backup

Once a load balancer has been fully deployed, configured, and is working as intended, it is very easy to take a full configuration backup.

1. From the WebUI, navigate to *Support > Technical Support Download*.
2. Click the **Generate Archive** button.
3. Wait for the archive to be generated.
4. Use the download link to save the archive to a safe location.

15.4. Restoring From a Backup

A support download archive contains an XML file (lb_config.xml) which describes the appliance's setup. This XML file can be imported into a load balancer to restore the backed up configuration.

1. From the WebUI, navigate to *Maintenance > Backup & Restore > Restore Tab*.
2. Use the **Upload XML file & Restore** function.
3. Wait for the success message that the configuration has been fully restored.

Note that a support download is also an invaluable tool for diagnosing problems with a load balancer, as it contains configuration and log files.

15.5. Performing Updates With Minimal Downtime

With a highly available pair of load balancers, it is possible to perform updates to the load balancers with a minimal amount of downtime.

Presented below are our full instructions on how to safely update a pair of load balancers with minimal disruption to load balanced services and clients.

Online and offline update options are both possible and are described below.

15.5.1. General Guidance for Performing Updates

- **Maintenance window**: since services may be restarted during the update process, we recommend performing the update during a maintenance window
- **Updates are incremental for versions prior to v8.8.0**: repeat the update process ignoring calls to restart services or reboot the appliance until you reach v8.8.0. Once at v8.8.0, when the update is next triggered the appliance is updated to the latest version in a single step.
- **Backups**: we recommend that you backup your XML configuration and firewall script (if changes have been made) before running an update. Do this by using the buttons in the web interface, under *Maintenance > Backup & Restore*.

15.5.2. Specific Guidance for Updating a Clustered Pair of Load Balancers

1. Perform the update on the passive appliance first, which is usually the Secondary. The updates are incremental for versions prior to v8.8.0, so repeat the update process ignoring calls to restart services or reboot the appliance until you reach v8.8.0. Once at v8.8.0, when the update is next triggered the appliance is updated to the latest version in a single step.
2. Next, restart services or reboot the appliance as directed.
3. Fail over to the updated appliance so that it becomes the active appliance.
4. Now update the other unit in the same way.
    - For a clustered pair, we strongly recommend fully testing & validating the Primary/Secondary fail-over process before going live. If testing was not carried out before go-live, we recommend scheduling a maintenance window to do this. For detailed steps, please refer to the *Administration Manual*.

15.5.3. Online Updates

To perform an online update:

- Ensure the load balancer can access the Internet (requires a valid default gateway and DNS server to be set).
- In the web interface, open *Maintenance > Software Update*
- Select *Online Update*
- If an update is available, information about the update will be displayed
- Click the *Online Update* button
- Once complete (the update can take several minutes depending on download speed and upgrade version), the message 'Update completed successfully' will be displayed

- If there are any specific post-upgrade requirements, such as a service restart, these will be displayed on the screen after the installation completes

15.5.4. Offline Updates

To perform an offline update:

- In the web interface, open *Maintenance > Software Update*
- Select *Offline Update*
- Select the correct matching update archive and checksum files
- Click *Upload and Install*

15.6. 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.

15.6.1. 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

WebUI Main Menu Option	Sub Menu Option	Description
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.

15.6.2. 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*.

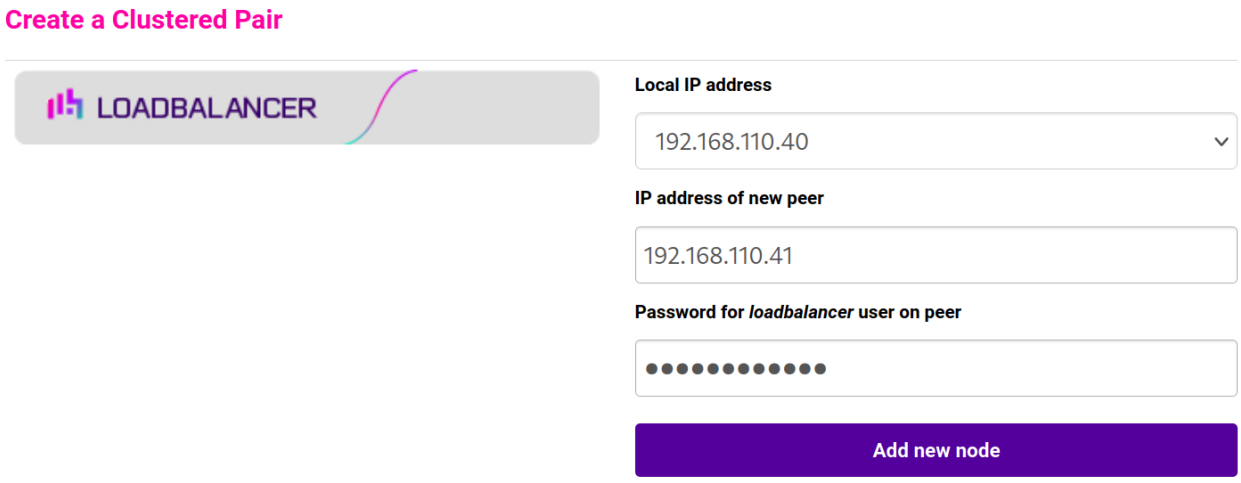

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

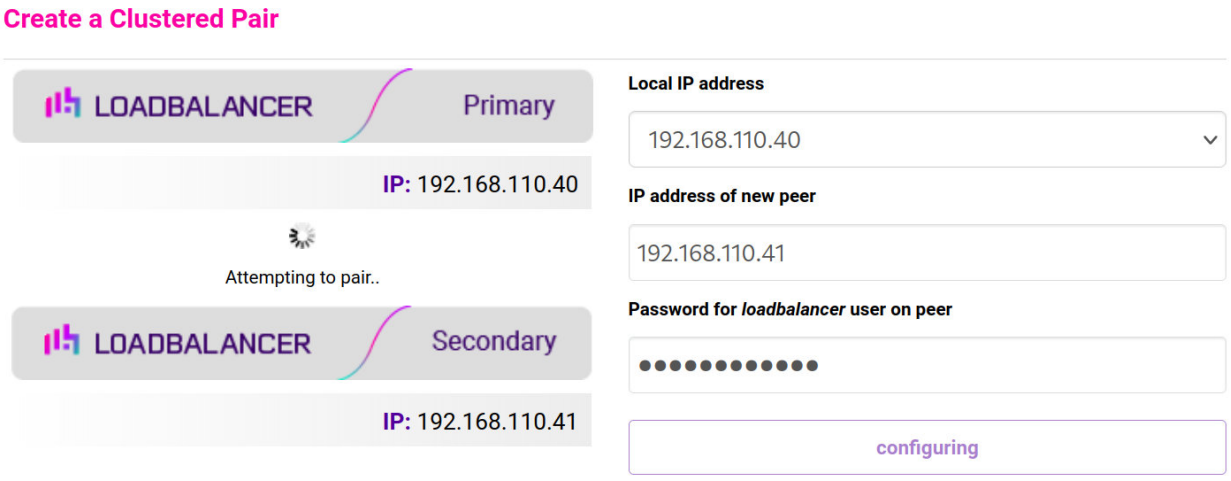

6. Once complete, the following will be displayed on the Primary appliance:

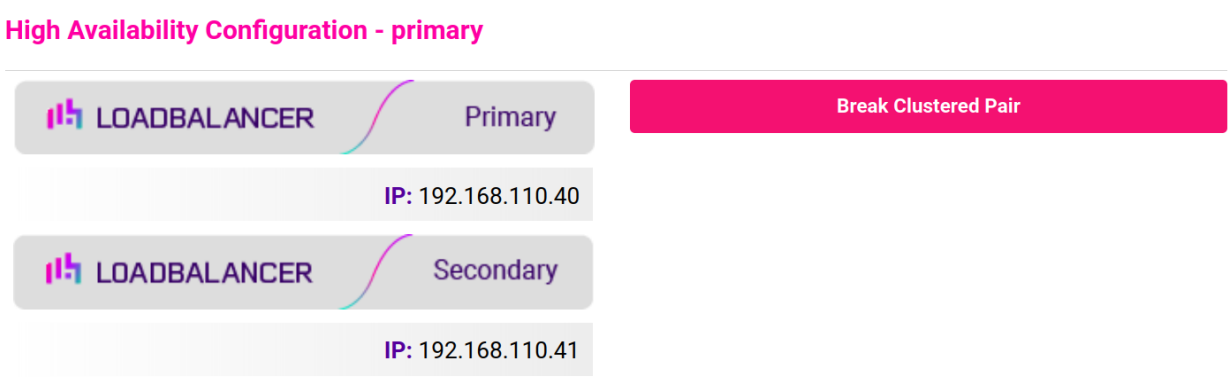

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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 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	10 May 2018	Initial version		AH
1.0.1	24 May 2018	Total re-write, based on Metaswitch's load balancing document sent over on the back of the initial version	Document re-write	AH
1.0.2	5 June 2018	Changed 'Metaswitch EAS SSS' to 'Metaswitch Virtual EAS SSS' at Metaswitch's request Added a paragraph saying that although the guide is for virtual it still applies to hardware	Required updates	AH
1.0.3	6 August 2018	Corrected the layer 7 persistence timeout instructions, as the units are minutes while Metaswitch customers will have a list of timeouts to input that are in seconds Added a new Appendix 1, 'Confirming the Gateway Settings on the EAS Servers' Changed 'Pool Configuration Test' paragraph to remove reference to using Python 3.x; Metaswitch have advised running the script with Python 2.7 only Corrected the System Overview screenshot under 'Useful Load Balancer Based Checks' where the MGMT_HTTP-80 VIP was erroneously in HTTP Mode	Required updates	AH
1.0.4	19 September 2018	Added section 6, "Sizing, Capacity, and Performance"	Required updates	AH
1.0.5	27 September 2018	Added section 7, "Using WAF Gateways"	Required updates	AH
1.0.6	6 December 2018	Added the new "Company Contact Information" page	Required updates	AH

Version	Date	Change	Reason for Change	Changed By
1.1.0	15 August 2019	Styling and layout Updated the advice in section 7, "Using WAF Gateways" Added a paragraph describing layer 4 health check timeout changes, to bring them in line with the layer 7 health check settings	General styling updates Required updates	AH
1.1.1	28 August 2020	New title page Updated Canadian contact details Amended instructions for setting persistence timeout options	Branding update Change to Canadian contact details Changes to the appliance WebUI	AH
1.2.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH,RJC,ZAC
1.2.1	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.2.2	5 January 2023	Combined software version information into one section Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section Removed references to the colour of certain UI elements	Housekeeping across all documentation	AH
1.2.3	2 February 2023	Updated screenshots	Branding update	AH
1.2.4	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.3.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH

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.