LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing Dell EMC ECS

Version 1.3.0

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. Dell EMC ECS . . . 4
4. Dell EMC ECS . . . 4
5. Load Balancing Dell EMC ECS . . . 4
    5.1. Persistence (aka Server Affinity) . . . 5
    5.2. Virtual Service (VIP) Requirements . . . 5
    5.3. Port Requirements . . . 5
    5.4. TLS/SSL Termination . . . 6
    5.5. Health Checks . . . 6
        5.5.1. S3, Atmos, and Swift Virtual Services . . . 6
        5.5.2. NFS Virtual Service . . . 6
6. Deployment Concept . . . 6
    6.1. Scenario 1 – Virtual Services for Each Protocol . . . 6
        6.1.1. Method A: Sorting by Port . . . 7
        6.1.2. Method B: Sorting by IP Address . . . 7
    6.2. Scenario 2 – Single Client-facing Virtual Service . . . 8
    6.3. Helping you Choose the Most Appropriate Deployment Type . . . 8
7. Loadbalancer.org Appliance – the Basics . . . 9
    7.1. Virtual Appliance . . . 9
    7.2. Initial Network Configuration . . . 10
    7.3. Accessing the Appliance WebUI . . . 10
        7.3.1. Main Menu Options . . . 11
    7.4. Appliance Software Update . . . 12
        7.4.1. Online Update . . . 12
        7.4.2. Offline Update . . . 12
    7.5. Ports Used by the Appliance . . . 13
    7.6. HA Clustered Pair Configuration . . . 14
8. Appliance Configuration for Dell EMC ECS – Scenario 1 . . . 14
    8.1. Changing the Global Layer 7 Settings . . . 14
    8.2. Configuring VIP 1 – S3 . . . 14
        8.2.1. Configuring the Virtual Service (VIP) . . . 14
        8.2.2. Defining the Real Servers (RIPs) . . . 15
    8.3. Configuring VIP 2 – Atmos . . . 16
        8.3.1. Configuring the Virtual Service (VIP) . . . 16
        8.3.2. Defining the Real Servers (RIPs) . . . 16
    8.4. Configuring VIP 3 – Swift . . . 17
        8.4.1. Configuring the Virtual Service (VIP) . . . 17
        8.4.2. Defining the Real Servers (RIPs) . . . 18
    8.5. Configuring VIP 4 – NFS . . . 18
        8.5.1. Configuring the Virtual Service (VIP) . . . 18
        8.5.2. Defining the Real Servers (RIPs) . . . 19
    8.6. Finalizing the Configuration . . . 20
9. Appliance Configuration for Dell EMC ECS – Scenario 2 . . . 20
    9.1. Changing the Global Layer 7 Settings . . . 20
    9.2. Configuring VIP 1 – S3 . . . 20

9.2.1. Configuring the Virtual Service (VIP) . . . 20
9.2.2. Defining the Real Servers (RIPs) . . . 21
9.3. Configuring VIP 2 – Atmos . . . 22
9.3.1. Configuring the Virtual Service (VIP) . . . 22
9.3.2. Defining the Real Servers (RIPs) . . . 22
9.4. Configuring VIP 3 – Swift . . . 23
9.4.1. Configuring the Virtual Service (VIP) . . . 23
9.4.2. Defining the Real Servers (RIPs) . . . 24
9.5. Configuring VIP 4 – NFS . . . 25
9.5.1. Configuring the Virtual Service (VIP) . . . 25
9.5.2. Defining the Real Servers (RIPs) . . . 26
9.6. Configuring VIP 5 – ECS Combined Service . . . 26
9.6.1. Configuring the Virtual Service (VIP) . . . 26
9.6.2. Setting Up the TLS/SSL Termination . . . 28
9.6.3. Finalizing the Configuration . . . 29
10. Testing & Verification . . . 29
10.1. Using System Overview . . . 29
11. Technical Support . . . 31
12. Further Documentation . . . 31
13. Appendix . . . 32
13.1. Multi-port NFS Health Check . . . 32
13.2. Configuring HA - Adding a Secondary Appliance . . . 32
13.2.1. Non-Replicated Settings . . . 33
13.2.2. Configuring the HA Clustered Pair . . . 33
14. Document Revision History . . . 36

1. About this Guide

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

- All versions

4. Dell EMC ECS

ECS (Elastic Cloud Storage) is an object storage solution developed by Dell EMC. It uses hardware 'nodes' to provide storage, and is designed to be flexible, resilient, and simple to deploy.

Dell recommend the use of load balancing in an ECS deployment, in order to distribute the inbound workload across all ECS nodes in an effort to maximise performance.

One of Dell EMC's approved and documented solutions for load balancing ECS is the free and open source HAProxy load balancer. HAProxy is a key component of the Loadbalancer.org appliance, making it a great fit for load balancing ECS deployments.

5. Load Balancing Dell EMC ECS

> **Note** It's highly recommended that you have a working Dell EMC ECS environment first before implementing the load balancer.

5.1. Persistence (aka Server Affinity)

Persistence is only recommended for NFS connections when load balancing a Dell EMC ECS deployment. This is due to the fact that caching occurs on the ECS servers when the NFS protocol is used. To maximize efficiency, a given NFS client should continue connecting to the same ECS server, so as to continue re-using the established cache.

5.2. Virtual Service (VIP) Requirements

To provide load balancing and HA for Dell EMC ECS, the following VIPs are usually required:

- S3 (for object access via the S3 protocol)
- Atmos (for object access via the Atmos protocol)
- Swift (for object access via the Swift protocol)
- NFS (for providing highly available NFS services)

Optionally, additional VIPs may be required as follows:

- ECS Combined Service (for scenario 2, where only a single IP address is client-facing)
- TLS/SSL termination service (for scenario 2, where HTTPS traffic must be decrypted for inspection)

5.3. Port Requirements

The following table shows the ports that are load balanced:

Port	Protocols	Use
80	TCP/HTTP	Object access via HTTP calls
443	TCP/HTTPS	Object access via HTTP calls (encrypted HTTPS)
9020	TCP/HTTP	Object access via the S3 protocol (HTTP)
9021	TCP/HTTPS	Object access via the S3 protocol (HTTPS)
9022	TCP/HTTP	Object access via the Atmos protocol (HTTP)
9023	TCP/HTTPS	Object access via the Atmos protocol (HTTPS)
9024	TCP/HTTP	Object access via the Swift protocol (HTTP)
9025	TCP/HTTPS	Object access via the Swift protocol (HTTPS)
2049	TCP/UDP/NFS	NFS service (mountd and nfsd)
111	TCP/UDP/ONC RPC	Port mapper service
10000	TCP/lockd	lockd NFS service

5.4. TLS/SSL Termination

Terminating TLS/SSL connections on the load balancer is not recommended, due to the significant computational overhead this introduces on the load balancer. Termination and decryption should continue to occur at the ECS servers, which are designed and best placed to perform this function.

It may be necessary to terminate and decrypt traffic at the load balancer, so that it may then be read as plaintext and sorted. This is required if sorting incoming traffic by protocol is not possible by using different ports or IP addresses. This is explained in detail in Deployment Concept.

5.5. Health Checks

5.5.1. S3, Atmos, and Swift Virtual Services

The S3 and Swift virtual services use protocol-specific health checks to query the readiness of a given ECS server to accept connections for those protocols.

The Atmos virtual service uses a standard 'connect to port' check, which examines whether the Atmos port is open on a given ECS server to determine whether the server is ready to accept connections using the Atmos protocol.

5.5.2. NFS Virtual Service

The NFS virtual service uses a standard 'connect to port' check by default, which examines whether the NFS port (2049) is open on a given ECS server to determine whether the server is ready to accept NFS connections.

It is possible to configure a custom health check for the NFS service, which will check the availability of all three ports related to NFS operation (111, 2049, and 10000) on the real servers. Only if all three ports are available will a real server be considered 'healthy' and ready to accept NFS connections.

Please refer to the Appendix section Multi-port NFS Health Check for instructions on how to configure such a custom health check.

6. Deployment Concept

There are two deployment scenarios when using Loadbalancer.org appliances as part of a Dell EMC ECS deployment.

6.1. Scenario 1 – Virtual Services for Each Protocol

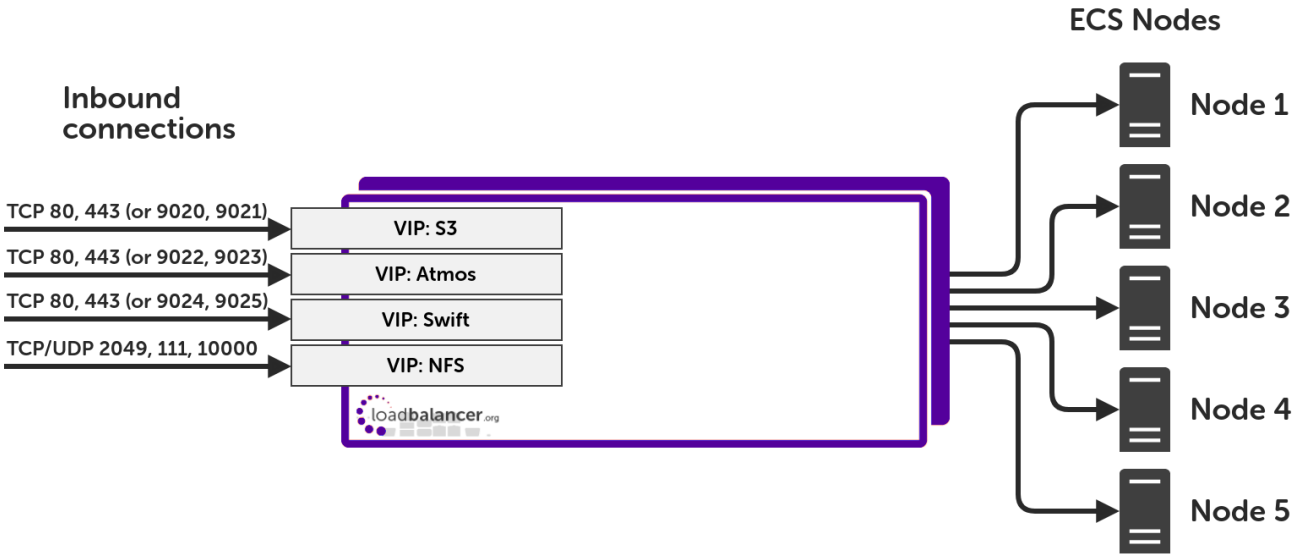

VIP = **V**irtual **IP** Address

> **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 appendix section Configuring HA - Adding a Secondary Appliance for more details on configuring a clustered pair.

This is the preferred scenario, and is the easiest to implement. Incoming traffic is not decrypted or modified in any way.

The different protocols are handled by different virtual services on the load balancer. Each virtual service is client-facing. **The traffic needs to be sorted by protocol (S3, Atmos, and Swift) by the time it reaches the load balancer**, either sorted by port or sorted by IP address. The S3 traffic needs to go to the S3 virtual service, the Atmos traffic needs to go to the Atmos virtual service, and the Swift traffic needs to go to the Swift virtual service.

6.1.1. Method A: Sorting by Port

This is the simplest way of sorting ECS traffic. It assumes that your clients are able to send request traffic using the correct protocol-specific ports. For example, an S3 client would send request traffic using ports 9020 and 9021. This is likely to be the case for an internal, non-public Internet facing ECS deployment.

The load balancer's S3, Atmos, and Swift virtual services would all use the same IP address, but would each listen on their respective ports in the 9020-9025 range.

6.1.2. Method B: Sorting by IP Address

This sorting method is a good alternative if sorting by port is not a possibility (for example if client traffic is being sent over the public Internet and ports 80 and 443 must be used for ***all*** traffic of all protocols).

A simple way of sorting the incoming traffic by protocol is to use multiple DNS records, one for each protocol. For example:

- os.website.org (FQDN for the S3 service) resolves to the IP address of the S3 VIP
- atmos.website.org (FQDN for the Atmos service) resolves to the IP address of the Atmos VIP

- swift.website.org (FQDN for the Swift service) resolves to the IP address of the Swift VIP

If the FQDNs in question need to resolve to public IP addresses, a valid setup would be to put the public IP addresses on an external facing firewall and then forward the traffic to the relevant load balancer VIPs.

6.2. Scenario 2 – Single Client-facing Virtual Service

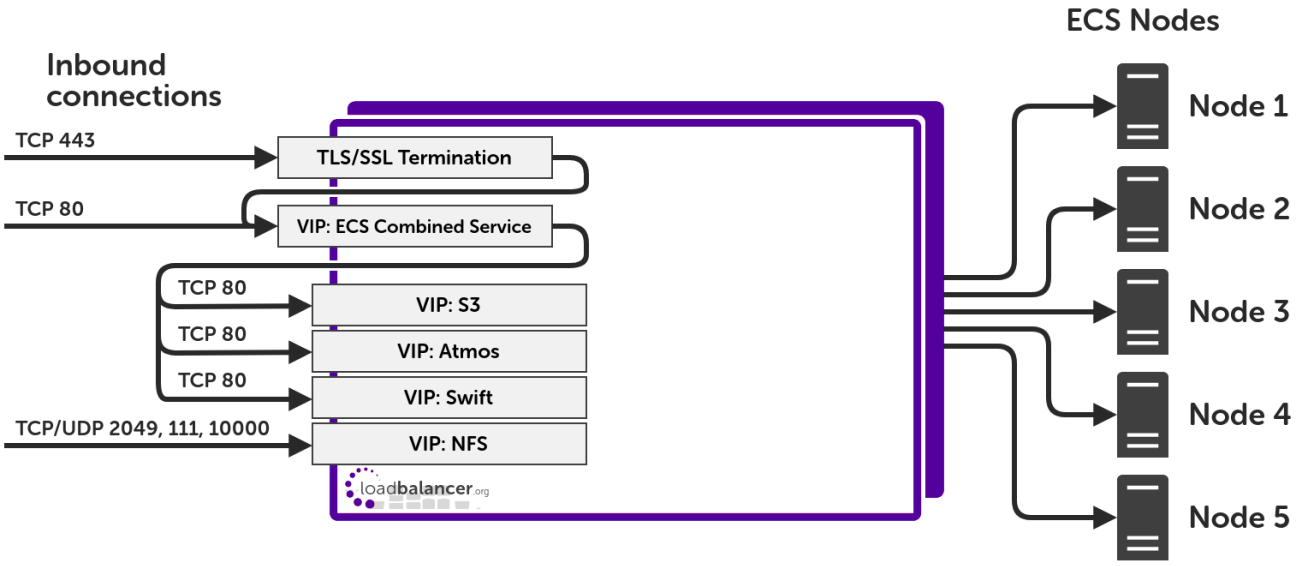

VIP = **V**irtual **IP** Address

> **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 appendix section Configuring HA - Adding a Secondary Appliance for more details on configuring a clustered pair.

This scenario involves inspecting all incoming HTTP(S) traffic and sorting it by FQDN, so that it may be forwarded to the correct protocol-specific virtual service, i.e. the S3, Atmos, or Swift virtual service.

This deployment type is useful when it is not possible to pre-sort traffic by port (for example if clients are on the public Internet and traffic needs to be sent using ports 80 and 443 only) or by IP address (for example if changing public DNS records is not possible).

The disadvantages of this setup are that it is more complex to set up than scenario 1 and that all incoming TLS/SSL encrypted traffic must be decrypted for inspection, which is CPU intensive on the load balancer.

6.3. Helping you Choose the Most Appropriate Deployment Type

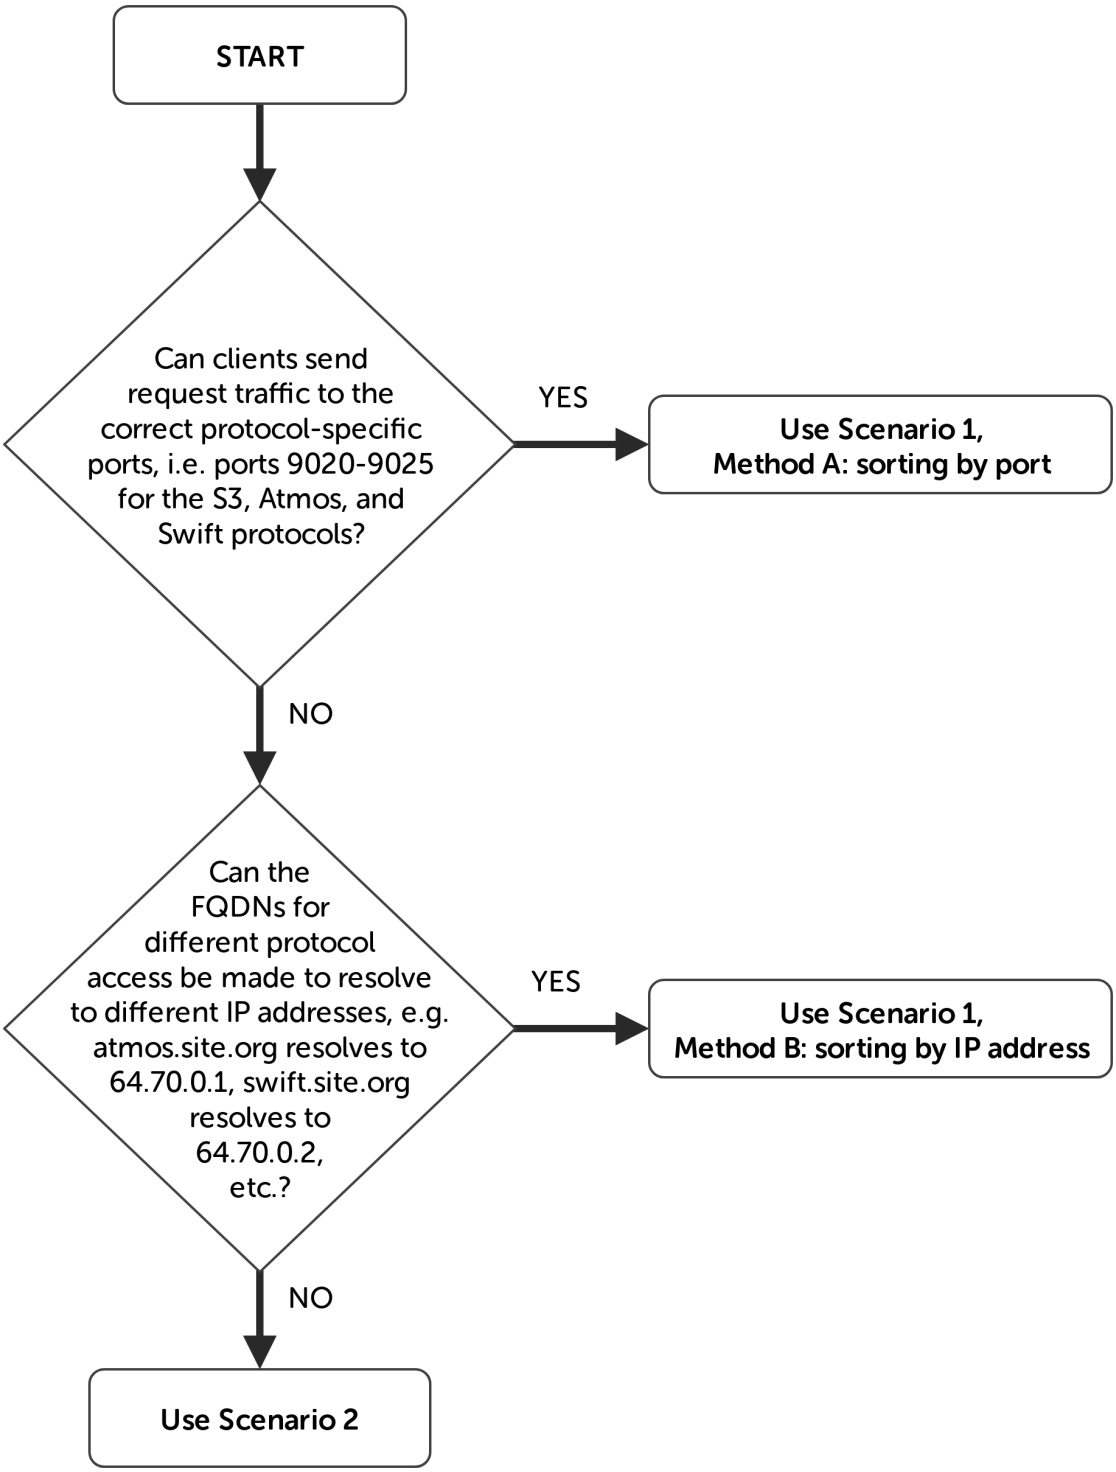

7. Loadbalancer.org Appliance – the Basics

7.1. 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.

7.2. 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.

7.3. 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:

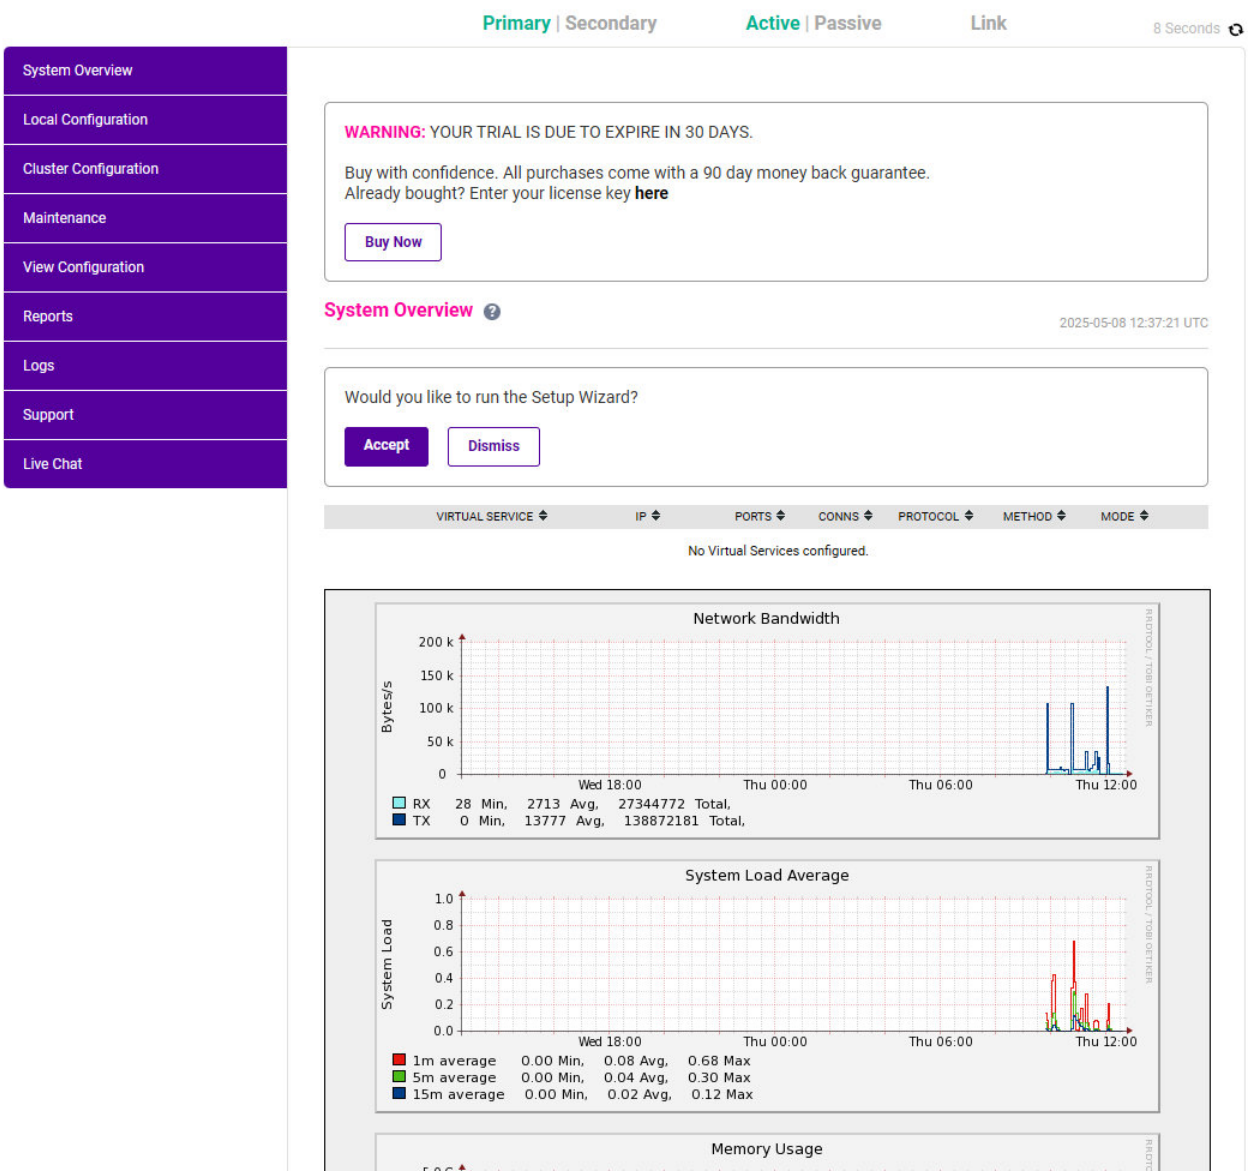

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

7.3.1. 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

7.4. 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.

7.4.1. 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:

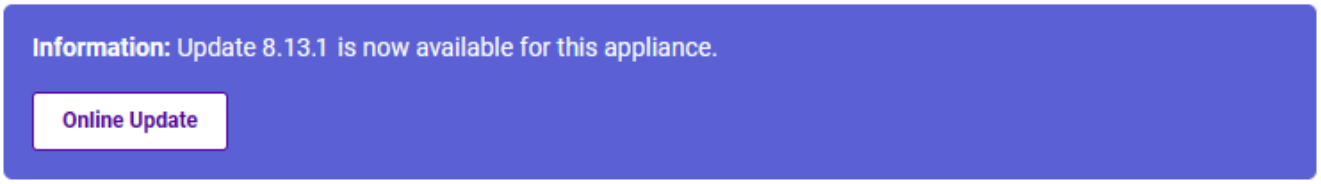

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

> **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:

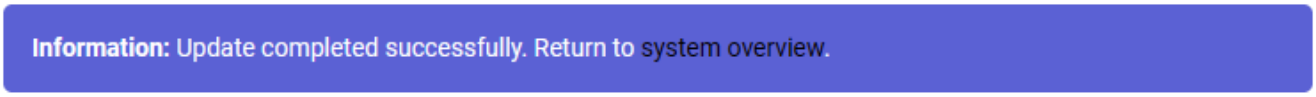

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

7.4.2. 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:

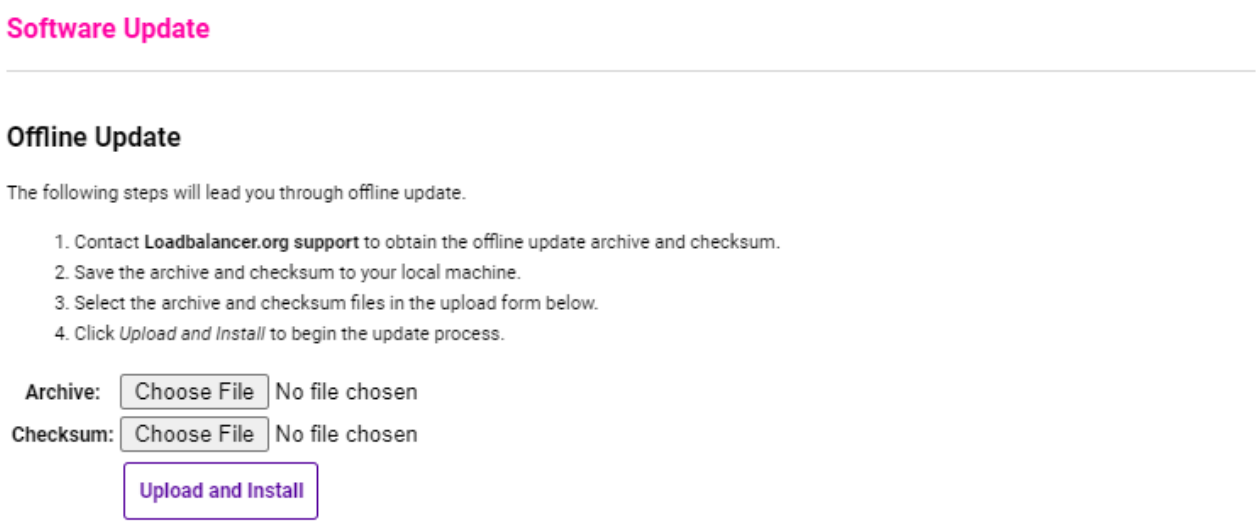

4. Select the ***Archive*** and ***Checksum*** files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

7.5. 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.

7.6. 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 appendix section Configuring HA - Adding a Secondary Appliance.

8. Appliance Configuration for Dell EMC ECS – Scenario 1

8.1. Changing the Global Layer 7 Settings

It is necessary to change some global layer 7 timeout settings when load balancing an ECS deployment.

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Advanced Configuration*.
2. Set the *Connection Timeout* value to **5000**.
3. Set the *Client Timeout* value to **50000**.
4. Set the *Real Server Timeout* to **50000**.
5. Click **Update** to apply the settings.

8.2. Configuring VIP 1 – S3

8.2.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-S3**.
3. Set the *Virtual Service IP* address field as required:
   - If using method A (sorting by port), use the same IP address for all virtual services
   - If using method B (sorting by IP address), use a unique IP address for the S3 virtual service
4. Set the *Virtual Service Ports* as required:
   - If using method A (sorting by port), use ports **9020,9021**
   - If using method B (sorting by IP address), use ports **80,443**
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

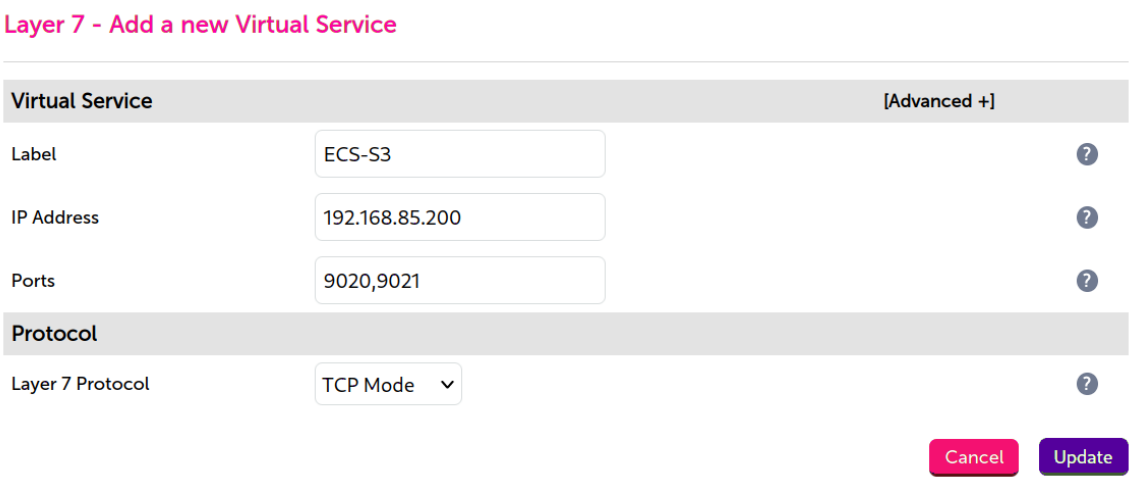

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. In the *Health Checks* section, click **Advanced** to show more options.
10. Set *Health Checks* to **Negotiate HTTP (GET)**.
11. Set *Request to send* to **/?ping**.
12. Set *Host Header* to **haproxy**.
13. Click **Update**.

8.2.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Leave the *Real Server Port* field blank.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

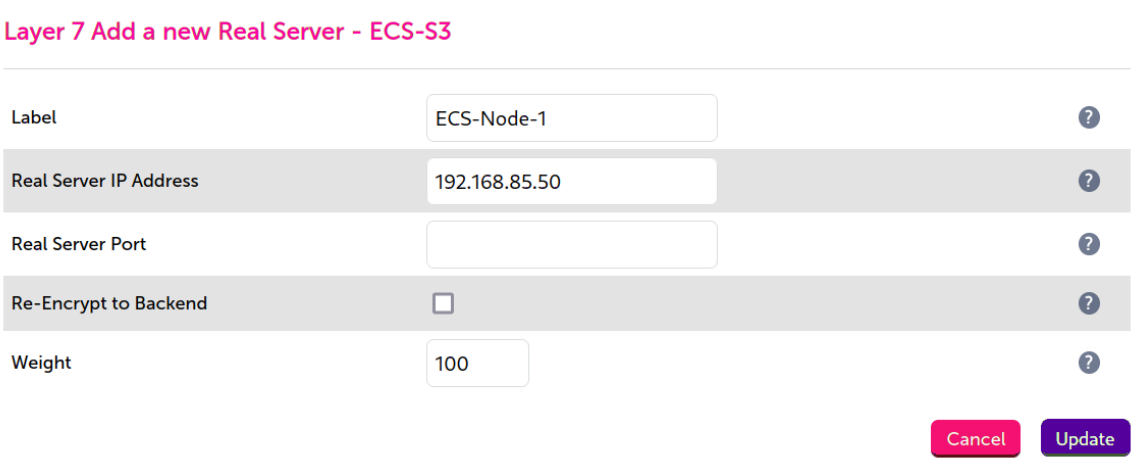

8.3. Configuring VIP 2 – Atmos

8.3.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-Atmos**.
3. Set the *Virtual Service IP* address field as required:
   - If using method A (sorting by port), use the same IP address for all virtual services
   - If using method B (sorting by IP address), use a unique IP address for the Atmos virtual service
4. Set the *Virtual Service Ports* as required:
   - If using method A (sorting by port), use ports **9022,9023**
   - If using method B (sorting by IP address), use ports **80,443**
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

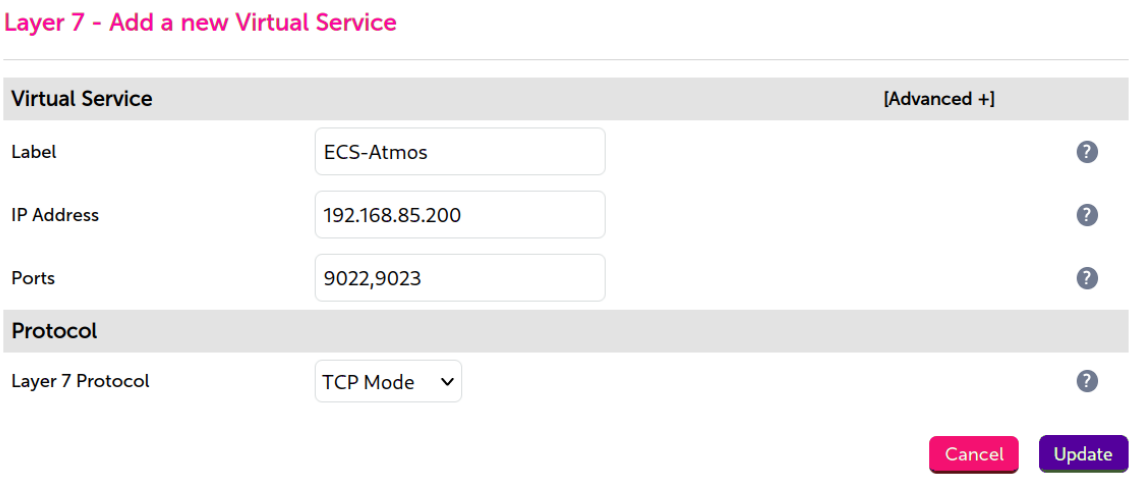

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Click **Update**.

8.3.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Leave the *Real Server Port* field blank.
5. Click **Update**.

6. Repeat these steps to add additional servers as required.

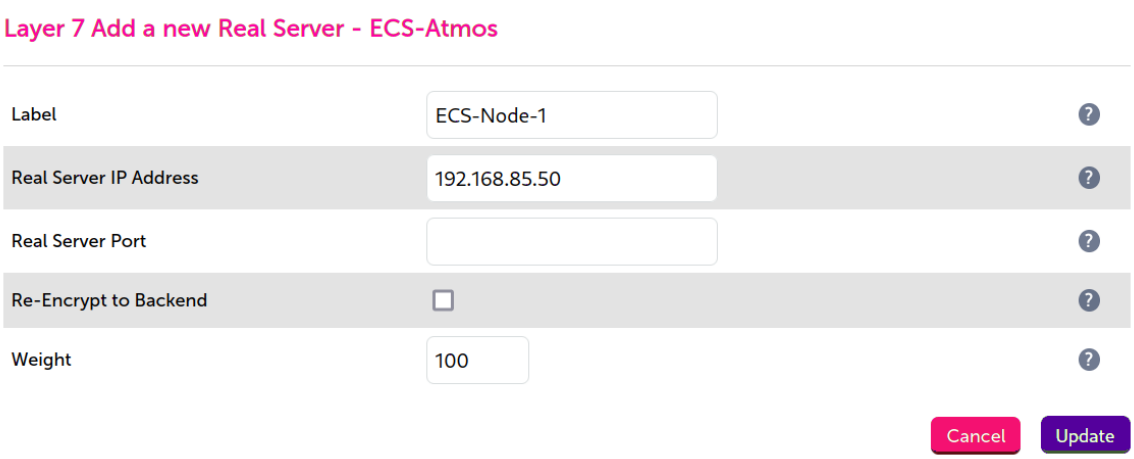

8.4. Configuring VIP 3 – Swift

8.4.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-Swift**.
3. Set the *Virtual Service IP* address field as required:
   - If using method A (sorting by port), use the same IP address for all virtual services
   - If using method B (sorting by IP address), use a unique IP address for the Swift virtual service
4. Set the *Virtual Service Ports* as required:
   - If using method A (sorting by port), use ports **9024,9025**
   - If using method B (sorting by IP address), use ports **80,443**
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

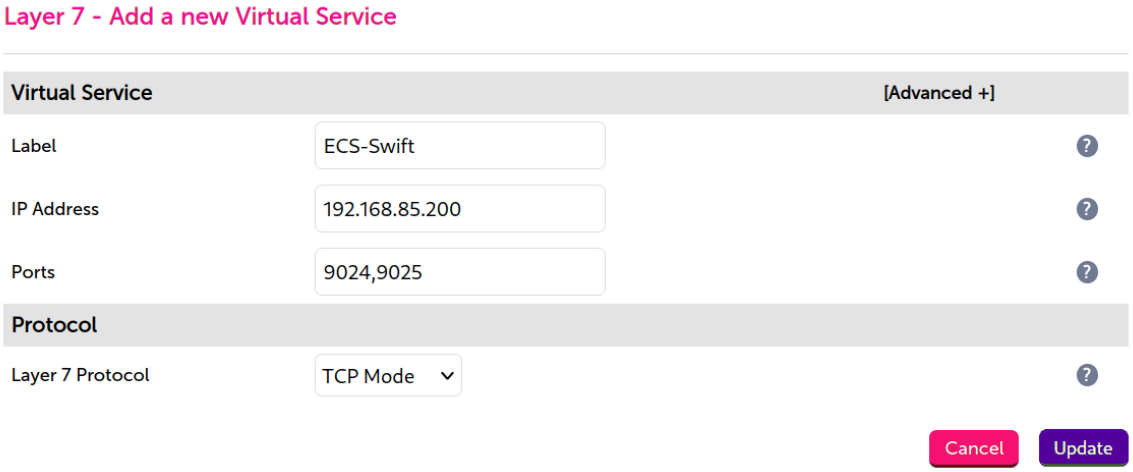

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Set *Health Checks* to **Negotiate HTTP (HEAD)**.
10. Set *Request to send* to **/healthcheck**.
11. Click **Update**.

8.4.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Leave the *Real Server Port* field blank.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

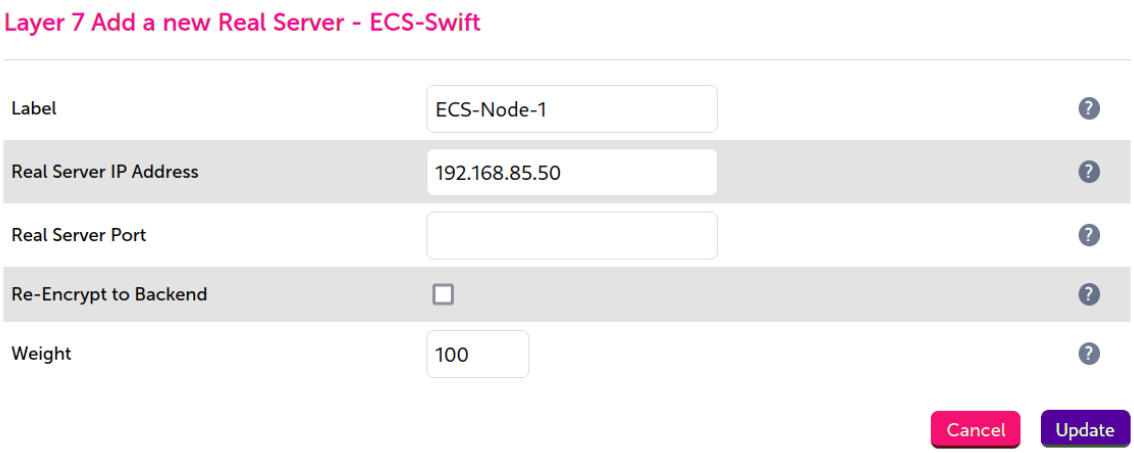

8.5. Configuring VIP 4 – NFS

8.5.1. 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. **ECS-NFS**.
3. Set the *Virtual Service IP* address field to the IP address to be used for NFS access, in this example **192.168.85.200**.
4. Set the *Virtual Service Ports* field to **111,2049,10000**.
5. Set the *Protocol* to **TCP/UDP**.
6. Set the *Forwarding Method* to **SNAT**.
7. Click **Update** to create the virtual service.

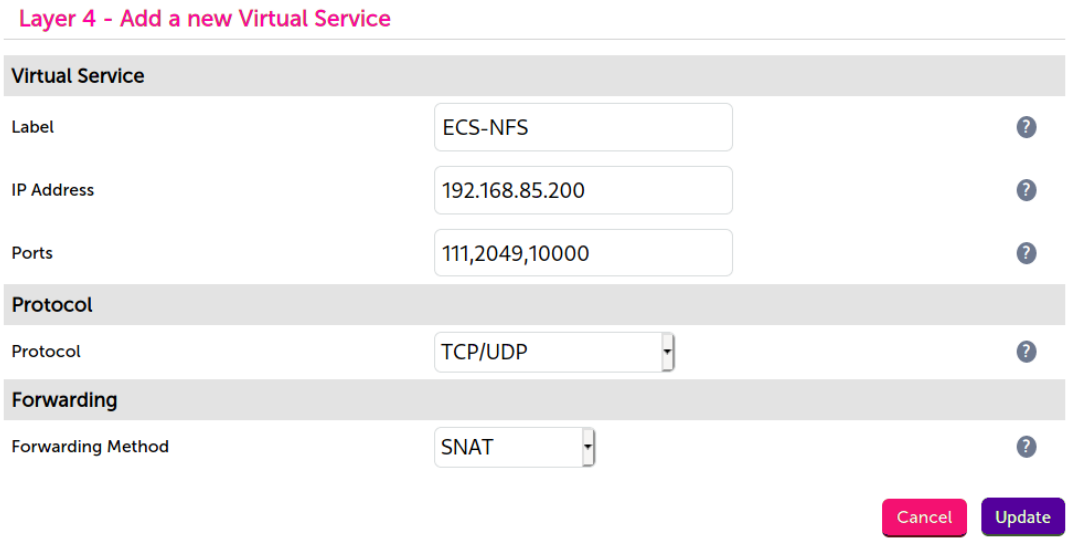

8. Click **Modify** next to the newly created VIP.
9. Ensure the *Persistence Enable* checkbox is enabled.
10. Set the *Persistence Timeout* field to **86400** (as the units are seconds, this equates to 24 hours).
11. Set *Check Type* to **Connect to port**.
12. Set *Check Port* to **2049**.

> **Note**
>
> In the default setup presented here, each ECS server will be checked on port 2049 only (the NFS port) to judge whether the server is ready to accept NFS connections.
>
> It is possible to configure a custom health check for the NFS service, which will check the availability of ***all three*** ports related to NFS operation (111, 2049, and 10000) on the real servers. In that case, only if all three ports are available will a real server be considered 'healthy' and ready to accept connections.
>
> Please refer to the appendix section Multi-port NFS Health Check for instructions on how to configure such a custom health check.

13. Click **Update**.

8.5.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Leave the *Real Server Port* field blank.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

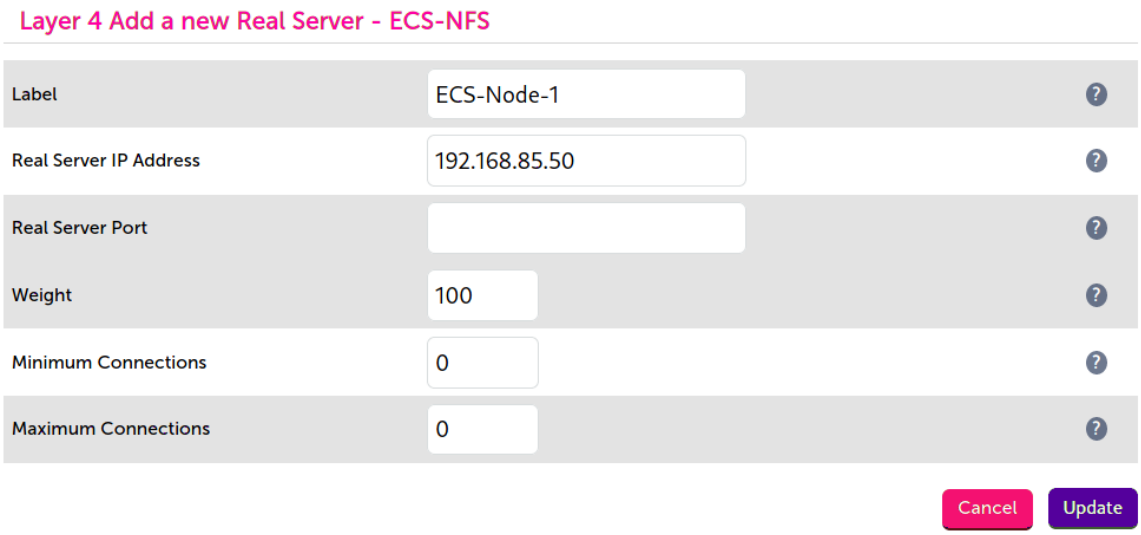

8.6. Finalizing the Configuration

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**.

9. Appliance Configuration for Dell EMC ECS – Scenario 2

9.1. Changing the Global Layer 7 Settings

It is necessary to change some global layer 7 timeout settings when load balancing an ECS deployment.

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Advanced Configuration*.
2. Set the *Connection Timeout* value to **5000**.
3. Set the *Client Timeout* value to **50000**.
4. Set the *Real Server Timeout* to **50000**.
5. Click **Update** to apply the settings.

9.2. Configuring VIP 1 – S3

9.2.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-S3**.
3. Set the *Virtual Service IP* address field to an unused IP address, e.g. **192.168.85.200**.
4. Set the *Virtual Service Ports* field to **80**.

5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

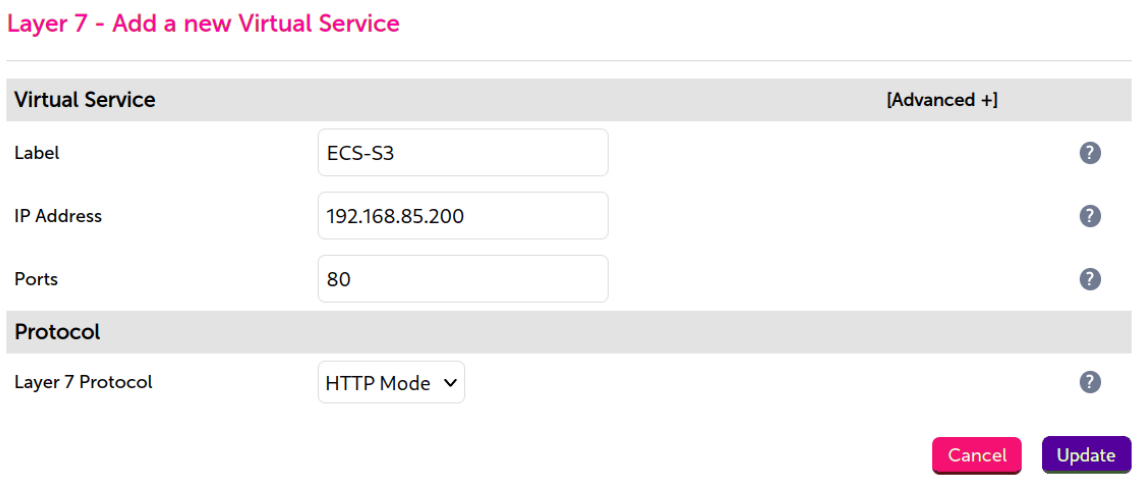

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. In the *Health Checks* section, click **Advanced** to show more options.
10. Set *Health Checks* to **Negotiate HTTP (GET)**.
11. Set *Request to send* to **/?ping**.
12. Set *Host Header* to **haproxy**.
13. Click **Update**.

9.2.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Set the *Real Server Port* field to **9020**.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

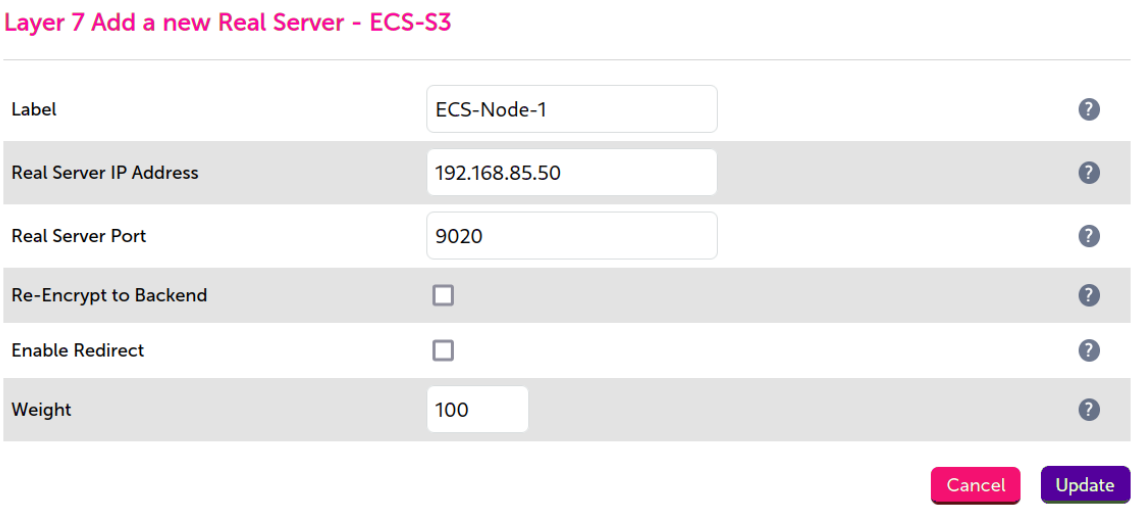

9.3. Configuring VIP 2 – Atmos

9.3.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-Atmos**.
3. Set the *Virtual Service IP* address field to an unused IP address, e.g. **192.168.85.201**.
4. Set the *Virtual Service Ports* field to **80**.
5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

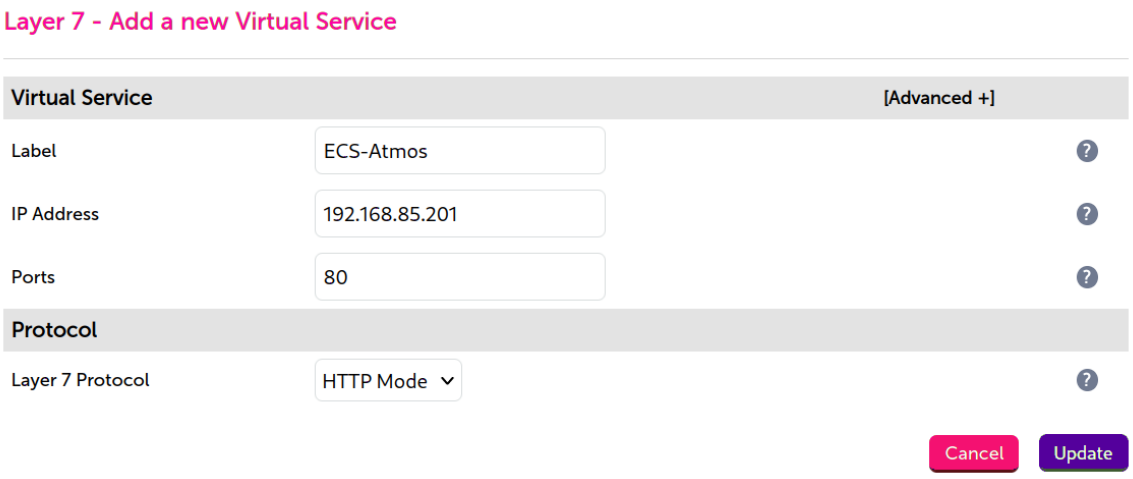

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Click **Update**.

9.3.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Set the *Real Server Port* field to **9022**.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

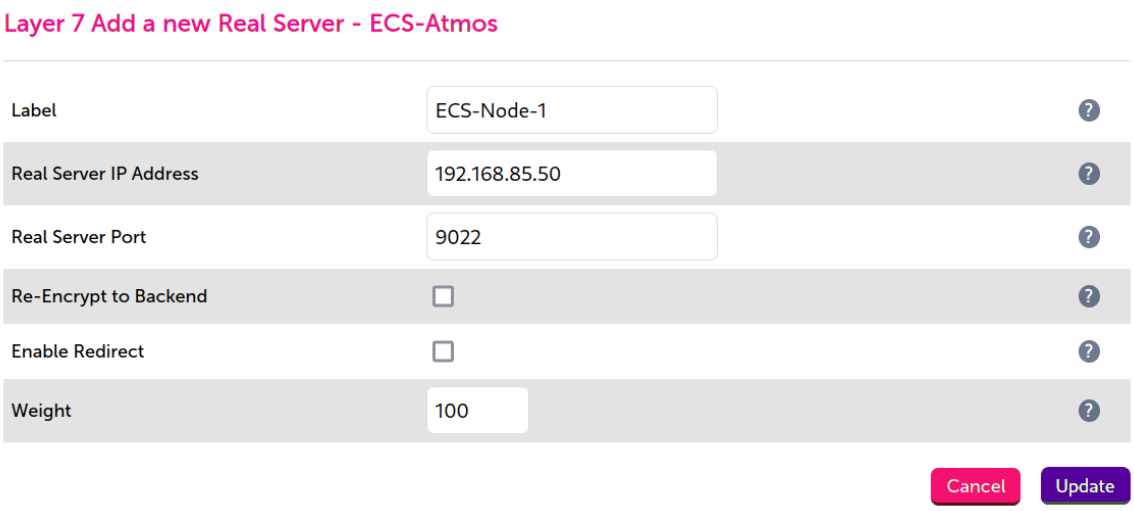

9.4. Configuring VIP 3 – Swift

9.4.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-Swift**.
3. Set the *Virtual Service IP* address field to an unused IP address, e.g. **192.168.85.202**.
4. Set the *Virtual Service Ports* field to **80**.
5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

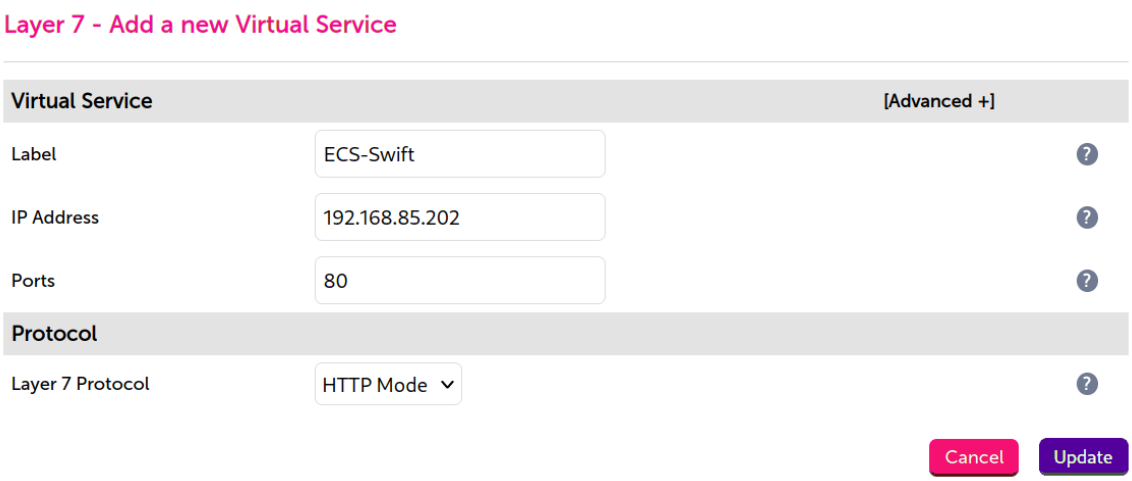

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Set *Health Checks* to **Negotiate HTTP (HEAD)**.
10. Set *Request to send* to **/healthcheck**.
11. Click **Update**.

9.4.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Set the *Real Server Port* field to **9024**.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

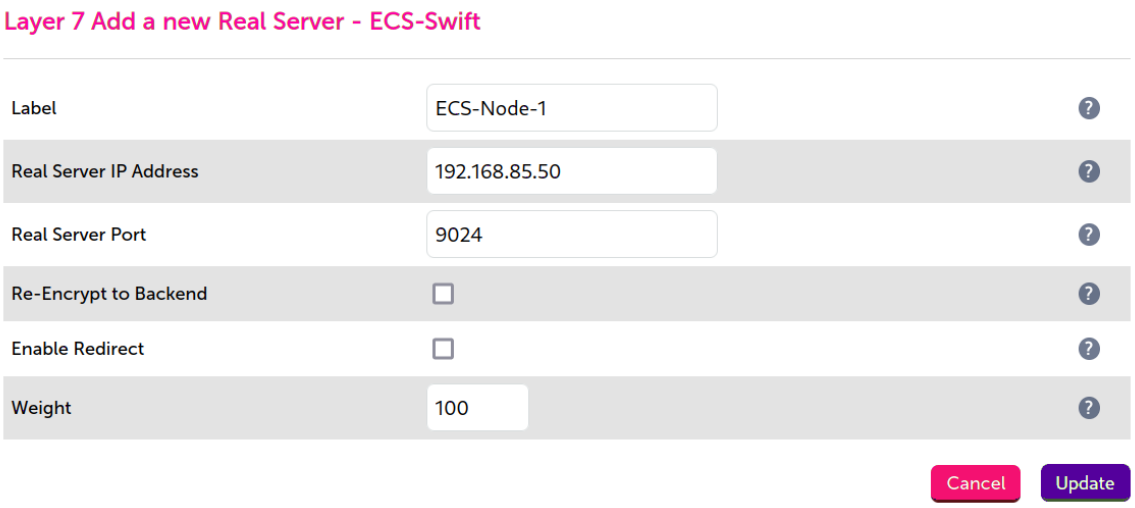

9.5. Configuring VIP 4 – NFS

9.5.1. 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. **ECS-NFS**.
3. Set the *Virtual Service IP* address field to the IP address to be used for NFS access, in this example **192.168.85.203**.
4. Set the *Virtual Service Ports* field to **111,2049,10000**.
5. Set the *Protocol* to **TCP/UDP**.
6. Set the *Forwarding Method* to **SNAT**.
7. Click **Update** to create the virtual service.

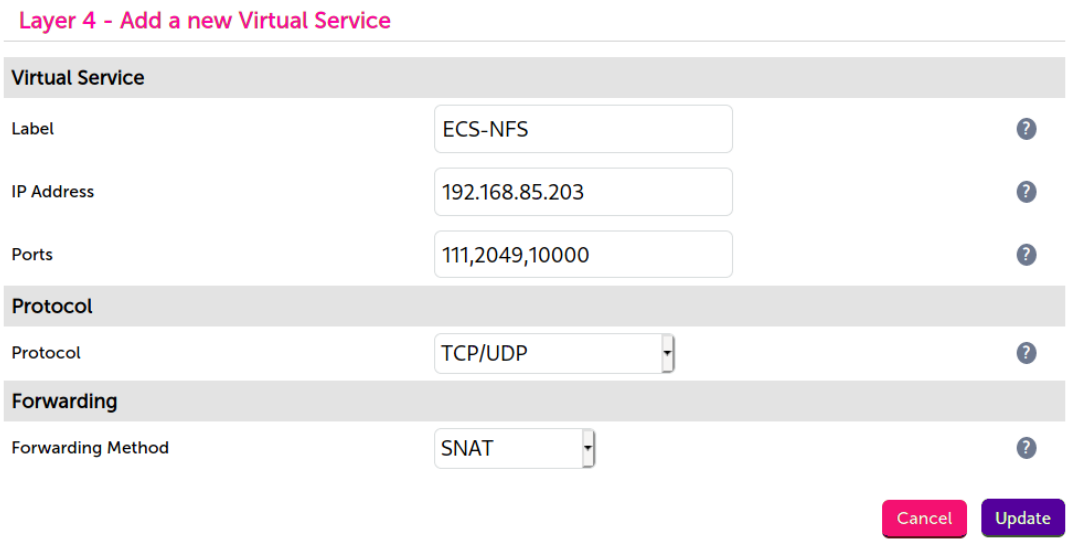

8. Click **Modify** next to the newly created VIP.
9. Ensure the *Persistence Enable* checkbox is enabled.
10. Set the *Persistence Timeout* field to **86400** (as the units are seconds, this equates to 24 hours).
11. Set *Check Type* to **Connect to port**.
12. Set *Check Port* to **2049**.

> **Note**
>
> In the default setup presented here, each ECS server will be checked on port 2049 only (the NFS port) to judge whether the server is ready to accept NFS connections.
>
> It is possible to configure a custom health check for the NFS service, which will check the availability of ***all three*** ports related to NFS operation (111, 2049, and 10000) on the real servers. In that case, only if all three ports are available will a real server be considered 'healthy' and ready to accept connections.
>
> Please refer to the appendix section Multi-port NFS Health Check for instructions on how

to configure such a custom health check.

13. Click **Update**.

9.5.2. Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter an appropriate name for the server in the *Label* field, e.g. **ECS-Node-1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.85.50**.
4. Leave the *Real Server Port* field blank.
5. Click **Update**.
6. Repeat these steps to add additional servers as required.

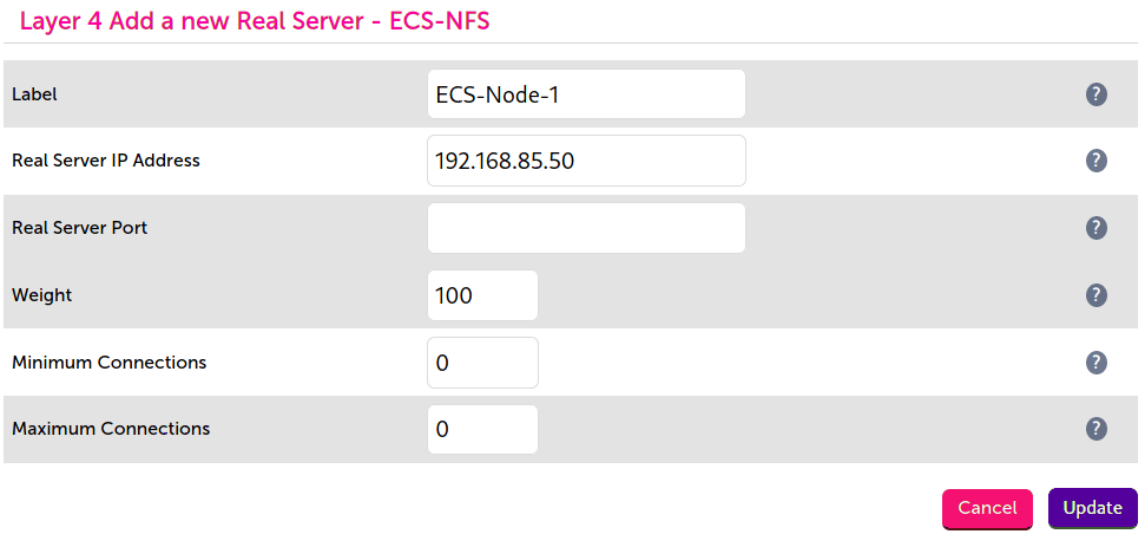

9.6. Configuring VIP 5 – ECS Combined Service

9.6.1. Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **ECS-Combined-Service**.
3. Set the *Virtual Service IP* address field to the IP address that all incoming client traffic will be arriving at on the load balancer. In the example presented here, all client traffic (regardless of port) is being sent to the IP address 192.168.85.150, and so the IP address used for the combined service is **192.168.85.150**.
4. Set the *Virtual Service Ports* field to **80**.
5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

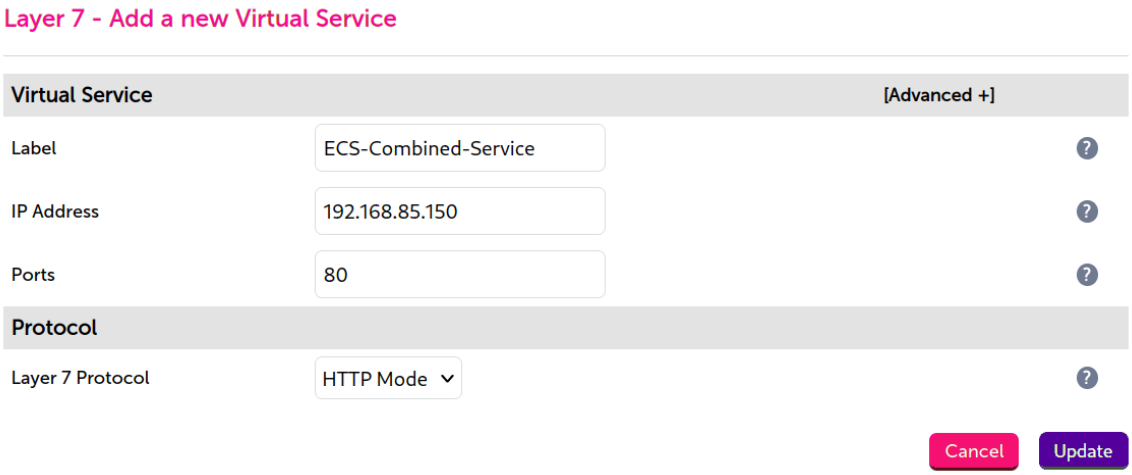

7. Click **Modify** next to the newly created VIP.
8. Under *ACL Rules* click **Add Rule**.
9. Create the first ACL rule, which will redirect S3 traffic to the S3 virtual service.
    a. Set *Type* to **hdr_host**.
    b. Set the *Bool* option to **Equals**.
    c. Set the *URL/Text* field to **-m dom**, to specify domain matching, followed by the domain in question. For example, use **-m dom os.website.org** (replacing "website.org" with the correct domain, e.g. the domain that clients' S3 protocol traffic is sent to). Note that this rule also picks up traffic to ***.os.website.org**, which accounts for S3 virtually hosted buckets.
    d. Set *Action* to **Use Backend**.
    e. Set the *Location/Value* to the name of the S3 virtual service created earlier, which in the example presented here is **ECS-S3**.
    f. Click the **Ok** button to add the ACL rule.

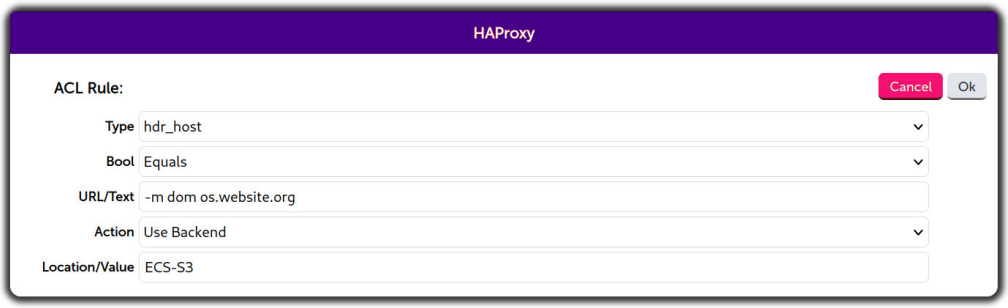

10. Repeat the steps above to create similar ACL rules to redirect Atmos and Swift protocol traffic to their respective virtual services in the same way. Use *URL/Text* values of the form **atmos.website.org** for Atmos traffic and **swift.website.org** for Swift traffic. The completed set of rules for the example presented here look like the following:

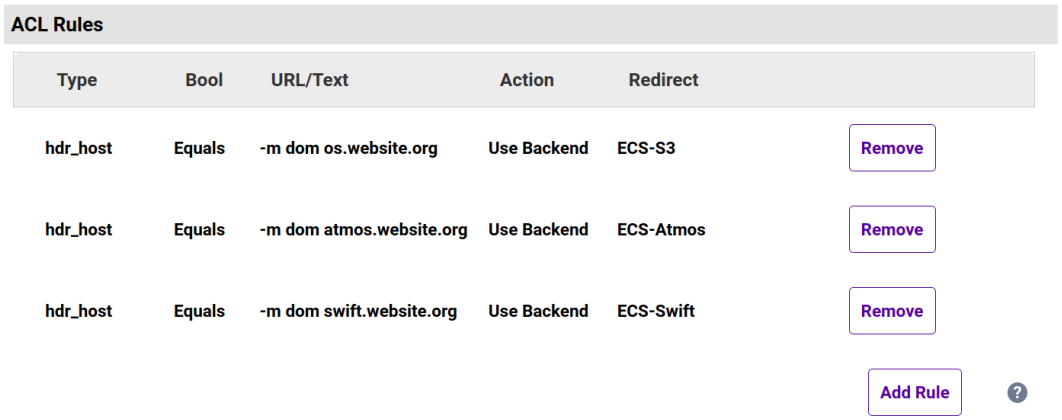

11. Click **Save** to save all of the added ACL rules.
12. Click **Update**.

9.6.2. Setting Up the TLS/SSL Termination

Incoming TLS/SSL encrypted traffic must be decrypted at the load balancer, so that it can then be read as plaintext HTTP traffic. This is necessary to separate the traffic by FQDN using the previously configured ACL rules, i.e. traffic destined for **os.website.org** (and ***.os.website.org**) goes to the S3 virtual service, traffic destined for **atmos.website.org** goes to the Atmos virtual service, and traffic destined for **swift.website.org** goes to the Swift virtual service.

Uploading the Certificate

The appropriate public certificate, including **both** the private key and public certificate parts, must be uploaded to the load balancer for TLS/SSL termination to work.

For more information on creating PEM certificate files and converting between certificate formats please refer to Creating a PEM File

The process for uploading a certificate is as follows:

1. Using the web user interface, navigate to *Cluster Configuration > SSL Certificate* and click on **Add a new SSL Certificate**.
2. Press the *Upload prepared PEM/PFX file* radio button.
3. Define the *Label* for the certificate as required. It may make sense to use the domain that the certificate is associated to, e.g. **website.org**.
4. Click on **Browse** and select the appropriate PEM or PFX style certificate.
5. If uploading a PFX certificate, enter the certificate's password in the *PFX File Password* field.
6. Click **Upload certificate**.

Creating the TLS/SSL Termination

1. Using the web user interface, navigate to *Cluster Configuration > SSL Termination* and click on **Add a new Virtual Service**.
2. From the *Associated Virtual Service* drop-down list, select the 'ECS Combined' service that was created

previously, e.g. **ECS-Combined-Service**.

3. Set the *Virtual Service Port* field to **443**.
4. From the *SSL Certificate* drop-down list, select the certificate for the service in question, which in this example is **website.org**.
5. Click **Update** to create the TLS/SSL termination service.

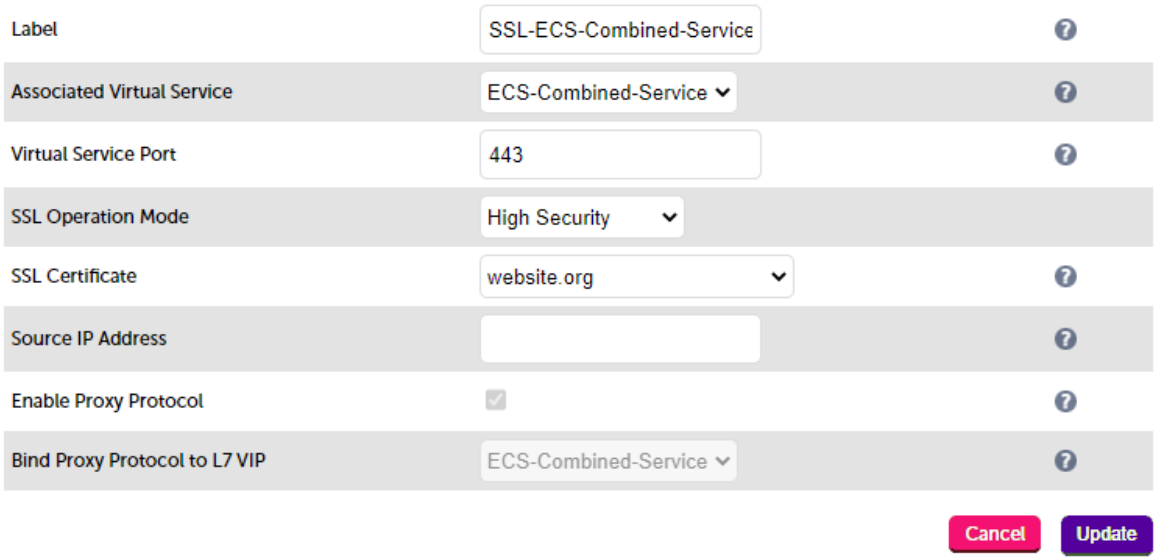

9.6.3. Finalizing the Configuration

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 WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.
3. Click **Reload STunnel**.

10. Testing & Verification

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

10.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the ECS servers) and shows the state/health of each server as well as the state of the cluster as a whole.

The example below shows a **scenario 1** style setup, where all five ECS servers are healthy and available to accept connections for each of the four protocol-specific virtual services:

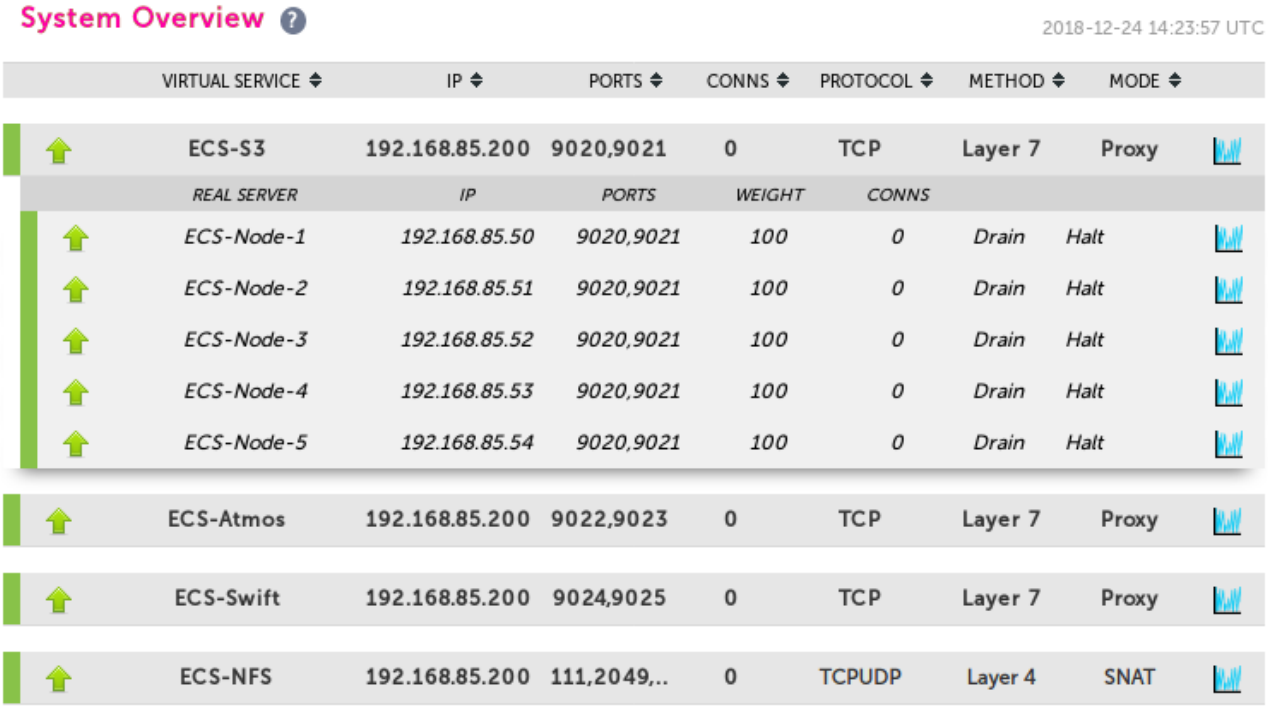

The example below shows a **scenario 2** style setup, where all five ECS servers are healthy and available to accept connections for each of the four protocol-specific virtual services.

Note that the 'ECS Combined Service' shows as red, as it does not have any healthy real servers (because it does not have any real servers defined). This is normal, as it is a 'dummy' service used only to redirect incoming traffic to the other four virtual services, based on the destination domain of incoming traffic.

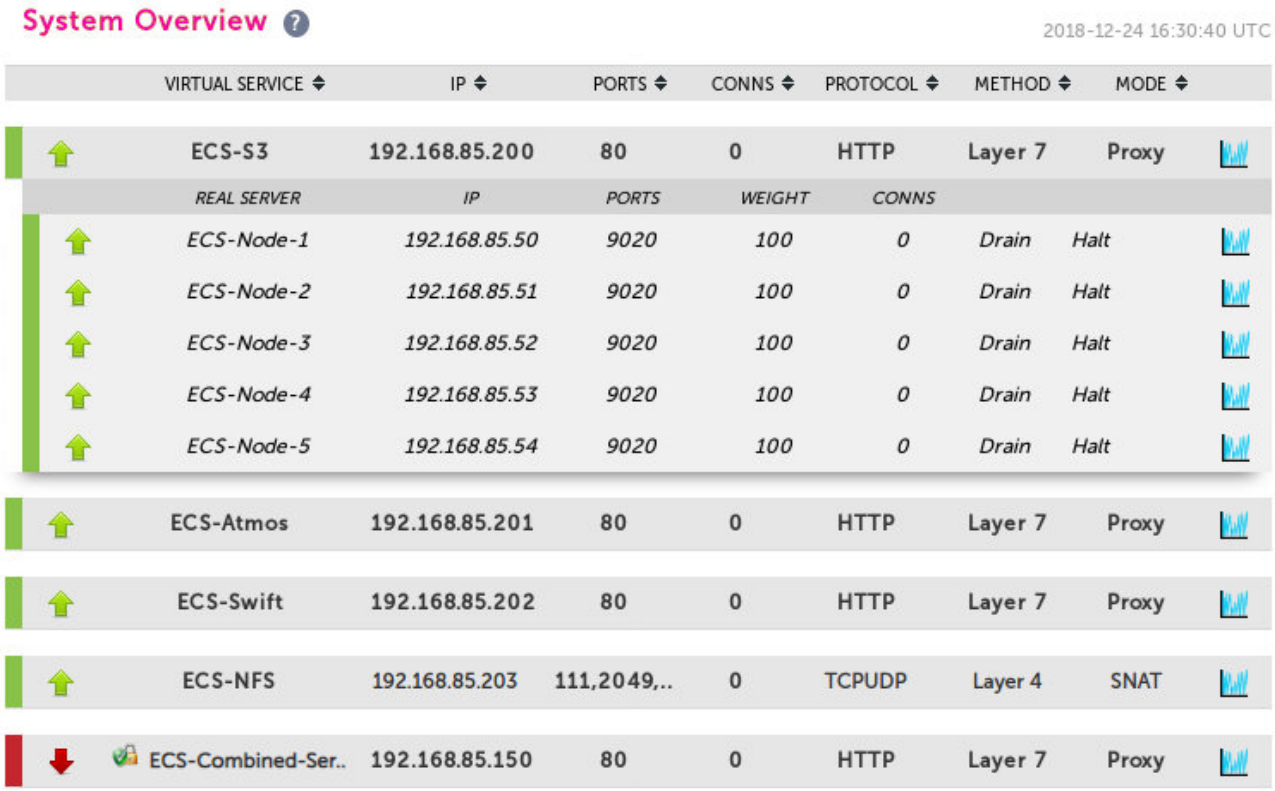

11. Technical Support

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.

12. Further Documentation

For additional information, please refer to the Administration Manual.

13. Appendix

13.1. Multi-port NFS Health Check

A custom health check can be used with the NFS virtual service. This will check the real servers and ensure that all three of the NFS ports (111, 2049, and 10000) are available before considering a server to be online and ready to accept connections. To configure this custom health check, there are 2 steps:

Step 1 - Configure the multi-port health check:

1. Using the WebUI, navigate to *Cluster Configuration > Health Check Scripts* and click **Add New Health Check**.

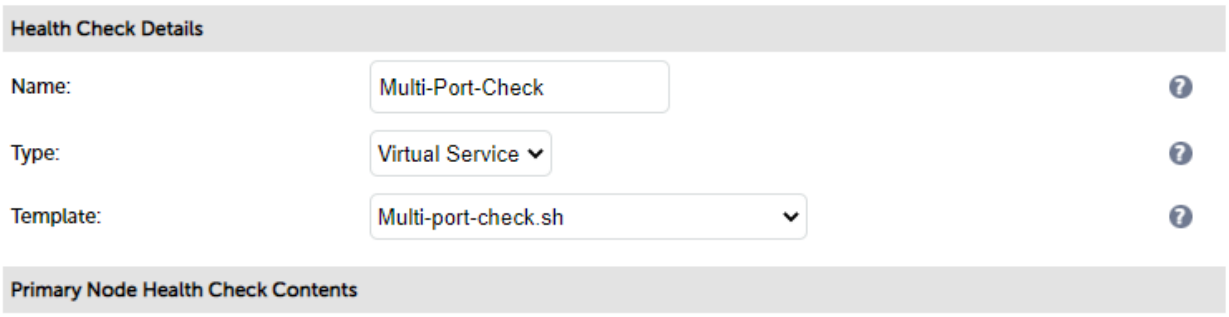

2. Specify an appropriate *Name* for the health check, e.g. **Multi-Port-Check**.
3. Set *Type* to **Virtual Service**.
4. Set *Template* to **Multi-port-check.sh**.
5. Using the editor window, change the CHECK_PORT definition so that it reads:

```
CHECK_PORT="111 2049 10000"
```

6. Click **Update**.

Step 2 - Configure the VIP to use the new health check:

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Modify** next to the NFS Virtual Service.
2. Scroll down to the *Health Checks* section.
   a. Set *Check Type* to **External script**.
   b. Set *External Script* to **Multi-port-check**.
3. Click **Update**.

13.2. 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.

13.2.1. Non-Replicated Settings

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

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	Interface IP addresses, bonding configuration and VLANs
Local Configuration	Routing	Default gateways and static routes
Local Configuration	System Date & time	Time and date related settings
Local Configuration	Physical – Advanced Configuration	Various appliance settings
Local Configuration	Portal Management	Portal management settings
Local Configuration	Security	Security settings
Local Configuration	SNMP Configuration	SNMP settings
Local Configuration	Graphing	Graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Backup & Restore	Local XML backups
Maintenance	Software Updates	Appliance software updates
Maintenance	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.

13.2.2. 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*.

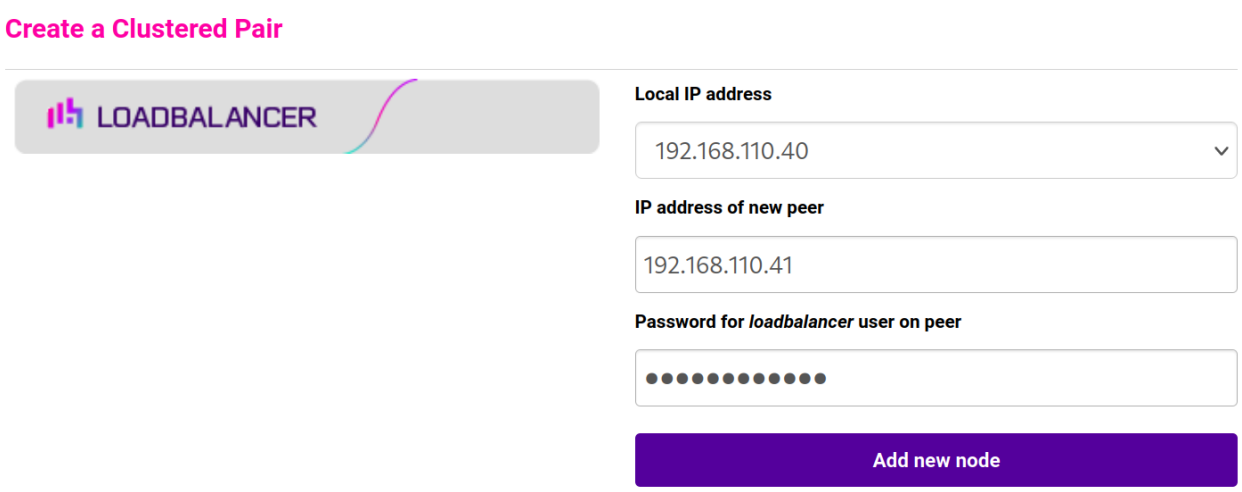

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

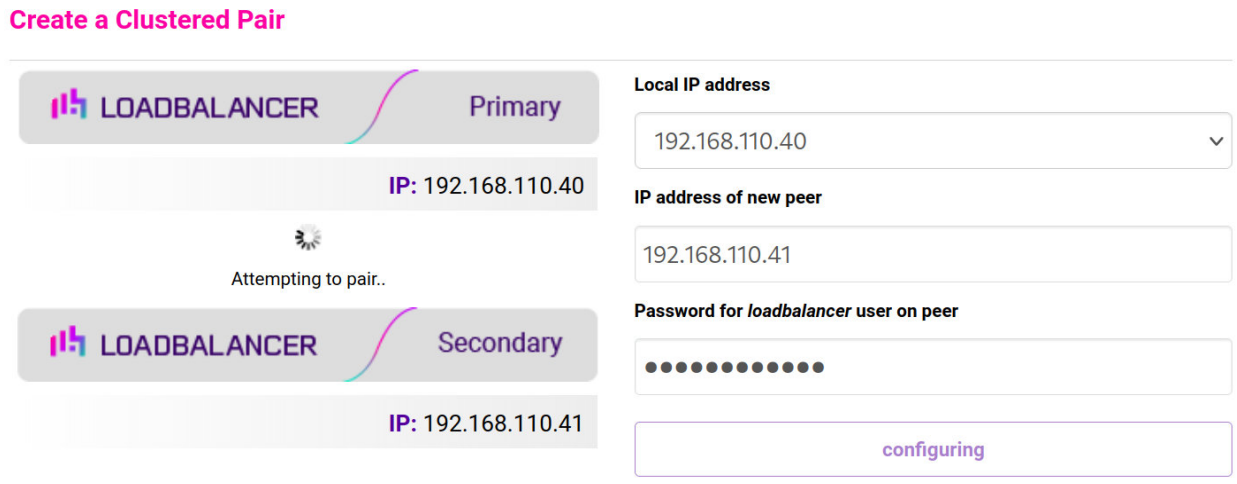

6. Once complete, the following will be displayed on the Primary appliance:

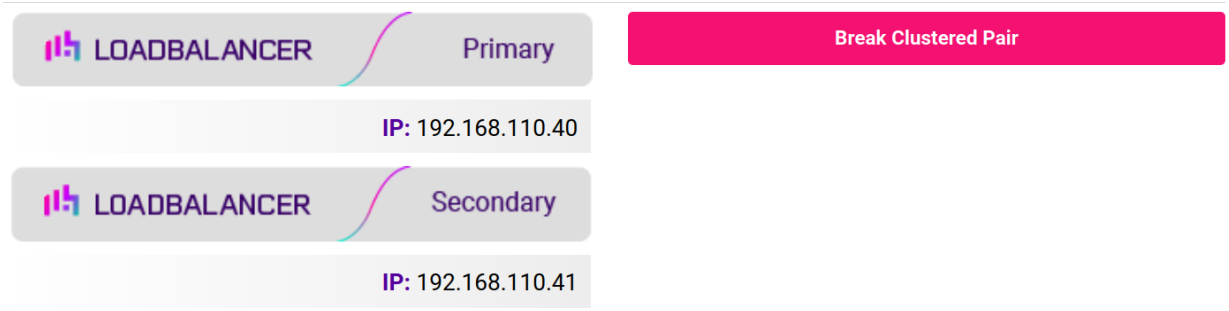

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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.

14. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	24 December 2018	Initial version		AH
1.1.0	30 August 2019	Styling and layout	General styling updates	AH
1.1.1	13 July 2020	Changed the NFS virtual service to be a layer 4 SNAT mode service Added explicit leading forward slashes to the Negotiate HTTP health check 'requests to send'	Some NFS applications require UDP traffic, which needs to be handled using a layer 4 service Reflects a change in the appliance's health check handling	AH
1.1.2	17 July 2020	New title page Updated Canadian contact details Added additional instructions for configuring health checks	Branding update Change to Canadian contact details Changes to the appliance WebUI	AH
1.2.0	1 November 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.2.1	13 April 2022	Updated ACL instructions	Changes to the appliance WebUI	AH
1.2.2	22 April 2022	Updated SSL related content to reflect latest software version	New software release	RJC
1.2.3	11 May 2022	Updated external health check related content to reflect latest software version	New software release	RJC
1.2.4	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH

Version	Date	Change	Reason for Change	Changed By
1.2.5	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
1.2.6	2 February 2023	Updated screenshots	Branding update	AH
1.2.7	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.