LOADBALANCER

MANUALS

Enterprise Quick Start Guide

Version 8.13.5 Revision 1.0.0

Table of Contents

1. About this Guide

This quick start guide provides an introduction to Enterprise and information on how to deploy the appliance, configure initial network settings and configure and verify a simple layer 7 test environment.

> **Note** Please also refer to the Administration Manual for much more detailed information on setting up the appliance and configuring a load balancing solution. For information on configuring the appliance for specific applications, please refer to our extensive library of Deployment Guides.

2. About Enterprise

Enterprise is Loadbalancer.org's first generation load balancer. The core software is based on LBOS-7 which is a customized Linux build maintained by Loadbalancer.org, LVS, Ldirectord, Linux-HA, HAProxy & STunnel. Full root access is provided which enables complete control of all settings.

The appliance is available in the following formats: hardware, virtual (VMware, HyperV, KVM, Nutanix & XEN) and cloud based (Amazon, Azure & GCP).

3. Appliance Configuration Overview

Initial network configuration is carried out at the console using the Network Setup Wizard. Once the wizard has been run, load balanced services can be configured using the WebUI - either using the Setup Wizard which can be used to configure layer 7 Virtual Services (VIPs) and the associated Real Servers (RIPs) or by manually defining the required services.

By default, the WebUI is accessible on HTTPS port **9443**, this can be changed if required. For more information, please refer to the "Appliance Security" section below.

We always recommend that where possible two appliances are deployed as a clustered pair for high availability and resilience, this avoids introducing a single point of failure to your network.

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. For more information on configuring an HA pair, please refer to Configuring HA - Adding a Secondary Appliance.

4. Appliance Security

> **Note** For full details of each security mode and all other security related features, please refer to Appliance Security Features.

4.1. Security Mode

To control how the appliance is accessed and which features are enabled, 3 security modes are provided:

- **Custom** - In this mode the security options can be configured to suit your requirements
- **Secure - (Default)** - In this mode:
  - All console access and SSH password access is disabled
  - WebUI connections are forced to use HTTPS
  - Access to the *Local Configuration > Execute shell command* menu option is disabled
  - The Firewall Script & the Firewall Lockdown Wizard Script cannot be edited
- **Secure - Permanent** - This mode is the same as **Secure** but once set it cannot be changed

> **Important** Setting the security mode to **Secure - Permanent** is irreversible.

To configure the Security Mode:

1. Using the WebUI, navigate to: *Local Configuration > Security*.
2. Select the required *Appliance Security Mode*. If **Custom** is selected, configure the required options.
3. Click **Update**.

4.2. Passwords

The password for the "root" user Linux account and the "loadbalancer" WebUI user account are set during the Network Setup Wizard.

> **Note** The passwords for the cloud products are either set to a default value or are configured during instance deployment. Also, for Enterprise AWS and Enterprise Azure it's not possible to directly log in as "root". For more details, please refer to the relevant Quick Start Configuration Guide.

4.2.1. The "root" Linux account

As explained in Security Mode above, all console access & SSH password access is disabled by default. Once enabled, the "root" user password can be changed at the console or via an SSH session using the following command:

```
# passwd
```

4.2.2. The "loadbalancer" WebUI account

This can be changed using the WebUI menu option: *Maintenance > Passwords*.

5. Deployment Concept

Once deployed, clients connect to the Virtual Service(s) (VIPs) on the load balancer rather than connecting directly to one of the load balanced servers. Requests are then distributed between the load balanced servers according to the load balancing algorithm selected.

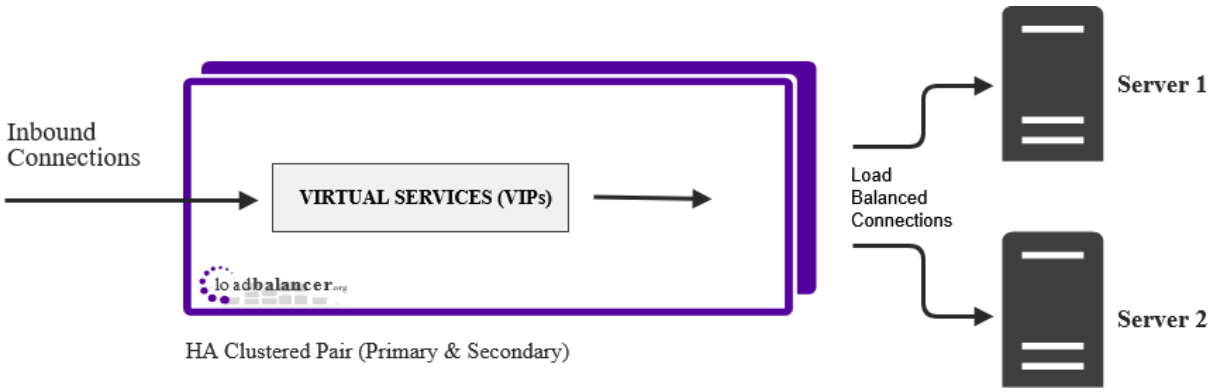

> **Note** We always recommend that two appliances are deployed as an active/passive HA pair. The Secondary appliance automatically takes over if the Primary appliance fails. For more information on configuring HA using two appliances please refer to Configuring HA - Adding a Secondary Appliance.

6. One-Arm and Two-Arm Topologies

The number of "arms" is a descriptive term for how many interfaces are used to connect a device to a network. It's common for a load balancer that uses a routing method (NAT) to have a two-arm configuration although one-arm is also supported. Proxy based load balancers commonly use a one-arm configuration although two-arm is also supported.

6.1. One-Arm

The VIP and the load balanced servers are located in a single subnet. The load balancer requires a single network interface.

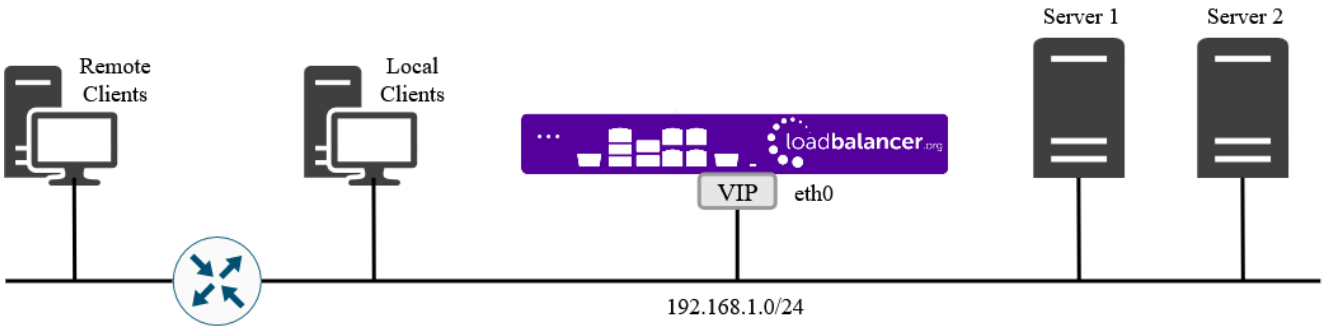

6.2. Two-Arm

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 as shown in the diagram below.

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

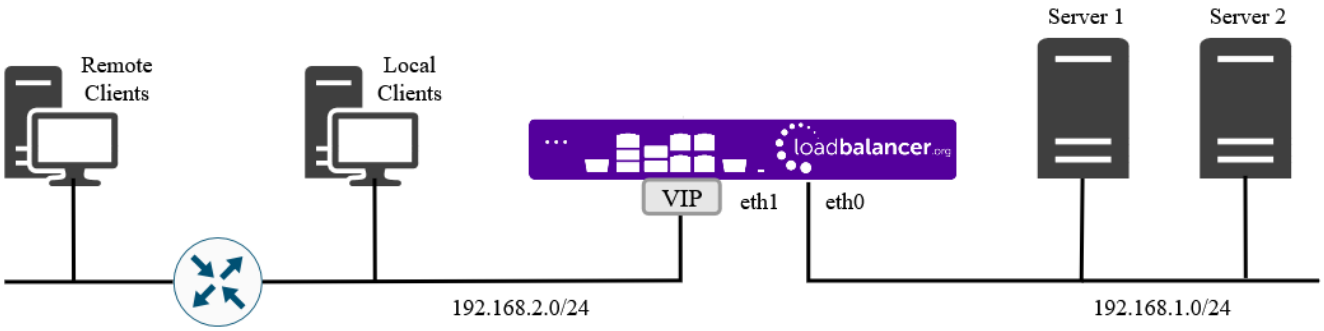

> **Note** Typically **eth0** is used as the internal interface and **eth1** is used as the external interface. This is not a requirement - each interface can be used for any purpose.

7. Supported Load Balancing Methods

The Loadbalancer.org appliance supports multiple load balancing methods/modes. These can be used at the same time or in combination with each other and are described in the table below.

Layer	Method/Mode	Comments	Topology	Note
Layer 4	DR Mode (*Direct Routing*)	Ultra-fast layer 4 load balancing • Requires the "ARP problem" to be solved on each Real Server - for details see DR Mode Considerations	One-Arm (*)	1
Layer 4	NAT Mode (*Network Address Translation*)	Fast Layer 4 load balancing • The appliance must be the default gateway for the Real Servers	One-Arm or Two-Arm	1
Layer 4	TUN Mode (*Tunneling*)	Similar to DR mode but works across IP encapsulated tunnels • Requires the "ARP problem" to be solved on each Real Server - for details see DR Mode Considerations	One-Arm	2
Layer 4	SNAT Mode (*Source Network Address Translation*)	Fast layer 4 load balancing • Requires no Real Server configuration changes	One-Arm or Two-Arm	3
Layer 7	SNAT Mode (*Source Network Address Translation*)	Not as fast as layer 4 methods, but offers greater flexibility, supports remote server load balancing and advanced functionality such as multiple persistence methods, header manipulation and URL rewriting • Based on HAProxy • Requires no Real Server configuration changes	One-Arm or Two-Arm	4

Layer	Method/Mode	Comments	Topology	Note
Layer 7	SSL/TLS Termination	Typically required to enable cookie persistence, header manipulation and URL rewriting in HTTPS streams • STunnel, Pound or HAProxy can be used as the terminator	One-Arm or Two-Arm	5

(*) DR mode can also be used in a multi-homed configuration where Real Servers are located in different subnets. In this case, the load balancer must have an interface in each subnet to enable layer 2 connectivity which is required for DR mode to operate.

Notes

1. Recommended for high performance fully transparent and scalable solutions.
2. Only required for DR mode implementations across routed networks (rarely used).
3. Useful when you want to load balance both TCP and UDP but you're unable to use DR mode or NAT mode due to network topology or Real Server related reasons.
4. Used across multiple environments including object storage, healthcare and various Microsoft application such as Exchange, IIS and RDS.
5. SSL/TLS termination is processor intensive - where possible, for a scalable solution terminating on the Real Servers is recommended.

8. Ports Used by the Appliance

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

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

> **Note** All ports listed above except port 123 (NTP) can be changed if required.

- To change the port used for heartbeat, refer to Configuring High Availability
- To change the port used for HAProxy replication, refer to Layer 7 - Advanced Configuration
- To change other ports, refer to Service Socket Addresses

9. Appliance Deployment

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. Hardware Appliance

For details of installing and connecting the appliance, please refer to Hardware Appliance Installation.

9.3. Cloud Appliances

9.3.1. AWS

For details of deploying and configuring the Amazon Web Services appliance, please refer to the AWS Configuration Guide.

9.3.2. Azure

For details of deploying and configuring the Microsoft Azure appliance, please refer to the Azure Configuration Guide.

9.3.3. GCP

For details of deploying and configuring the Google Cloud Platform appliance, please refer to the GCP Configuration Guide.

10. Configuring Initial Network Settings

After power up, the following startup message is displayed on the appliance console:

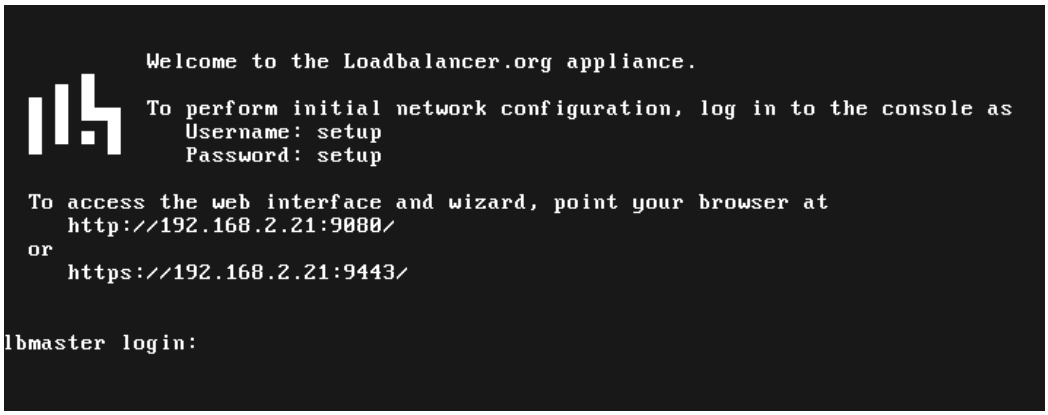

As mentioned in the text, to perform initial network configuration, login as the "setup" user at the appliance console.

Once logged in, the Network Setup Wizard will start automatically. This will enable you to configure the management IP address and other network settings for the appliance.

login to the console:

Username: setup
Password: setup

A series of screens will be displayed that allow network settings to be configured:

In the **Configure Management IP** screen, leave **Yes** selected and hit **Enter** to continue.

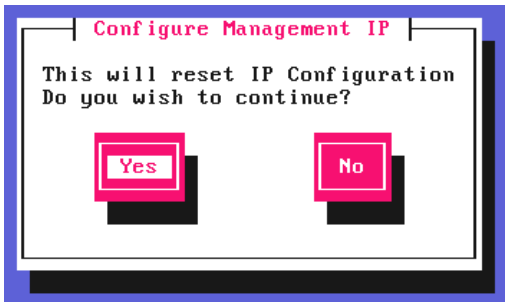

In the **Peer Recovery** screen, leave **No** selected and hit **Enter** to continue.

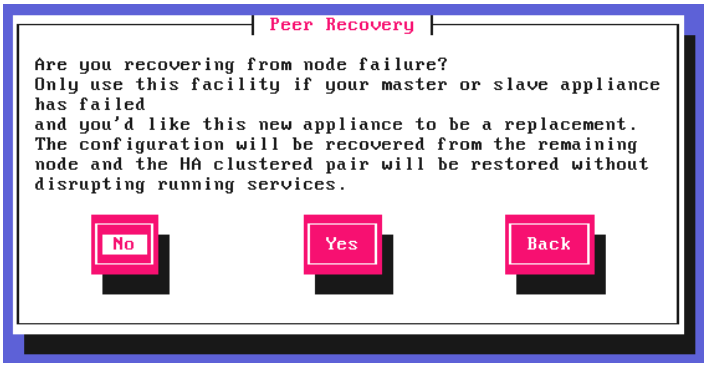

> **Note** For more details on node recovery using this option please refer to Disaster Recovery After Node (Primary or Secondary) Failure.

In the **Centralized Management** screen, if you would like to enroll the appliance with a management server (typically **portal.loadbalancer.org**), select **Yes**, otherwise leave **No** selected, then hit **Enter** to continue. If you select **Yes**, you'll be asked to confirm the server's details and provide login credentials at the end of this setup process.

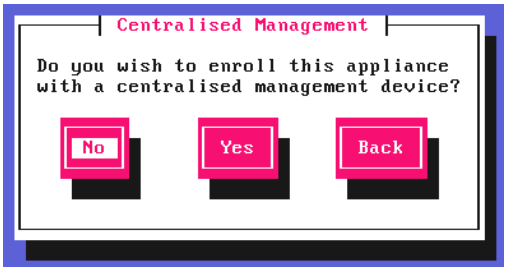

> **Note** For information on how to modify Centralized Management settings via the WebUI, please refer to Portal Management & Appliance Adoption.

In the **Available Interfaces** screen, a list of available interfaces will be displayed, hit **Enter** to continue.

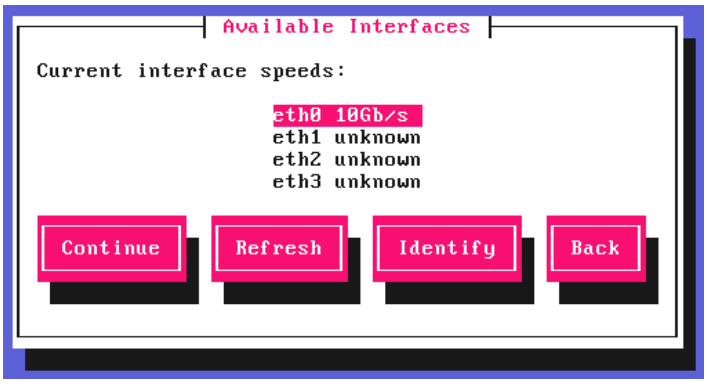

In the **Configure Bonding** screen, select **Yes** if you want to configure a bonded interface, if not leave **No** selected, then hit **Enter** to continue.

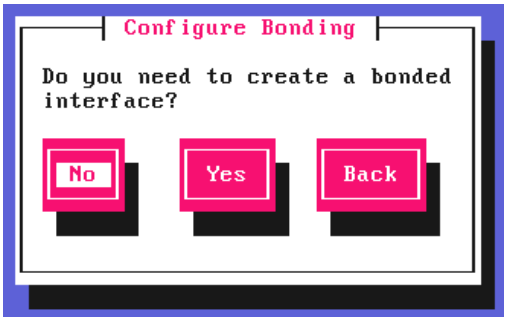

If you select **Yes**, the **Select Interfaces** screen will be displayed. Using the space bar, select the interfaces you'd like to include in the bond, select **Create** and hit **Enter** to continue.

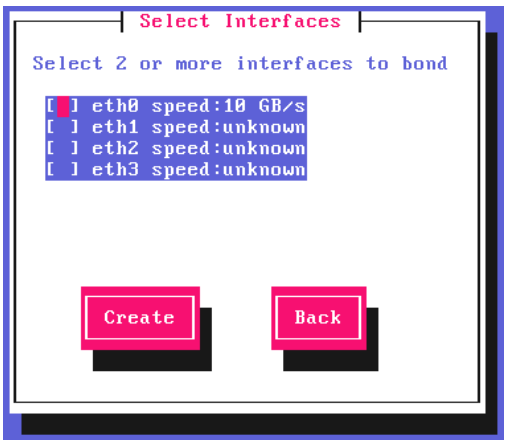

You'll then be prompted to select the bond mode. The following options are available:

Mode 1 - Active/Backup. This places one of the adapters in a backup state and will only become active if the link is lost to the active adapter. This mode provides fault tolerance.

Mode 4 - 802.3ad. Dynamic link aggregation mode. This mode requires a switch that supports IEEE 802.3ad.

In the **Configure a VLAN** screen, select **Yes** if you want to configure a VLAN, if not leave **No** selected, then hit **Enter** to continue.

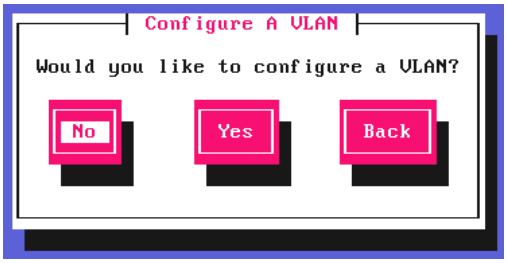

If you select **Yes** you'll be prompted to enter a VLAN Tag ID.

In the **Configure Management IP** screen, select the interface that'll be used to manage the appliance, then hit **Enter** to continue.

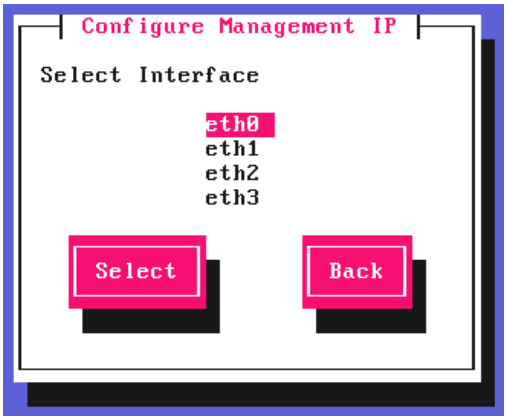

In the **Set IP address** screen, either enter the required *Static IP Address* & *CIDR Prefix* and select **Done** or select **Use DHCP**, then hit **Enter** to continue.

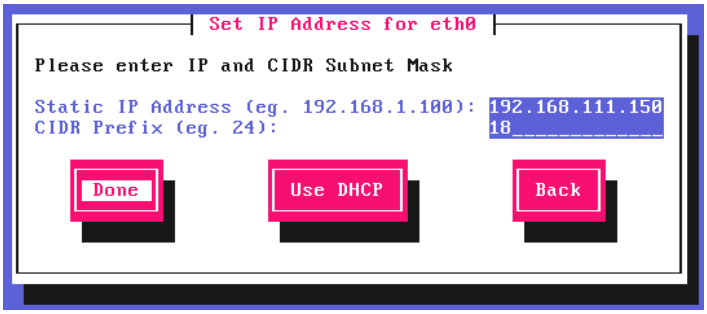

> **Note** A subnet mask such as 255.255.255.0 is not valid, in this case enter 24 instead.

In the **Configure Default Gateway** screen, enter the required *Default Gateway IP Address*, select **Done** and hit **Enter** to continue.

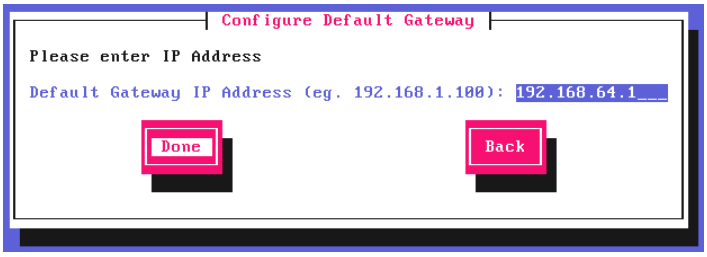

In the **Configure DNS Servers** screen, configure the required DNS server(s), select **Done** and hit **Enter** to continue.

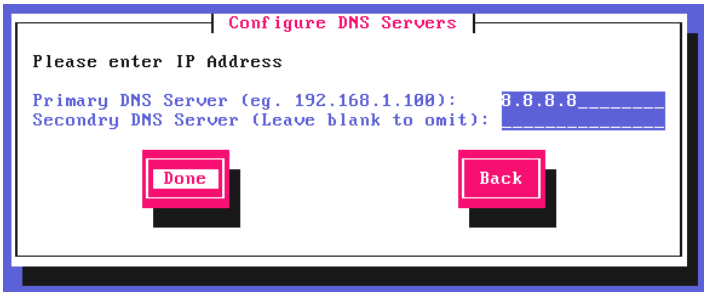

In the **Set Password** screen, hit **Enter** to continue.

Enter the *Password* you'd like to use for the **loadbalancer** WebUI user account and the **root** Linux user account. Repeat the password, select **Done** and hit **Enter** to continue.

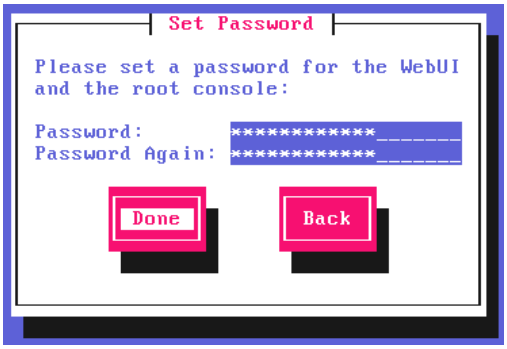

If you selected **Yes** when asked if you want to enroll for Centralized Management, you'll now be prompted for the details. Default values for the *Host* and *Port* are set and can be changed if required. Enter the *Username* and *Password* for the management server account you'd like the appliance to be associated with, select **Done** and hit **Enter** to continue.

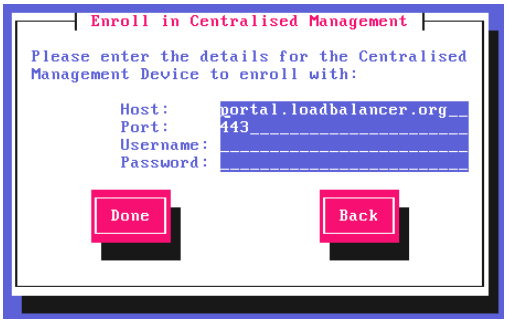

In the **Summary** screen, check all settings. If everything is correct, leave **Configure** selected and hit **Enter** to continue. All settings will be applied. If you need to change a setting, use the **Back** button.

> **Note** For v8.13.2 and later, once the settings have been applied the appliance will check if a software update is available. If an update is found, it will be installed automatically.

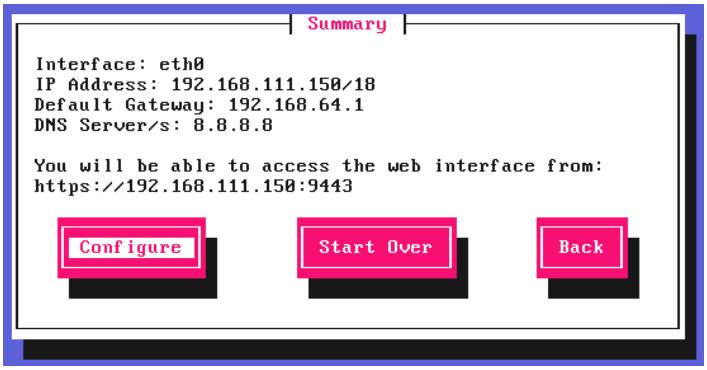

Once the configuration has been written, the **Configuration Complete** screen and message will be displayed. Click **OK** to exit the wizard and return to the command prompt.

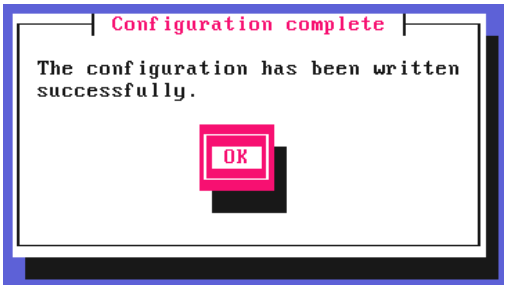

11. Accessing the Appliance WebUI

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:

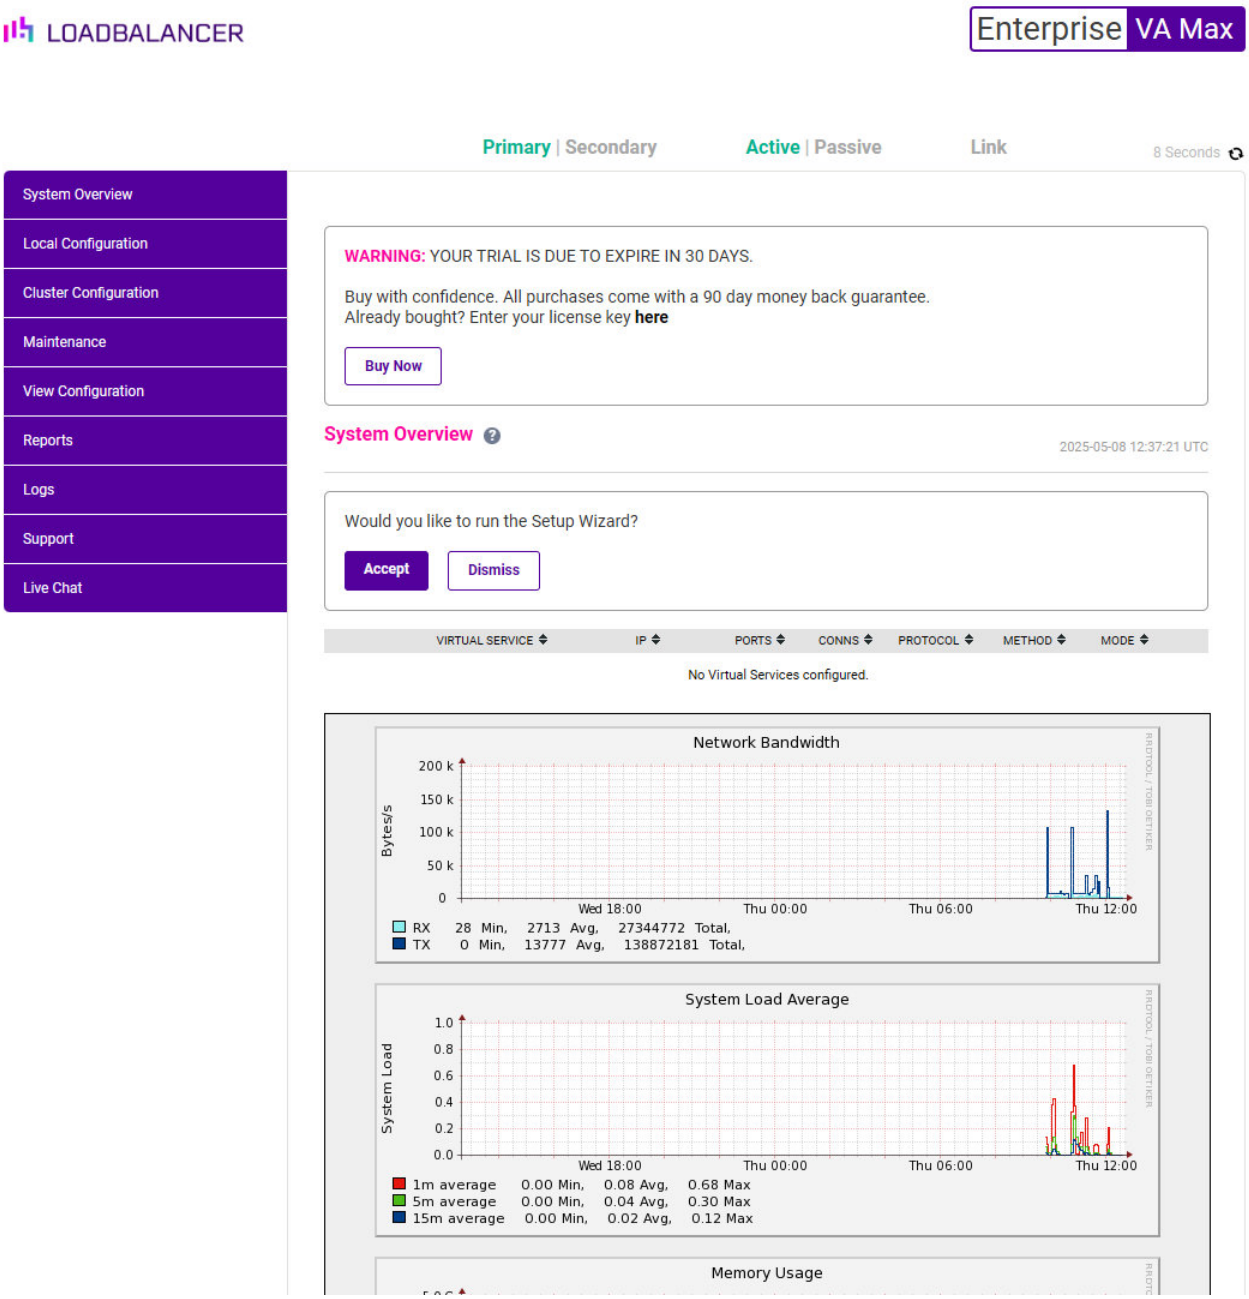

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

11.1. 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

12. Installing the License Key

The appliance can be used completely unrestricted for 30 days without installing a license key. After 30 days, the appliance continues to work but it's no longer possible to make configuration changes.

> **Note** if you're conducting a PoC (Proof of Concept) using the VA and require more time to complete your evaluation, please contact sales@loadbalancer.org who will be able to provide guidance on how to extend the trial.

For an unlicensed VA, the following message is displayed:

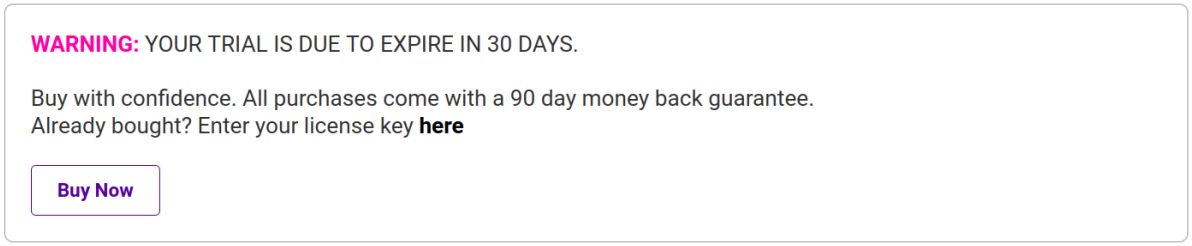

For an unlicensed hardware appliance, the following message is displayed:

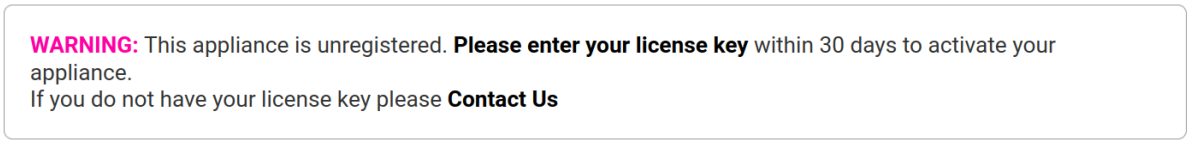

To install the license key:

1. Using the WebUI, navigate to: *Local Configuration > License Key*.

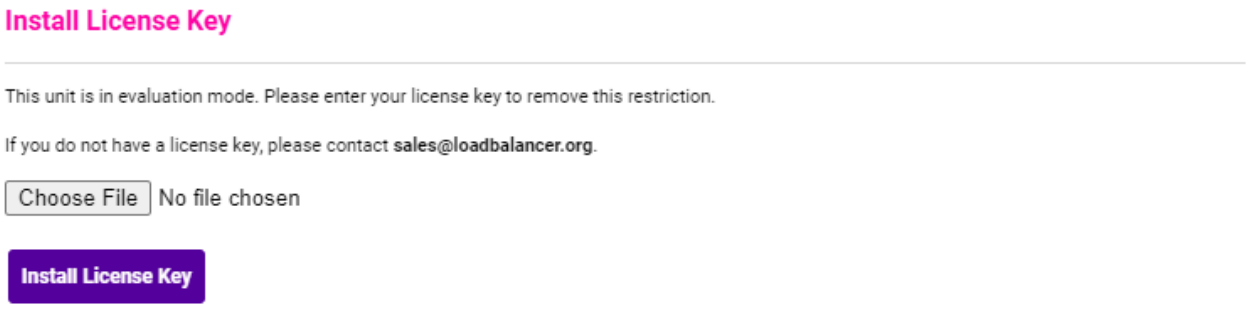

2. Click **Choose File** then browse to and select the license file provided when the appliance was purchased.

3. Click **Install License Key**.

> **Note** Once the license is applied, these warning messages will no longer be displayed.

13. 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.

13.1. 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:

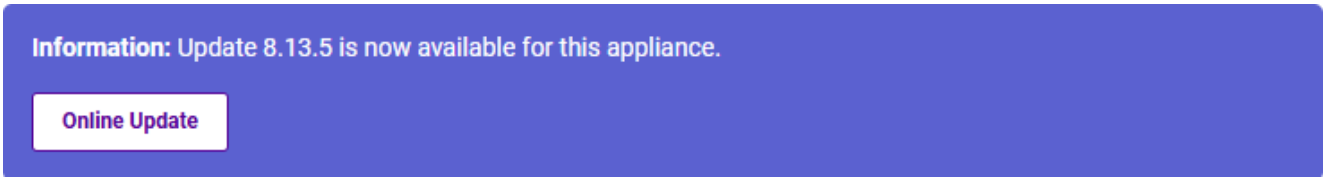

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

> **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:

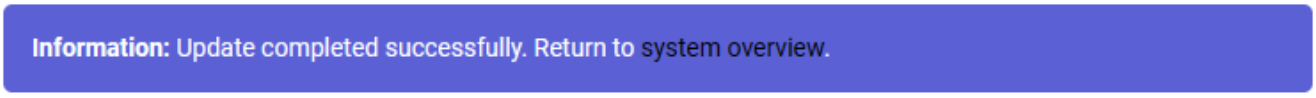

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

13.2. 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:

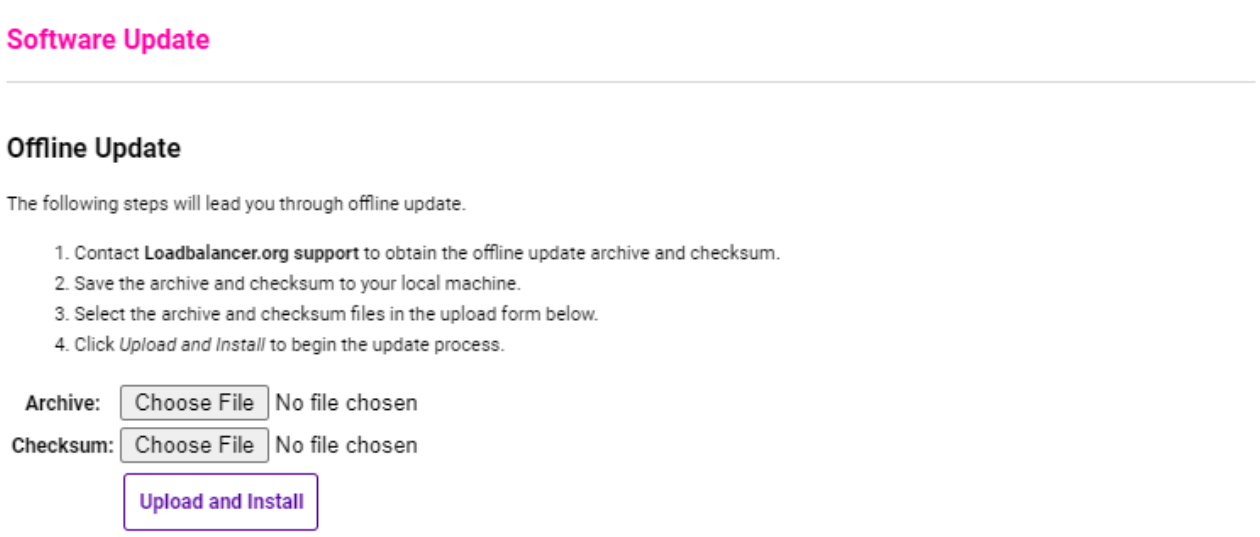

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

14. Configuring & Testing a Simple Load Balanced Test Environment

This configuration example illustrates how to configure a simple layer 7 load balanced test environment.

> **Note** Layer 7 SNAT mode is used in the example. As mentioned earlier, this is not the fastest mode but is very simple to deploy and requires no mode specific configuration changes to the Real Servers.

The following table & diagram describe the environment:

IP Address	Device	Notes
192.168.1.10	Test Client	
192.168.1.20	Load Balancer	the load balancer's own IP address
192.168.1.25	Load Balancer	the Virtual IP address (VIP), the IP address clients connect to
192.168.1.30	Web Server 1	the first Real Server (RIP)

IP Address	Device	Notes
192.168.1.40	Web Server 2	the second Real Server (RIP)

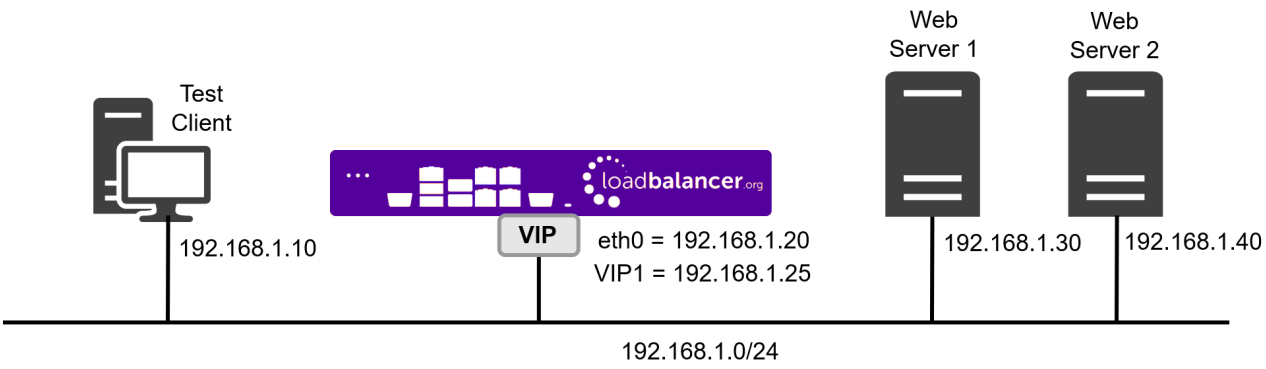

14.1. STEP 1 - Deploy the Appliance

Please refer to Appliance Deployment.

14.2. STEP 2 - Run the Network Setup Wizard

Please refer to Configuring Initial Network Settings.

14.3. STEP 3 - Configure the Virtual Service (VIP) & Associated Real Servers (RIPs)

14.3.1. Virtual Service Configuration

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

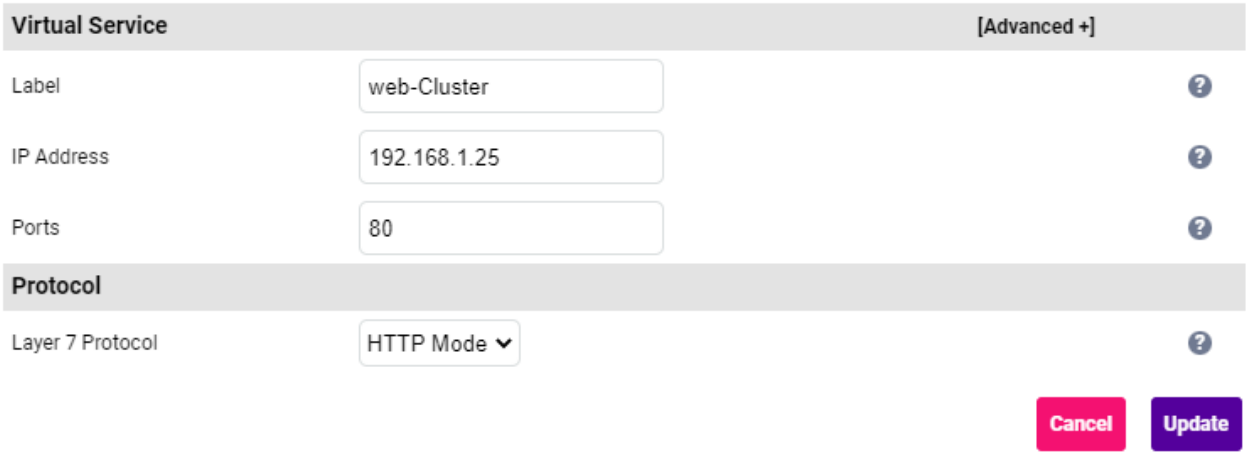

- Specify an appropriate *Label* (name) for the Virtual Service, e.g. **Web-Cluster**.
- Set the *IP Address* field to the required address, e.g. **192.168.1.25**.
- Set the *Ports* field to the required port, e.g. **80**.

- Leave the *Protocol* set to **HTTP Mode**.

3. Click **Update** to create the Virtual Service.

14.3.2. Real Server Configuration

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

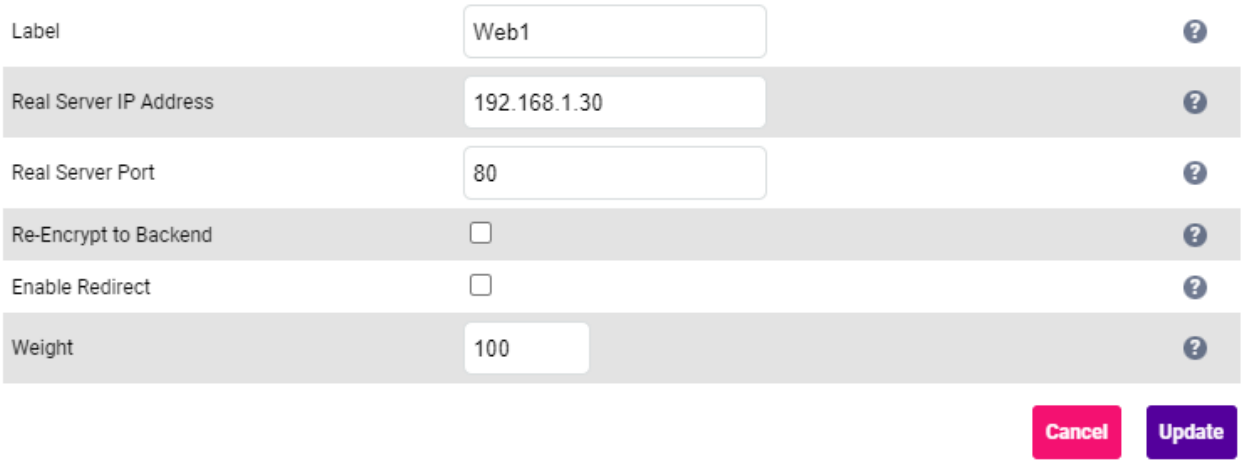

- Specify an appropriate *Label* for the RIP, e.g. **Web1**.
- Set the *Real Server IP Address* field to the required address, e.g. **192.168.1.30**.
- Set the *Real Server Port* field to the required port, e.g. **80**.

3. Click **Update** to add the Real Server.
4. Now repeat these steps to add the second Real Server, e.g. **Web2**.

> **Note** By default, Real Server health-checks are set to use a TCP port connect. If you need a more robust check, this can be changed by modifying the configuration as explained below. For more information, please refer to Real Server Health Monitoring & Control.

14.4. STEP 4 - Finalizing the Configuration

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

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

14.5. STEP 5 - Viewing & Modifying the Configuration

1. The VIP can be viewed using the WebUI menu option: *Cluster Configuration > Layer 7 - Virtual Services* as shown below:

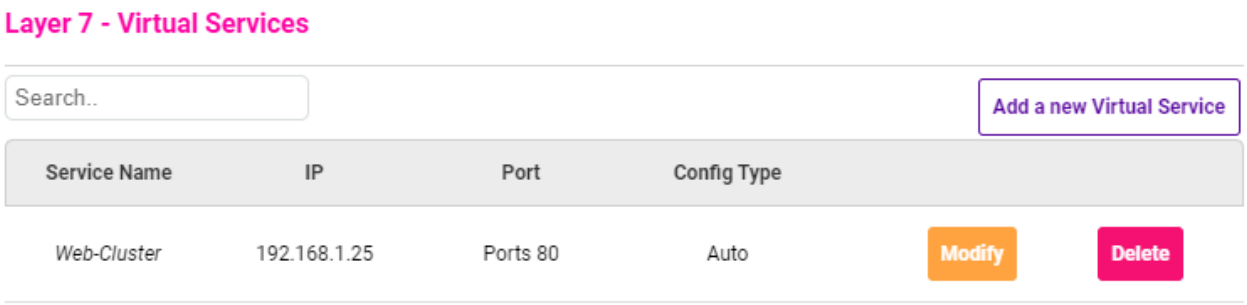

- Clicking the **Modify** button allows all VIP settings to be modified.
- If changes are made, click the **Update** button to save the changes, then use the **Reload HAProxy** button in the "Commit changes" box at the top of the screen to apply the changes.

2. The RIP(s) can be viewed using the WebUI menu option: *Cluster Configuration > Layer 7 - Real Services* as shown below:

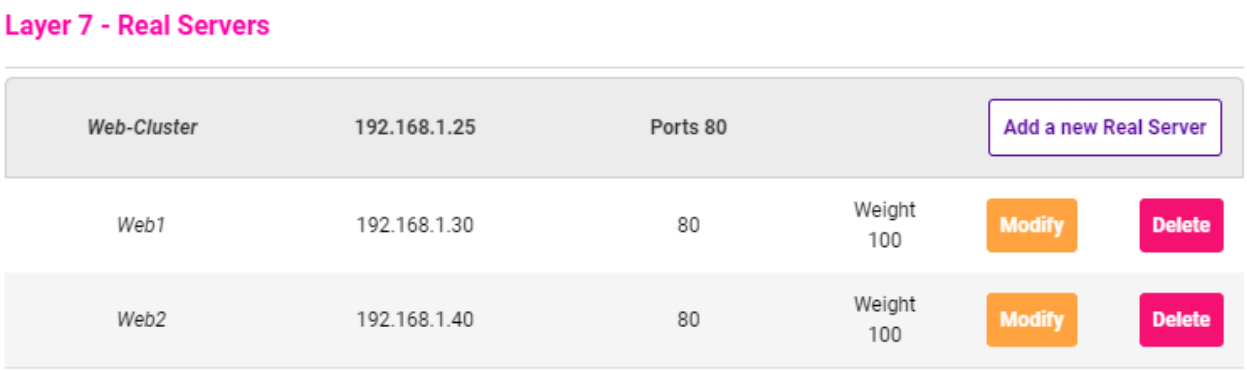

- Clicking the **Modify** button allows all RIP settings to be modified.
- If changes are made, click the **Update** button to save the changes, then use the **Reload HAProxy** button in the "Commit changes" box at the top of the screen to apply the changes.

14.6. STEP 6 - Checking the Status using System Overview

1. Using the WebUI, navigate to: *System Overview* to view the newly created VIP & RIPs. Green indicates that the associated RIPs are passing their health checks:

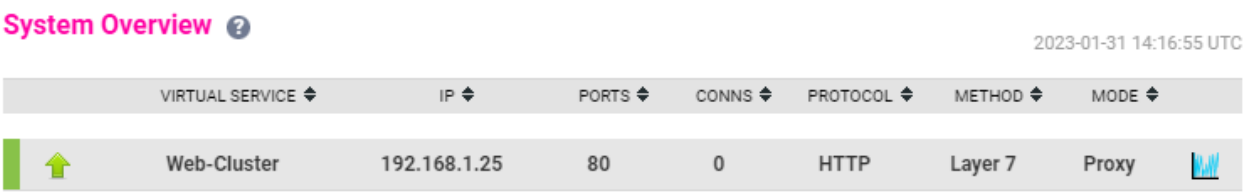

2. Click anywhere on the VIP's horizontal grey area to expand the VIP and view the RIPs:

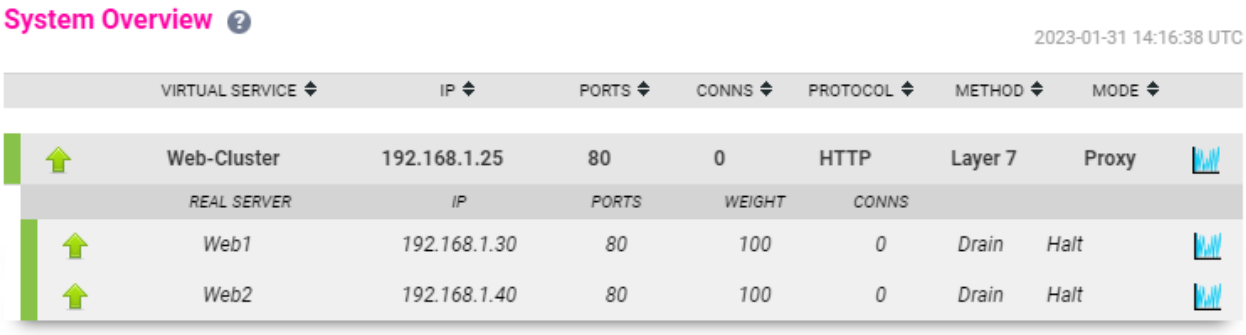

14.7. STEP 7 - Verification & Testing

1. Using the System Overview, verify that both Real Servers are up.

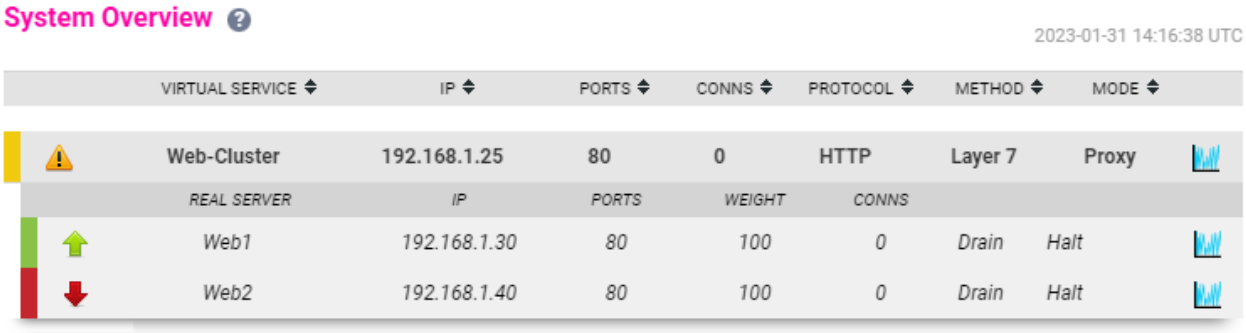

2. In the example above, Web2 is failing its health-check. This should be investigated and corrected - possible steps include:
   - Verify that the application/service is running on the Real Server.
   - Make sure you can ping the Real Server from the load balancer - either from the console, via an SSH session or using the WebUI menu option: *Local Configuration > Execute Shell Command*.

> **Note**
>
> To enable shell commands to be run from the WebUI, the appliance Security Mode mode must be set to **Custom**:
>
> 1. Using the WebUI, navigate to: *Local Configuration > Security*.
> 2. Set *Appliance Security Mode* to **Custom**.
> 3. Click **Update**.

> **Note**
>
> If you run ping from the WebUI, use the form:
>
> ```
> ping -c 4 192.168.1.40
> ```
>
> The **-c 4** means ping 4 times then stop.

   - Verify that the application/service is up and available when accessed from the load balancer - various methods can be used:

- Using **telnet** at the console or via an SSH session:

```
telnet 192.168.1.40 80
```

The following shows a successful connection to port 80:

```
Trying 192.168.1.40...
Connected to 192.168.1.40.
Escape character is '^]'.
```

- Using **nmap** at the console, via an SSH session or using the WebUI option: *Local Configuration > Execute Shell Command*:

```
nmap 192.168.1.40
```

The following is displayed for a working server:

```
Starting Nmap 5.51 ( http://nmap.org ) at 2022-10-25 14:18 UTC
Nmap scan report for 192.168.1.40
Host is up (0.00056s latency).
Not shown: 995 filtered ports
PORT     STATE SERVICE
80/tcp   open  http
135/tcp  open  msrpc
443/tcp  open  https
445/tcp  open  microsoft-ds
3389/tcp open  ms-term-serv
MAC Address: 00:50:56:82:0B:D3 (VMware)

Nmap done: 1 IP address (1 host up) scanned in 4.87 seconds
```

This shows that the server is listening on port 80.

- Using **curl** at the console, via an SSH session or using the WebUI option: *Local Configuration > Execute Shell Command*:

```
curl http://192.168.1.40
or
curl http://host.mydomain.com
```

For a working web server listening on port 80, the default page is returned.

3. Once both servers are up (shown green) browse to the VIP address and verify that you see the web page from each Real Server:
   - Halt Web1 using the *Halt* option for Web1 in the System Overview and verify that content is served by Web2 on a browser refresh.

- Bring Web1 back online using the ***Online (Halt)*** option for Web1, then halt Web2 and verify that content is served by Web1 on a browser refresh.

> **Note** For more configuration examples using Layer 7 SNAT mode as well as other modes, please refer to Configuration Examples.

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

15. 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.

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

15.1. 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

WebUI Main Menu Option	Sub Menu Option	Description
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.

15.2. Configuring the HA Clustered Pair

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

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

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

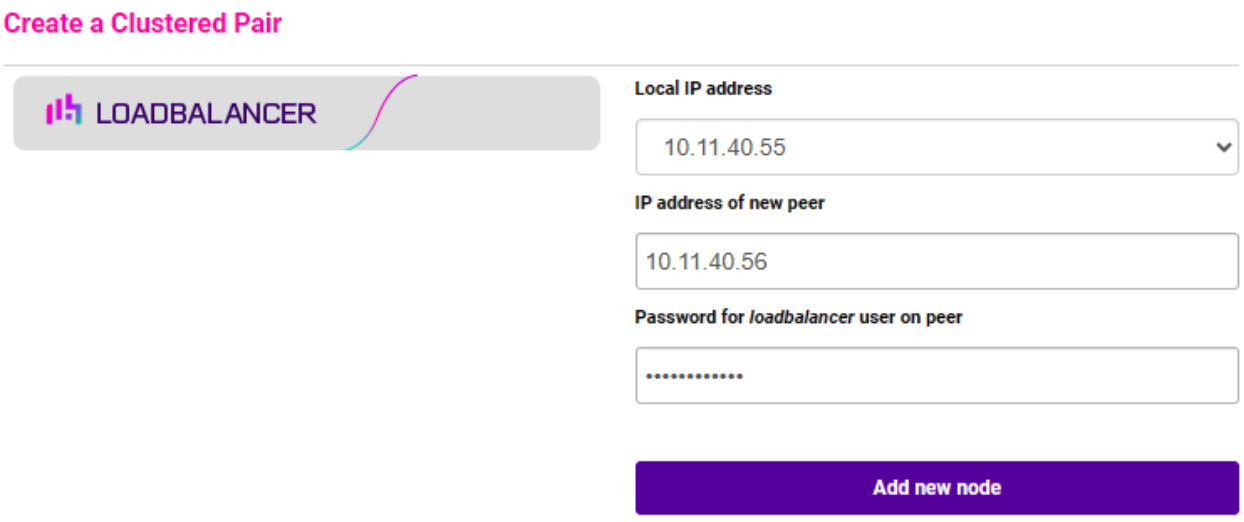

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

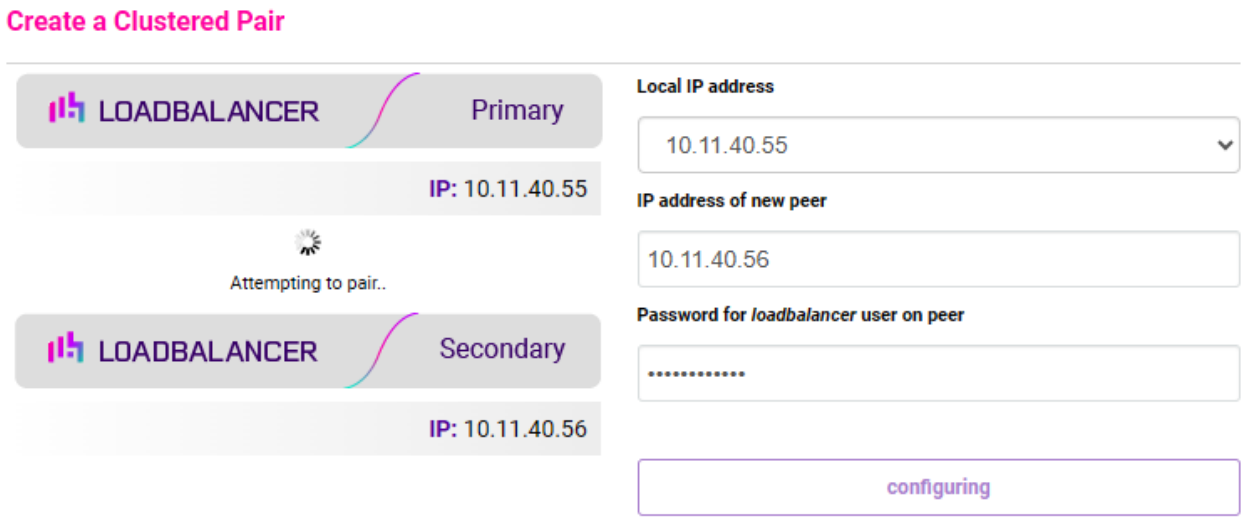

6. Once complete, the following will be displayed on the Primary appliance:

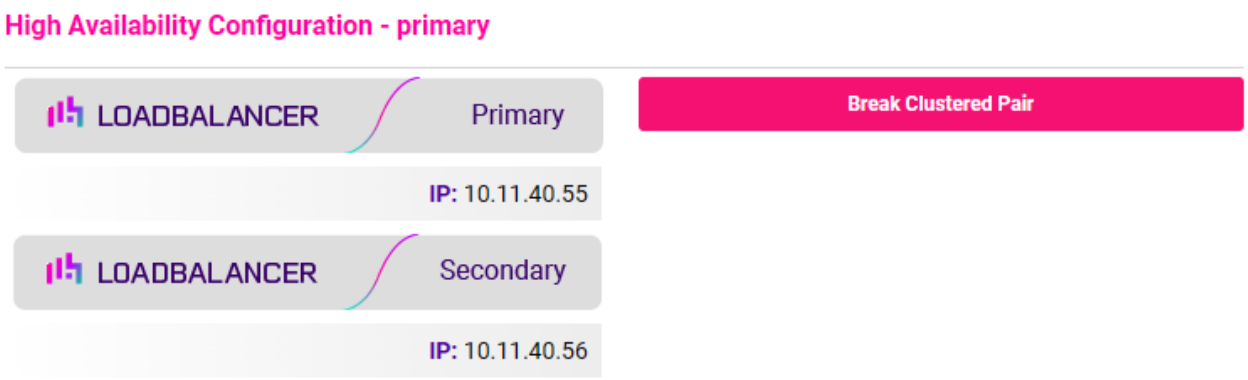

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **Note** Clicking the **Restart Heartbeat** button on the Primary appliance will also automatically restart heartbeat on the Secondary appliance.

> **Note** For more details on configuring HA with 2 appliances, please refer to Configuring High Availability.

> **Note** For details on testing and verifying HA, please refer to Clustered Pair Diagnostics.

16. More Information

Please refer to our website for all the latest Manuals and Deployment Guides.

17. Loadbalancer.org Technical Support

Our highly experienced Support Engineers are on hand to help 24 hours a day, 365 days a year.

17.1. Contacting Support

If you have any questions regarding the appliance or need assistance with load balancing your application, please don't hesitate to contact support@loadbalancer.org.

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.