



Appliance Quick Start Guide

Version 8.7.0 Revision 1.0.0

Table of Contents

1. About this Guide	3
2. Appliance Configuration Overview	3
3. Appliance Security	3
Security Mode	3
Passwords	4
4. Deployment Concept	4
5. One-Arm and Two-Arm Topologies	5
6. Supported Load Balancing Methods	5
7. Ports Used by the Appliance	7
8. Appliance Deployment	7
Virtual Appliance	7
Hardware Appliance	7
Cloud Appliance	7
AWS	7
Azure	7
Google Cloud Platform	8
9. Configuring Initial Network Settings	8
10. Accessing the WebUI	12
Main Menu Options	13
11. Licensing	14
12. Software Updates	14
13. Configuring & Testing a Simple Load Balanced Test Environment	14
STEP 1 - Deploy the Appliance	15
STEP 2 - Run the Network Setup Wizard	15
STEP 3 - Configure the Virtual Service (VIP) & Associated Real Servers (RIPs)	15
VIP Configuration	15
RIP Configuration	16
STEP 4 - Finalizing the Configuration	16
STEP 5 - Viewing & Modifying the Configuration	16
STEP 6 - Checking the Status using System Overview	17
STEP 7 - Verification & Testing	17
14. Configuring HA - Adding a Secondary Appliance	18
Non-Replicated Settings	19
15. More Information	21
16. Loadbalancer.org Technical Support	21

1. About this Guide

This quick start guide provides enough information to deploy the appliance, configure a simple load balanced test environment and test and verify its functionality.

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. 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 (for Layer 7 services) or by manually defining the Virtual Services (VIPs) and associated Real Servers (RIPs).

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 should be fully configured first, then the Secondary appliance should be added to create an HA pair. Once the HA pair is configured, load balanced services will be automatically synchronized from the Primary to Secondary appliance. Load balanced services should then be configured & modified on the Primary appliance and the Secondary will be automatically kept in sync. For more information on configuring an HA pair, please refer to [Configuring HA - Adding a Secondary Appliance](#).

3. Appliance Security

Note

For full details of all security related features, please refer to [Appliance Security Features](#).

Security Mode

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

- **Secure - (default)** - in this mode:
 - the WebUI is accessible on HTTPS port **9443**. If you attempt to access the WebUI on HTTP port **9080** you will be redirected to HTTPS port **9443**
 - access to the *Local Configuration > Execute shell command* menu option is disabled
 - the ability to edit the firewall script & the lockdown wizard is disabled
 - 'root' user console & SSH password access are disabled
- **Custom** - in this mode, the security options can be configured to suit your requirements
- **Secure - Permanent** - this mode is the same as **Secure** but the change is *irreversible*

Important

Only set the security mode to **Secure - Permanent** if you are 100% sure this is what you want!

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 additional options according to your requirements.
3. Configure the *HTTPS Port for Web User Interface*, *Web Interface SSL Certificate* and *Ciphers to use* according to your requirements.
4. Click **Update**.

Passwords

The password for the 'loadbalancer' WebUI user account and the 'root' Linux user account are set during the Network Setup Wizard. These can be changed at any time.

1 - the 'root' Linux account

As explained above, 'root' user console & SSH password access are disabled by default. If enabled, the 'root' password can be changed at the console, or via an SSH session using the following command:

```
# passwd
```

Note

For the AWS and Azure cloud products it's not possible to directly login as root. If root access is required, once you've logged into the console/SSH session using the credentials defined during instance deployment, run the following command:

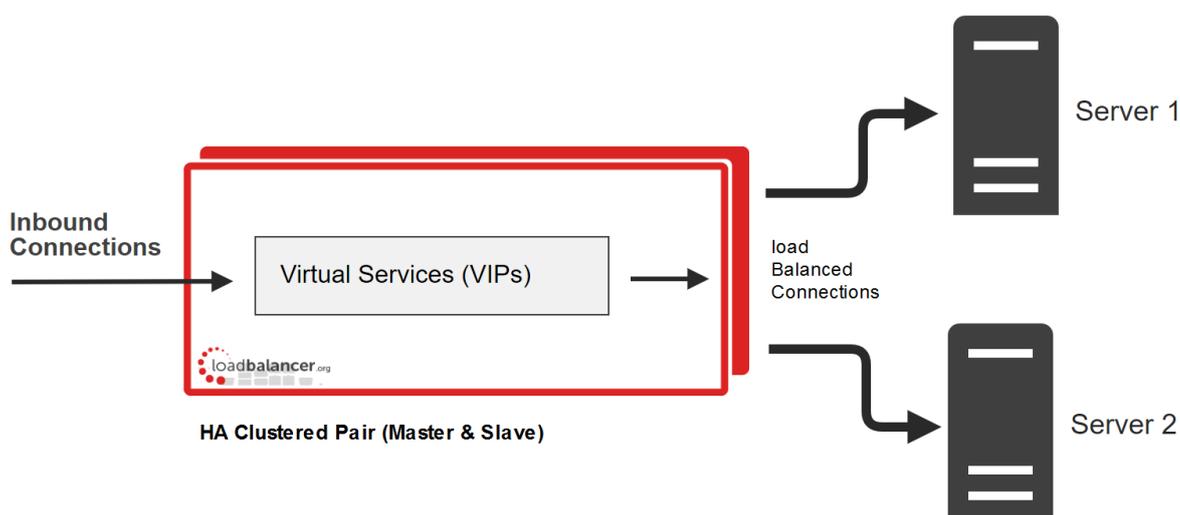
```
$ sudo su
```

2 - the 'loadbalancer' WebUI account

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

4. 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

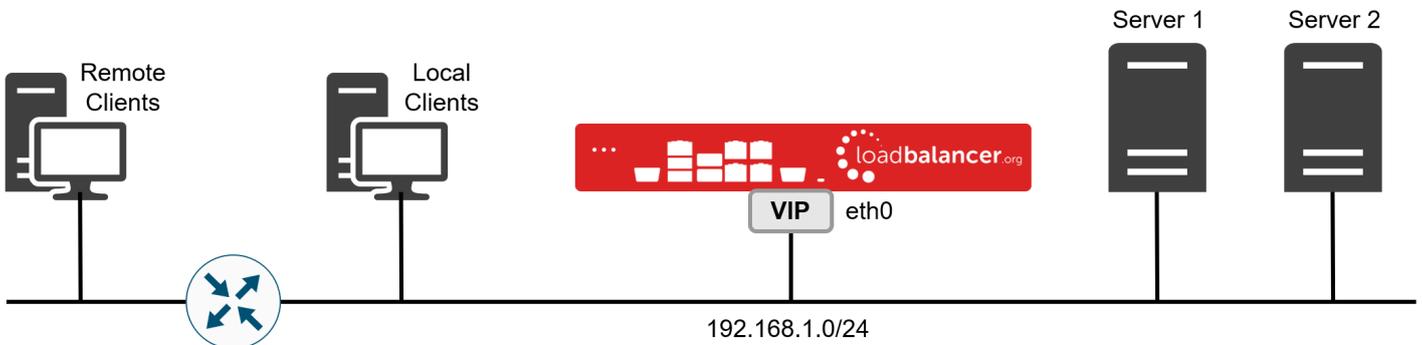
We always recommend that 2 appliances are deployed as an active/passive HA pair. The Secondary appliance automatically takes over if the Primary unit fails. For more information on configuring HA using 2 appliances please refer to [Configuring HA - Adding a Secondary Appliance](#).

5. 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 (SNAT) commonly use a one-arm configuration although two-arm is also supported.

One-Arm

The VIP and the load balanced servers are located in a single subnet. The load balancer requires a single network interface adapter - eth0 in the diagram below.

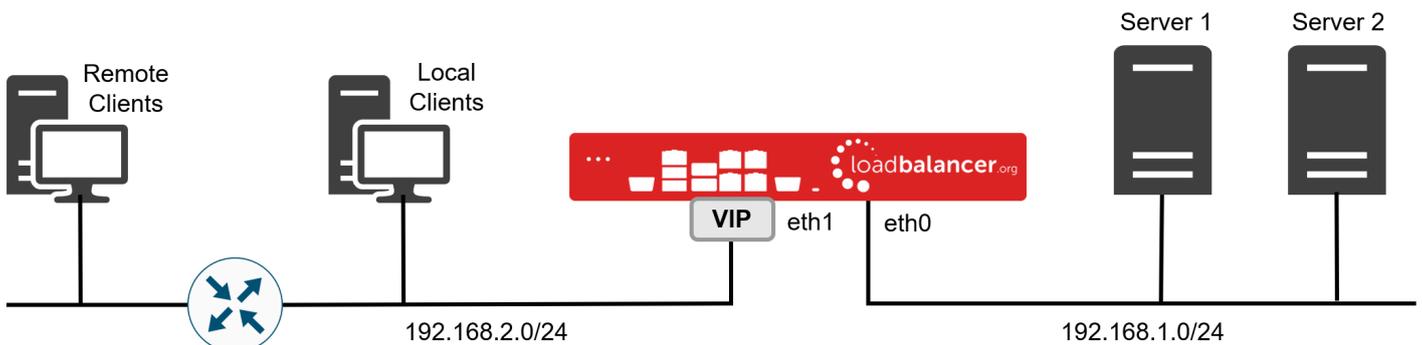


Two-Arm

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.

6. Supported Load Balancing Methods

The Loadbalancer.org appliance is one of the most flexible load balancers available. The design allows different load balancing modules to utilize the core high availability framework of the appliance. Multiple load balancing

methods can be used at the same time or in combination with each other. The table below describes the methods supported by the appliance.

Layer	Method	Comments	Topology	Note
Layer 4	DR (Direct Routing)	Ultra-fast local server based load balancing <ul style="list-style-type: none"> • Requires the ARP Problem to be solved on each Real Server - for more details please refer to DR Mode Considerations 	One-Arm (*)	1
Layer 4	NAT (Network Address Translation)	Fast Layer 4 load balancing <ul style="list-style-type: none"> • The appliance must be the default gateway for the Real Servers 	One or Two-Arm	1
Layer 4	TUN	Similar to DR but works across IP encapsulated tunnels	One-Arm	2
Layer 4	SNAT (Source Network Address Translation)	Fast layer 4 SNAT supporting both TCP & UDP <ul style="list-style-type: none"> • Very simple to implement • Requires no Real Server configuration changes 	One or Two-Arm	3
Layer 7	SSL Termination (STunnel & Pound)	Usually required in order to process cookie persistence in HTTPS streams on the load balancer <ul style="list-style-type: none"> • SSL Termination is processor intensive 	One or Two-Arm	4
Layer 7	SNAT (Source Network Address Translation using HAProxy)	Layer 7 allows greater flexibility including full SNAT and remote server load balancing, cookie insertion and URL switching <ul style="list-style-type: none"> • Very simple to implement • Requires no Real Server configuration changes • Not as fast as Layer 4 methods 	One or Two-Arm	4

(*) 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 the same 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 Direct Routing implementation across routed networks (rarely used).
3. Recommended 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. Recommended if HTTP cookie persistence is required, also used for several Microsoft applications such as Exchange, SharePoint & Remote Desktop Services and for overall deployment simplicity since Real Servers can be on any accessible subnet and no Real Server changes are required.

7. Ports Used by the Appliance

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

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

8. Appliance Deployment

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.

Hardware Appliance

For details of all hardware models and for information on installing and connecting the appliance, please click [here](#).

Cloud Appliance

AWS

For details of deploying and configuring the Amazon Web Services (AWS) appliance please refer to the [AWS Configuration Guide](#).

Azure

For details of deploying and configuring the Microsoft Azure appliance please refer to the [Azure Configuration Guide](#).

Google Cloud Platform

For details of deploying and configuring the Google Cloud appliance please refer to the [GCP Configuration Guide](#).

9. Configuring Initial Network Settings

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

```
Welcome to the Loadbalancer.org appliance.

To perform initial network configuration, log in to the console as
Username: setup
Password: setup

To access the web interface and wizard, point your browser at
http://192.168.2.21:9080/
or
https://192.168.2.21:9443/

lbmaster login: _
```

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

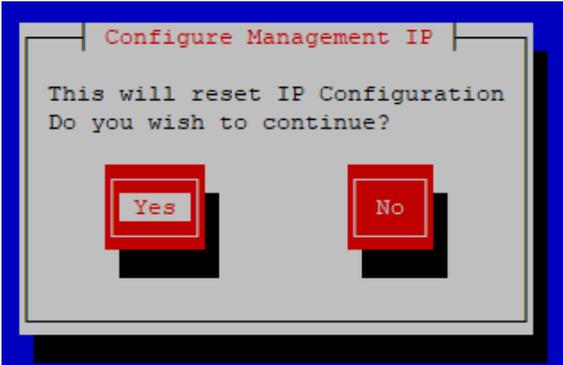
Once logged in, the Network Setup Wizard will start automatically. This will enable you to configure the management IP address and other network settings for the appliance.

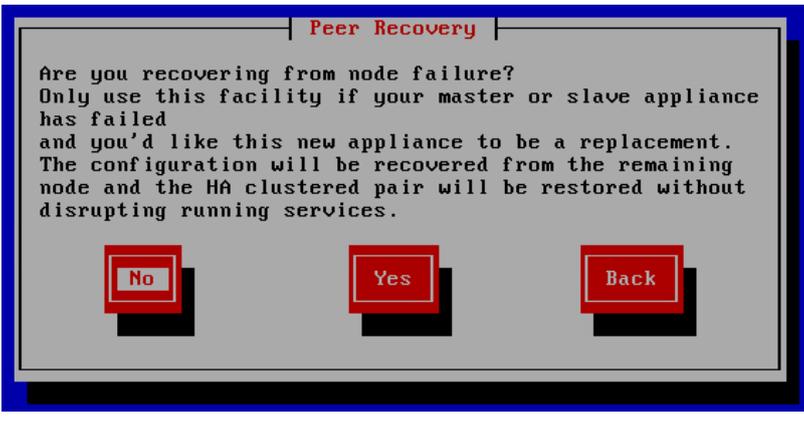
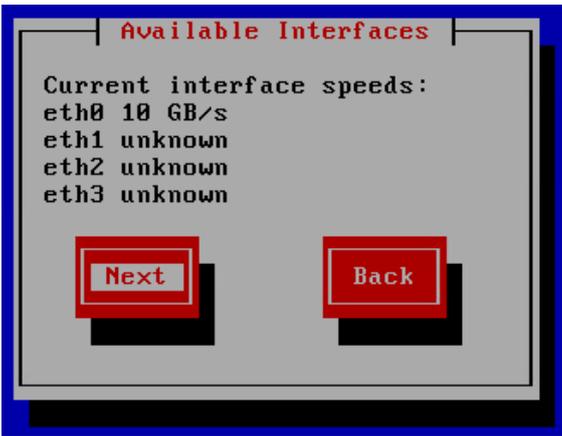
login to the console:

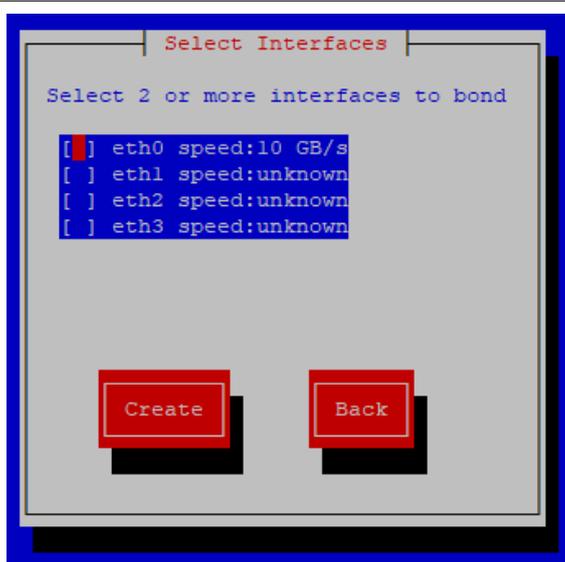
Username: setup

Password: setup

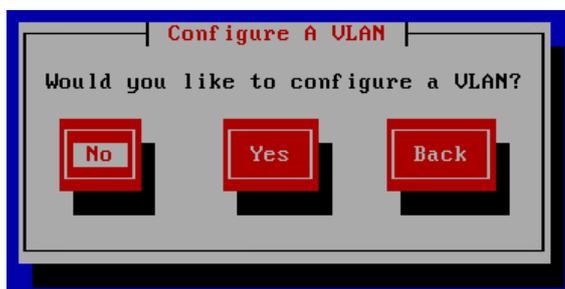
A series of screens will be displayed that allow network settings to be configured:

	<p>To continue with the Network Setup Wizard select Yes and hit <ENTER> to continue.</p>
--	---

 <p style="text-align: center;">Peer Recovery</p> <p>Are you recovering from node failure? Only use this facility if your master or slave appliance has failed and you'd like this new appliance to be a replacement. The configuration will be recovered from the remaining node and the HA clustered pair will be restored without disrupting running services.</p> <p style="text-align: center;"> <input type="button" value="No"/> <input type="button" value="Yes"/> <input type="button" value="Back"/> </p>	<p>Next, you'll be asked if you're recovering from node (i.e. Primary or Secondary) failure. If you're simply deploying a new appliance, select No and hit <ENTER> to continue.</p> <p>Note</p> <p>For more details on node recovery using this option please refer to Disaster Recovery After Node (Primary or Secondary) Failure.</p>
 <p style="text-align: center;">Centralised Management</p> <p>Do you wish to enroll this appliance with a centralised management device?</p> <p style="text-align: center;"> <input type="button" value="No"/> <input type="button" value="Yes"/> <input type="button" value="Back"/> </p>	<p>Next, you'll be asked if you want to enroll the appliance with a Central Management Server. if you have been provided with this information select Yes, otherwise leave No selected, then hit <ENTER> to continue.</p>
 <p style="text-align: center;">Available Interfaces</p> <p>Current interface speeds: eth0 10 GB/s eth1 unknown eth2 unknown eth3 unknown</p> <p style="text-align: center;"> <input type="button" value="Next"/> <input type="button" value="Back"/> </p>	<p>A list of available interfaces will be displayed, hit <ENTER> to continue.</p>
 <p style="text-align: center;">Configure Bonding</p> <p>Do you need to create a bonded interface?</p> <p style="text-align: center;"> <input type="button" value="No"/> <input type="button" value="Yes"/> <input type="button" value="Back"/> </p>	<p>Select Yes If you want to configure a bonded interface, if not leave No selected, then hit <ENTER> to continue.</p> <p>If you select Yes, the screen shown below will be displayed:</p>

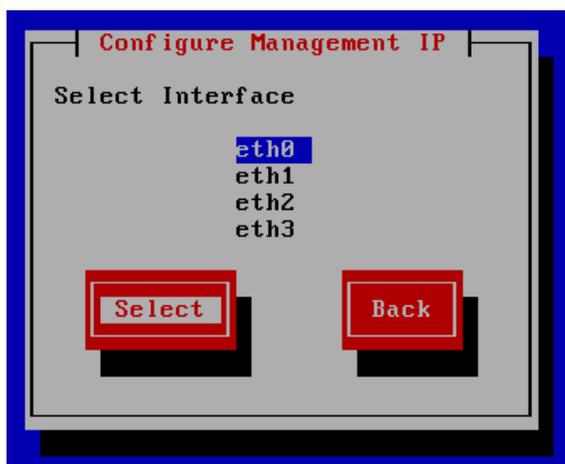


Using the space bar, select the interfaces you'd like to include in the bond, then click **Create**.



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.



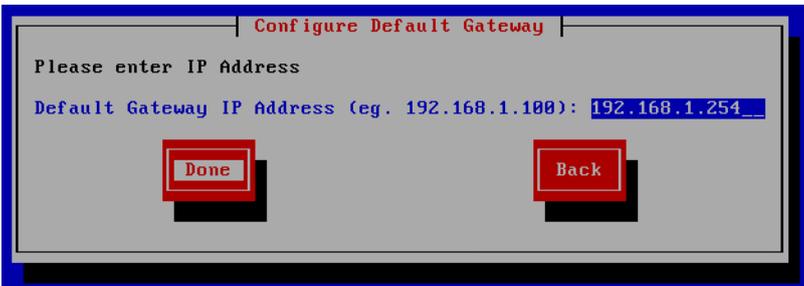
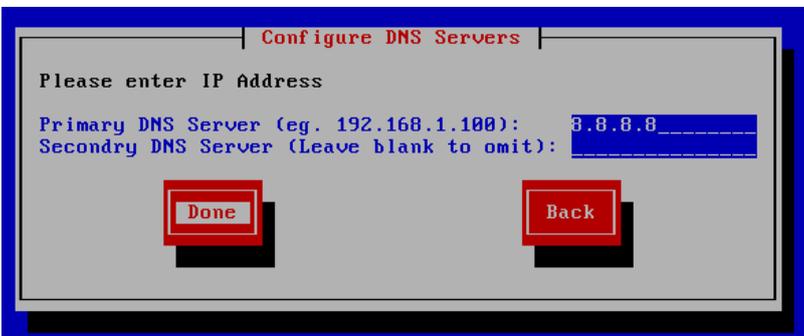
Select the interface that will be used to manage the appliance, select **Select** and hit <ENTER> to continue.

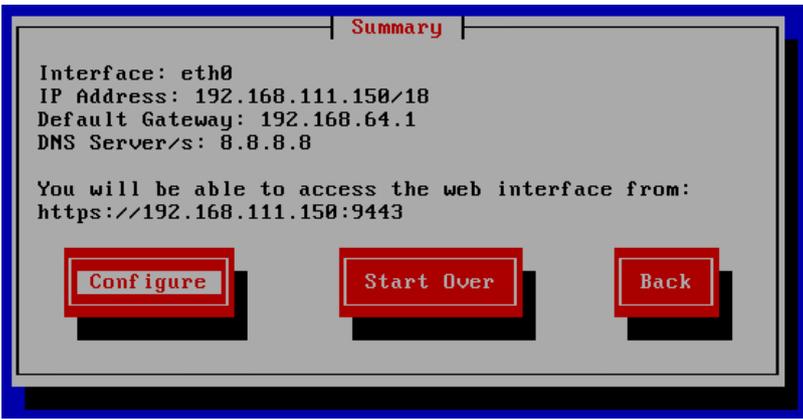


Either enter the required management IP address & CIDR prefix and select **Done** or select **Use DHCP** to request an address and then hit <ENTER> to continue.

Note

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

	<p>Enter the default gateway address, select Done and hit <ENTER> to continue.</p>
	<p>Configure the required DNS server(s), select Done and hit <ENTER> to continue.</p>
	<p>Hit <ENTER> to continue.</p>
	<p>Enter the password you'd like to use for the 'loadbalancer' WebUI user account and the 'root' Linux user account, select Done and hit <ENTER> to continue.</p>
	<p>If you selected Yes when asked in you want to configure a Central Management Server, you'll now be prompted for the details. Enter the required information and select Done and hit <ENTER> to continue.</p>

	<p>A summary of all settings is displayed, if everything looks good select Configure and hit <ENTER> to continue, all settings will then be applied.</p>
	<p>Once the configuration has been written, the Configuration Complete message will be displayed. Click OK to exit the wizard and return to the command prompt.</p>

10. Accessing the WebUI

The WebUI is accessed using a web browser. By default, user authentication is based on local Apache .htaccess files. User administration tasks such as adding users and changing passwords can be performed using the WebUI menu option: *Maintenance > Passwords*.

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

Note | If required, users can also be authenticated against LDAP, LDAPS, Active Directory or Radius. For more information please refer to [External Authentication](#).

1. Using a browser, access the WebUI using the following URL:

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

2. Log in to the WebUI:

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:

- System Overview
- Local Configuration
- Cluster Configuration
- Maintenance
- View Configuration
- Reports
- Logs
- Support
- Live Chat

WARNING: YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS.

Buy with confidence. All purchases come with a 90 day money back guarantee.
Already bought? Enter your license key [here](#)

Buy Now

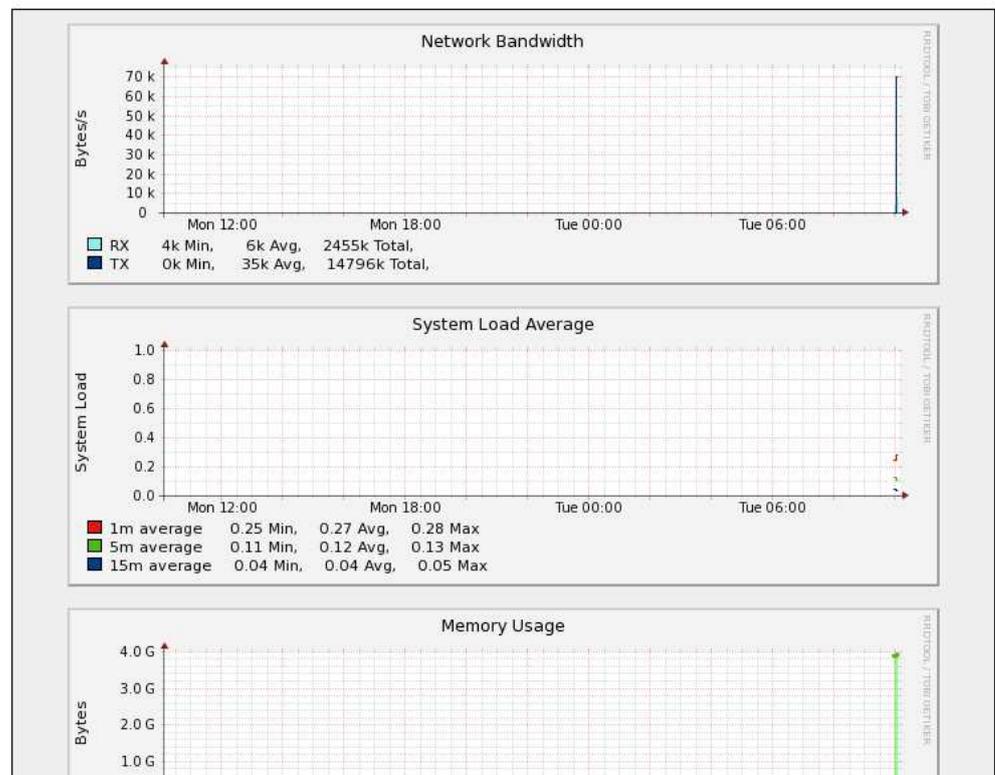
System Overview

2022-06-14 10:07:30 UTC

Would you like to run the Setup Wizard?

VIRTUAL SERVICE IP PORTS CONNS PROTOCOL METHOD MODE

No Virtual Services configured.



Note

The WebUI for the VA is shown, the hardware and cloud appliances are very similar. The yellow licensing related message is platform & model dependent.

- You'll be asked if you want to run the Setup Wizard. If you click **Accept** the Layer 7 Virtual Service configuration wizard will start. If you want to configure the appliance manually, simply click **Dismiss**.

Main Menu Options

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

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

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

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

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

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

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

11. Licensing

The trial runs for 30 days and is completely unrestricted during this time. After 30 days, the appliance continues to work but it's no longer possible to make changes to the configuration. If you need more time to complete your evaluation, please contact sales@loadbalancer.org who will be able to provide guidance on how to extend the trial using a simple command.

When a license is purchased, you'll be provided with a license key file by our sales team. You can then simply apply this license to your appliance.

To install the license:

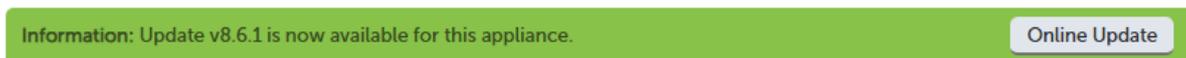
1. Using the WebUI, navigate to: *Local Configuration > License Key*.
2. Browse to the license file provided when the appliance was purchased.
3. Click **Install License Key**.

12. Software Updates

Loadbalancer.org continually develop and add new and improved features to the appliance. These updates can be applied during the trial period to ensure you have the very latest version of our software for your evaluation.

By default the appliance periodically contacts the Loadbalancer.org update server (update.loadbalancer.org) and checks for updates.

If an update is found, a message will be displayed at the top of the screen as shown in the following example:



To initiate the update process click the **Online Update** button.

Note

If your appliance does not have Internet access, please contact support@loadbalancer.org for details of how to obtain the offline update files (if available).

13. Configuring & Testing a Simple Load Balanced Test Environment

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

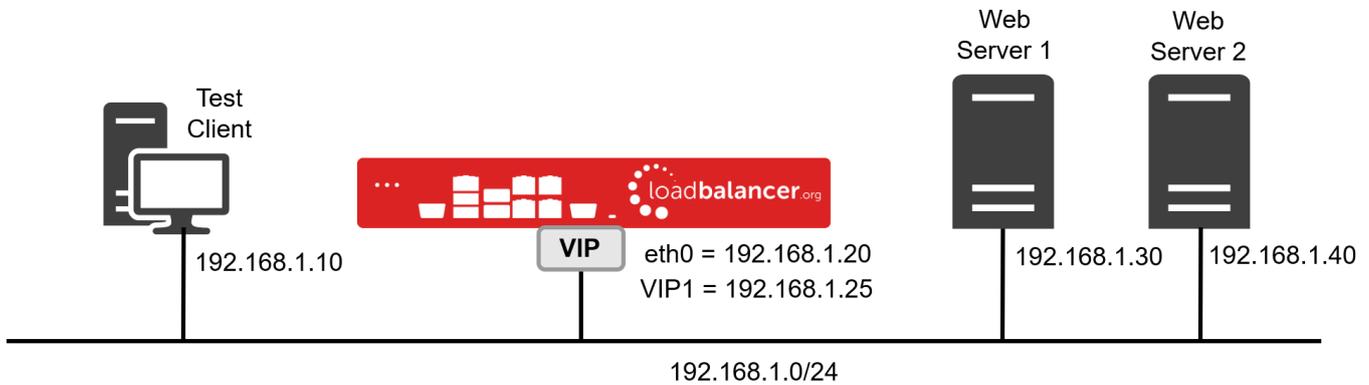
Note

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 configuration changes to the Real Servers.

The following table & diagram describe the environment:

IP Address	Device	Notes
192.168.1.10	Test Client	

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



STEP 1 - Deploy the Appliance

Please refer to [Appliance Deployment](#).

STEP 2 - Run the Network Setup Wizard

Please refer to [Configuring Initial Network Settings](#).

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

VIP Configuration

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

Virtual Service		[Advanced +]
Label	<input type="text" value="Web-Cluster"/>	?
IP Address	<input type="text" value="192.168.1.25"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?

- 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.

RIP 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:

Label	<input type="text" value="Web1"/>	
Real Server IP Address	<input type="text" value="192.168.1.30"/>	
Real Server Port	<input type="text" value="80"/>	
Re-Encrypt to Backend	<input type="checkbox"/>	
Enable Redirect	<input type="checkbox"/>	
Weight	<input type="text" value="100"/>	

- 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](#).

STEP 4 - Finalizing the Configuration

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the blue 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**.

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:

Layer 7 - Virtual Services

Search..				Add a new Virtual Service	
Service Name	IP	Port	Config Type		
Web-Cluster	192.168.1.25	Ports 80	Auto	Modify	Delete

- 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 blue 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:

Layer 7 - Real Servers

Web-Cluster		192.168.1.25		Ports 80		Add a new Real Server	
Web1	192.168.1.30	80	Weight 100	Modify	Delete		

- 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 blue box at the top of the screen to apply the changes.

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:

SYSTEM OVERVIEW ?							2022-04-07 13:31:41 UTC	
	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
↑	Test-VIP	192.168.1.25	80	0	TCP	Layer 7	Proxy	

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

↑	Test-VIP	192.168.1.25	80	0	TCP	Layer 7	Proxy	
REAL SERVER								
	IP	PORTS	WEIGHT	CONNS				
↑	Web1	192.168.1.30	80	100	0	Drain	Halt	
↑	Web2	192.168.1.40	80	100	0	Drain	Halt	

STEP 7 - Verification & Testing

1. Verify that both Real Servers are up. In the example below, Web2 is failing its health-check:

Test-VIP		192.168.1.25	80	0	TCP	Layer 7	Proxy	
REAL SERVER	IP	PORTS	WEIGHT	CONNS				
	Web1	192.168.1.30	80	100	0	Drain	Halt	
	Web2	192.168.1.40	80	100	0	Drain	Halt	

- This should be investigated and corrected, possible steps include:
 - Check that the application/service is running on the Real Server
 - Make sure you can ping the Real Server from the load balancer
 - Verify that you can connect to the application port from the load balancer. This can be done using telnet at the console or via an SSH session:

```
telnet 192.168.1.40 80
```

2. 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 (CTRL-F5)
- Halt Web2 using the *Halt* option for Web2 in the System Overview and verify that content is served by Web1 on a browser refresh (CTRL-F5)

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](#).

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

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

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

Non-Replicated Settings

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

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

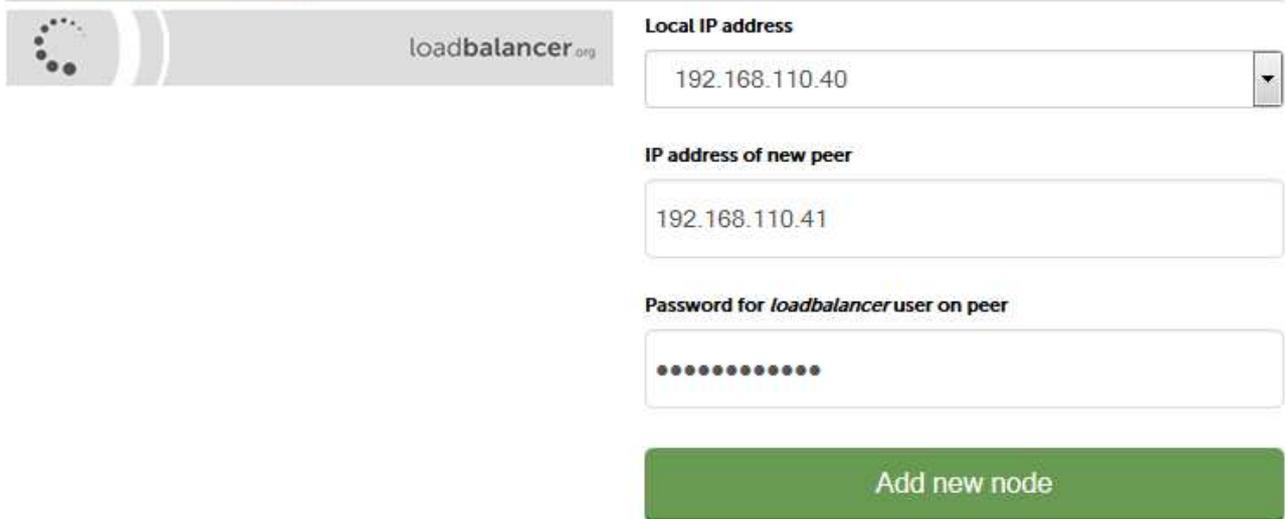
Important | Make sure that if these settings/updates have been configured on the Primary appliance, they're also configured on the Secondary appliance.

To add a Secondary node - i.e. create a highly available clustered pair:

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

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

CREATE A CLUSTERED PAIR



loadbalancer.org

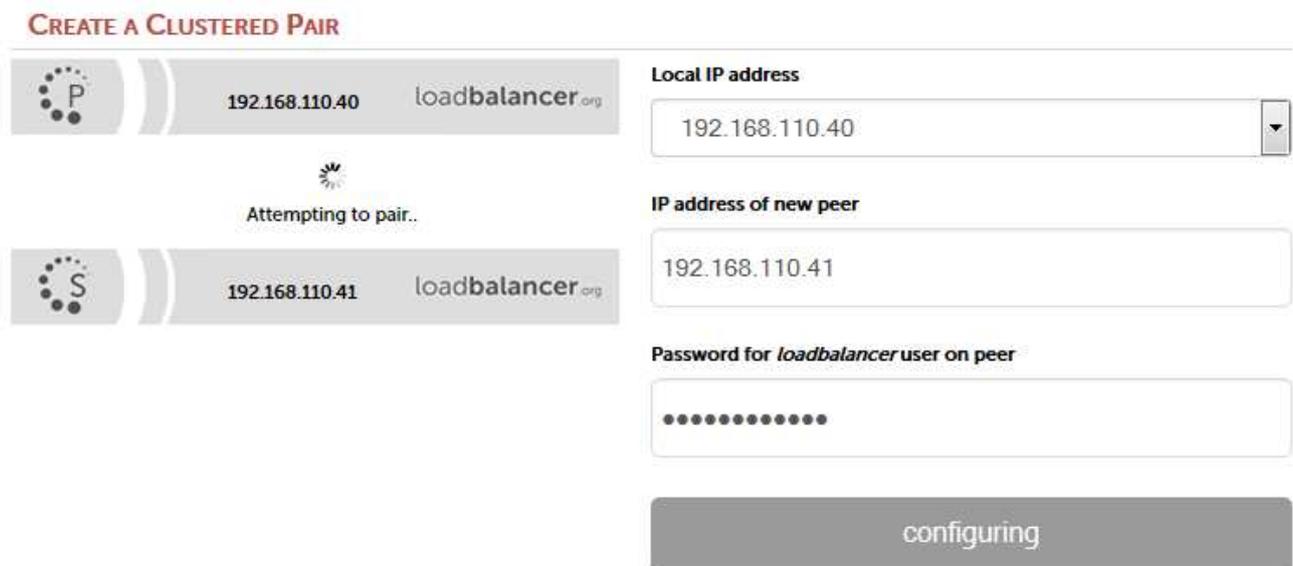
Local IP address
192.168.110.40

IP address of new peer
192.168.110.41

Password for *loadbalancer* user on peer
.....

Add new node

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



CREATE A CLUSTERED PAIR

192.168.110.40 loadbalancer.org

Attempting to pair..

192.168.110.41 loadbalancer.org

Local IP address
192.168.110.40

IP address of new peer
192.168.110.41

Password for *loadbalancer* user on peer
.....

configuring

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



High Availability Configuration - primary

192.168.110.40 loadbalancer.org

192.168.110.41 loadbalancer.org

Break Clustered Pair

7. To finalize the configuration, restart heartbeat and any other services as prompted in the blue 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](#).

15. More Information

Please refer to our website for all the latest [Manuals](#) and [Deployment Guides](#).

16. Loadbalancer.org Technical 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.

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.



United Kingdom

Loadbalancer.org Ltd.
Compass House, North Harbour
Business Park, Portsmouth, PO6 4PS
UK: +44 (0) 330 380 1064
sales@loadbalancer.org
support@loadbalancer.org

Canada

Loadbalancer.org Appliances Ltd.
300-422 Richards Street, Vancouver,
BC, V6B 2Z4, Canada
TEL: +1 866 998 0508
sales@loadbalancer.org
support@loadbalancer.org

United States

Loadbalancer.org, Inc.
4550 Linden Hill Road, Suite 201
Wilmington, DE 19808, USA
TEL: +1 833.274.2566
sales@loadbalancer.org
support@loadbalancer.org

Germany

Loadbalancer.org GmbH
Tengstraße 2780798,
München, Germany
TEL: +49 (0)89 2000 2179
sales@loadbalancer.org
support@loadbalancer.org