



Appliance Quick Start Guide **v8.3.6**

Rev. 1.0.2

Table of Contents

1. About this Guide.....	3
2. Appliance Configuration Overview.....	3
3. Appliance Security.....	3
4. Deployment Concept.....	4
5. One-Arm and Two-Arm Topologies.....	4
6. Load Balancing Methods.....	5
7. Appliance Deployment.....	6
Virtual Appliance.....	6
Hardware Appliance.....	7
Cloud Appliance.....	7
AWS.....	7
Azure.....	7
8. Initial Network Interface Configuration.....	7
9. Accessing the Web User Interface (WebUI).....	9
10. Ports Used by the Appliance.....	10
11. Licensing.....	10
12. Software Updates.....	11
13. Configuring & Testing a Simple Load Balanced Test Environment.....	11
STEP 1 – Deploy the Appliance.....	12
STEP 2 – Run the Network Setup Wizard.....	12
STEP 3 – Run the WebUI Setup Wizard.....	12
STEP 4 – Viewing & Modifying the Configuration.....	13
STEP 5 – Checking the Status using System Overview.....	14
STEP 6 – Verification & Testing.....	14
14. Configuring HA - Adding a Slave Appliance.....	15
15. More Information.....	16
16. Loadbalancer.org Technical Support.....	16
17. Company Contact Information.....	17

1. About this Guide

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

Note:

Please refer to the [Administration Manual](#) for much more detailed information on setting up the appliance and configuring a load balancing solution.

2. Appliance Configuration Overview

Initial network configuration can be carried out in a number of ways. The recommended method is to use the Network Configuration Wizard at the console of the VA. Please refer to page [7](#) for more details.

Once the network is configured, load balanced services can be configured using the WebUI, either with the Setup Wizard (for Layer 7 services) or manually. The WebUI is accessible using **HTTP** on port **9080** and **HTTPS** on port **9443**. Please refer to page [9](#) for more details.

We always recommend that where possible two load balancer 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 master is fully configured first, then the slave should be added. For more information on configuring an HA pair please refer to page [15](#).

Note:

Please refer to the [Administration Manual \(page 77\)](#) for details on limiting WebUI access to HTTPS only.

3. Appliance Security

We strongly recommend that default passwords are changed as soon as the appliance is deployed. Passwords for the following user accounts should be changed:

1 - the 'root' Linux account:

The password can be changed at the console, or via an SSH session using the following command:

```
passwd
```

2 - the 'loadbalancer' WebUI account:

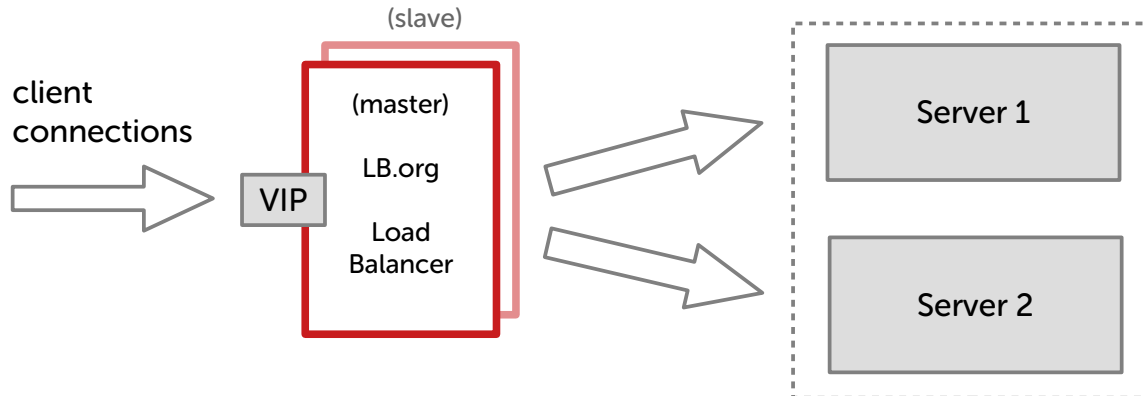
This can be changed using the WebUI menu option: *Maintenance > Passwords*

Note:

The appliance also includes a security lockdown command (lbsecure) that enables passwords to be set, network access to be locked down and SSH key regeneration in one simple step. This command can be run on a single appliance or an HA pair. For more details please refer to the [Administration Manual \(page 75\)](#).

4. Deployment Concept

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



VIP = Virtual IP Address

load balanced servers
(aka : "Real Servers")
(aka : "RIPs")

Note:


We recommend that 2 appliances are deployed as an active/passive HA pair as shown above. The slave appliance automatically takes over if the master unit fails. Please refer to page [15](#) for more information on setting up an HA pair of Loadbalancer.org appliances.

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. Proxy based load balancers (SNAT) commonly use a one-arm configuration.

One Arm

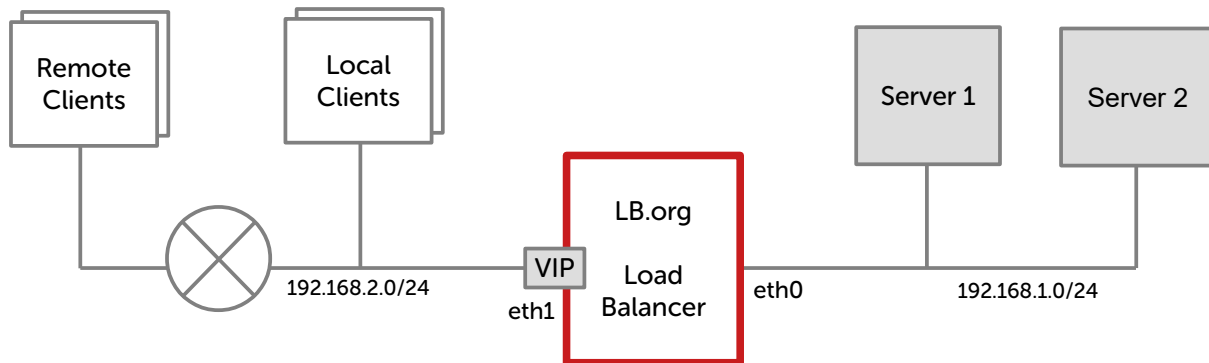
In this mode, the VIP and the load balanced servers are located in a single subnet. The load balancer requires a single network interface adapter – eth0 in the diagram below.



Note:
The example configuration on page [11](#) of this guide uses a one-arm topology.

Two Arm

In this mode, 2 subnets are used. The VIP is located in one subnet and the load balanced servers are located in the other subnet. The load balancer requires 2 interfaces – eth0 and eth1 in the diagram below. Note that this can be achieved by using two network adapters, or by creating VLAN's on a single adapter.



6. Load Balancing Methods

The Loadbalancer.org appliance is one of the most flexible load balancers on the market. 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.

Layer 4	DR (Direct Routing)	Ultra-fast local server based load balancing - Requires solving the 'ARP problem' on the Real Servers - please refer to the Administration Manual (page 87) for more details	One-Arm (*)
Layer 4	NAT (Network Address Translation)	Fast Layer 4 load balancing - The appliance must be the default gateway for the Real Servers	One or Two-Arm
Layer 4	TUN	Similar to DR but works across IP encapsulated tunnels	One-Arm
Layer 4	SNAT (Source Network Address Translation)	Fast layer 4 SNAT supporting both TCP & UDP - Requires no Real Server changes	One or Two-Arm
Layer 7	SSL Termination (Pound & STunnel)	Usually required in order to process cookie persistence in HTTPS streams on the load balancer - SSL Termination is processor intensive	One or Two-Arm

Layer 7	SNAT (Source Network Address Translation using HAProxy)	Layer 7 allows great 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 changes - Not as fast as Layer 4 methods 	One or Two-Arm
---------	--	--	----------------

(*) 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.

Key

- Recommended for high performance fully transparent and scalable solutions
- 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
- Only required for Direct Routing implementation across routed networks (rarely used)
- 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

Note:

Layer 7 SNAT mode is generally the simplest most flexible method to use. As mentioned above, it's not the fastest method but enables a robust load balancing solution to be rapidly deployed. This is illustrated by the simple example on page [11](#) of this guide.

Note:

Please refer to the [Administration Manual \(pages 26-32\)](#) for more detailed information on each load balancing method.

7. Appliance Deployment

VIRTUAL APPLIANCE

The VA is currently available for VMware, Virtual Box, Hyper-V, KVM, Nutanix and XEN and has been optimized for each Hypervisor. By default, the VA is allocated 1 CPU, 2GB of RAM and has an 8GB virtual disk.

Note:

The Virtual Appliance can be downloaded [here](#).

Note:

Please refer to the [Administration Manual \(page 35\)](#) and the ReadMe.txt text file included in each VA download for more detailed information on deploying the VA using various Hypervisors.

HARDWARE APPLIANCE

For details of all hardware models and information on installing and connecting the appliance, please refer to the Hardware Installation Guide available on our [Website](#).

CLOUD APPLIANCE

AWS

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

AZURE

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

8. Initial Network Interface Configuration

By default the load balancer is pre-configured with the following IP address & subnet mask:

192.168.2.21/24 (192.168.2.21/255.255.255.0)

This can be easily changed using the Network Setup Wizard at the console. To run the wizard, login to the console of the appliance as the 'setup' user. This is explained in the initial console start-up message as shown below:

```
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: _
```

login to the console:

```
Username: setup
Password: setup
```

Once logged in, enter the IP address, mask, VLAN tag ID, default gateway & DNS servers at the prompts as shown in the example below:

Loadbalancer.org basic network set up

This will overwrite the current configuration.
If you do not wish to proceed please enter CTRL + c.

```
Static IP address (eg. 192.168.0.26)      :   192.168.1.20
Interface netmask (eg. 24)                :   24
VLAN tag ID (Press enter to skip) (eg. 10) :   120
Default gateway (eg. 192.168.0.1)        :   192.168.1.254

DNS Servers
  Primary (eg. 192.168.0.250)             :   8.8.8.8
  Secondary (Leave blank to omit)          :   -
```

After the required settings have been entered, a summary will be presented along with details of how to access the WebUI as shown below:

Summary of settings

```
Static IP address:      192.168.1.20/24
Default gateway:       192.168.1.254
VLAN ID:               120
DNS servers:           8.8.8.8
```

You may now connect the eth0 network interface to your switch, and continue configuration through the web interface on:

```
http://192.168.1.20:9080/lbadmin/
```

As mentioned in the text above, the IP address is now configured for interface eth0.

At this stage you will also be asked if you're recovering from node (i.e. master or slave) failure as shown below:

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

(If you are simply deploying a new appliance, hit N)

Do you want to continue? [y/N]

-

As mentioned in the text, if you're simply deploying a new appliance, click "N"

Note:

If you set a VLAN tag ID, and later want to remove this, you'll need to first restore default settings using the WebUI option: *Maintenance > Backup & Restore* and clicking **Restore Manufacturer's Defaults**, then run through the Network Setup Wizard again.

9. Accessing the Web User Interface (WebUI)

- Using a web browser, access the WebUI using the following URL:

http://192.168.2.21:9080/lbadmin/

(replace with your IP address if it's been changed)

or via HTTPS:

https://192.168.2.21:9443/lbadmin/

(replace with your IP address if it's been changed)

- Login to the WebUI:

Username: loadbalancer

Password: loadbalancer

Note:

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

The screenshot shows the loadbalancer.org web interface. At the top left is the logo and 'loadbalancer.org'. At the top right, it says 'Enterprise VA MAX'. Below the logo, there are status indicators: 'Master | Slave', 'Active | Passive', 'Link', and '12 Seconds'. A navigation menu on the left includes: System Overview, Local Configuration, Cluster Configuration, Maintenance, View Configuration, Reports, Logs, and Support. The main content area features a yellow warning banner: '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'. Below the banner is the 'System Overview' section with a timestamp '2019-01-30 10:36:02 UTC'. A dark grey dialog box asks 'Would you like to run the Setup Wizard?' with 'Accept' and 'Dismiss' buttons. Below the dialog is a table for 'VIRTUAL SERVICE' with columns: VIRTUAL SERVICE, IP, PORTS, CONNS, PROTOCOL, METHOD, and MODE. The table is empty, with the text 'No Virtual Services configured.' below it. At the bottom, there is a 'Network Bandwidth' graph with a y-axis labeled 'Bytes/s' ranging from 0 to 200k and an x-axis for time.

The WebUI for the VA is shown. The hardware and cloud appliances are very similar, but have

different startup messages (shown yellow above).

- Once logged in, you'll be asked if you want to run the web based setup wizard. If you click **Accept** the Layer 7 Virtual Service configuration wizard will start (please refer to page [12](#) for details of using the wizard). 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

10. Ports Used by the Appliance

The appliance uses the following ports by default:

Protocol	Port	Purpose
TCP	22	SSH
TCP	9080	WebUI - HTTP
TCP	9081	Nginx fallback page
TCP	9443	WebUI - HTTPS
TCP	7777	HAProxy statistics page
TCP	7778	HAProxy persistence table replication
UDP	6694	Heartbeat between master & slave appliances in HA mode

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:

- Using the WebUI, navigate to: *Local Configuration > License Key*
- Browse to the license file provided when the appliance was purchased
- 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 to ensure you have the very latest version of our software for your evaluation.

To run Software Update:

1. Using the WebUI, navigate to: *Maintenance > Software Update*
2. Choose **Online Update** if the appliance has Internet access
3. If updates are available, you'll be presented with a list of changes, click the **Online Update** button at the bottom of the page to start the update

Note:

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

13. Configuring & Testing a Simple Load Balanced Test Environment

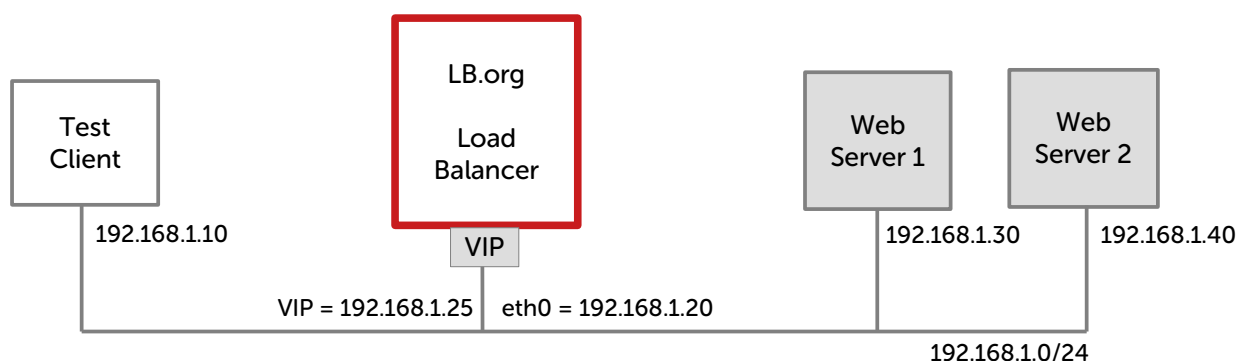
This example illustrates how to quickly configure a simple load balanced test environment using the Network Setup Wizard at the console to configure network settings, and the Setup Wizard from the WebUI to configure the layer 7 virtual service.

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 changes to the Real Servers.

The following table and diagram describe the environment:

IP Address	Device	Notes
192.168.1.10	Test Client	
192.168.1.20	Load Balancer	the load balancers own IP address
192.168.1.25	Load Balancer	the Virtual IP address (VIP), the IP address the clients connect to
192.168.1.30	Web Server 1	
192.168.1.40	Web Server 2	



STEP 1 – DEPLOY THE APPLIANCE

- Please refer to page [6](#)

STEP 2 – RUN THE NETWORK SETUP WIZARD

- Please refer to page [7](#)

STEP 3 – RUN THE WEBUI SETUP WIZARD

1. Open the WebUI and start the wizard by clicking the **Accept** button, or by using the WebUI menu option: *Cluster Configuration > Setup Wizard* and clicking **General Layer 7 Virtual Service**
2. Define the required Virtual Service settings as shown in the example below:

Setup Wizard - General Layer 7 Virtual Service

Load balancer configuration

	Master	Slave
Hostname	lbmaster	<i>Not configured</i>
Static IP Addresses eth0	192.168.1.20/24	
Floating IP Addresses		

Create a new Layer 7 Virtual Service

Label

Virtual Service	IP Address	<input style="width: 100%;" type="text" value="192.168.1.25"/>
	Ports	<input style="width: 100%;" type="text" value="80"/>

Layer 7 Protocol

Select the Layer 7 protocol to be handled by this Virtual Service.
Advanced options may be set by editing this Virtual Service once it has been created.

3. Click **Create Virtual Service**
4. Now continue and add the associated load balanced servers (Real Servers) as shown below:

Attach Real Servers

Label	IP Address	Port	Weight	
<input type="text" value="Web1"/>	<input type="text" value="192.168.1.30"/>	<input type="text" value="80"/>	<input type="text" value="100"/>	
<input type="text" value="Web2"/>	<input type="text" value="192.168.1.40"/>	<input type="text" value="80"/>	<input type="text" value="100"/>	<input type="button" value="✘"/>

- Use the **Add Real Server** button to define additional Real Servers and use the red cross to delete Real Servers
- Once you're happy, click **Attach Real Servers** to create the new Virtual Service & Real Servers
- A confirmation message will be displayed as shown in the example below:

Information: Real Server Web1 added.

Information: Real Server Web2 added.

Information: Virtual Service configured successfully

5. Click **Continue**
6. Finally, reload HAProxy using the **Reload HAProxy** button in the blue box at the top of the screen or by using the WebUI menu option: *Maintenance > Restart Services* and clicking **Reload HAProxy**

Note:

Running the wizard again will permit additional Layer 7 VIPs and associated RIPs to be defined.

Note:

To restore manufacturer's settings use the WebUI menu option: *Maintenance > Backup & Restore > Restore Manufacturer's Defaults*. This will reset the IP address to 192.168.2.21/24.

Note:

By default Real Server health-checks set as a TCP port connect. If you need a more robust check, this can be changed by modifying the configuration as explained below. Please refer to the [Administration Manual \(page 188\)](#) for more information on configuring health-checks.

STEP 4 – VIEWING & MODIFYING THE CONFIGURATION

1. The VIP created by the wizard can be seen 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	
Test-VIP	192.168.1.25	Ports 80	Auto	Modify Delete

2. Clicking the **Modify** button allows all VIP setting to be modified
3. If changes are made, click the **Update** button to save the changes, then use the **Reload HAProxy** button at the top of the screen to apply the changes
4. Additional VIPs can be added by running the Setup Wizard again, or by clicking the **Add a new Virtual Service** button to define the VIP manually

Note:
 Real Servers can be added manually using the WebUI menu option: *Cluster Configuration > Layer 7 – Real Servers*.

STEP 5 – CHECKING THE STATUS USING SYSTEM OVERVIEW

1. Using the WebUI, navigate to: *System Overview* to view the newly created VIP & RIPs:

SYSTEM OVERVIEW 2016-10-14 12:26:41 UTC

VIRTUAL SERVICE	IP	PORTS	CONN	PROTOCOL	METHOD	MODE
↑ Test-VIP	192.168.1.25	80	0	TCP	Layer 7	Proxy

2. To view the RIPs, click anywhere on the horizontal gray area to expand the VIP as shown below:

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

STEP 6 – 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	CONN		
↑ 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:

Please refer to the [Administration Manual \(page 232-247\)](#) for more configuration examples using Layer 7 SNAT mode and also Layer 4 DR mode, Layer 4 NAT mode & Layer 4 SNAT mode.

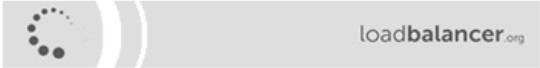
14. Configuring HA - Adding a Slave Appliance

As mentioned earlier, 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 master is fully configured first, then the slave should be added. The clustered HA pair uses Heartbeat to determine the state of the other appliance. Should the active device (normally the master) suffer a failure, the passive device (normally the slave) will take over.

To add a slave node – i.e. create a highly available clustered pair:

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

CREATE A CLUSTERED PAIR



Local IP address

IP address of new peer

Password for *loadbalancer* user on peer

[Add new node](#)

3. Specify the IP address and the *loadbalancer* users password (the default is 'loadbalancer') for the slave (peer) appliance as shown above

- Click **Add new node**
- The pairing process now commences as shown below:

- Once complete, the following will be displayed:

- 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 master appliance will also automatically restart heartbeat on the slave appliance.

Note:

Please refer to the [Administration Manual \(page 208-220\)](#) for more detailed information on configuring HA with 2 appliances.

15. More Information

Please refer to our website for the latest administration manual, deployment guides and all other documentation: <http://www.loadbalancer.org/resources/manuals>

16. Loadbalancer.org Technical Support

If you have any questions regarding the appliance or how to load balance your application, please don't hesitate to contact our support team using the following email address: support@loadbalancer.org

17. Company Contact Information

Website	URL: www.loadbalancer.org
North America (US)	<p>Loadbalancer.org, Inc. 4550 Linden Hill Road, Suite 201 Wilmington, DE 19808 USA</p> <p>Tel: +1 833.274.2566 Email (sales): sales@loadbalancer.org Email (support): support@loadbalancer.org</p>
North America (Canada)	<p>Loadbalancer.org Appliances Ltd. 300-422 Richards Street Vancouver, BC V6B 2Z4 Canada</p> <p>Tel: +1 302.213.0122 Email (sales): sales@loadbalancer.org Email (support): support@loadbalancer.org</p>
Europe (UK)	<p>Loadbalancer.org Ltd. Compass House North Harbour Business Park Portsmouth, PO6 4PS UK</p> <p>Tel: +44 (0)330 380 1064 Email (sales): sales@loadbalancer.org Email (support): support@loadbalancer.org</p>
Europe (Germany)	<p>Loadbalancer.org GmbH Tengstraße 27 80798 München Germany</p> <p>Tel: +49 (0)89 2000 2179 Email (sales): vertrieb@loadbalancer.org Email (support): support@loadbalancer.org</p>