

Load Balancing Verint WorkForce Optimization

Version 1.0.0



Table of Contents

1. About this Brief	3
2. Loadbalancer.org Appliances Supported	3
3. Software Versions Supported	3
3.1. Loadbalancer.org Appliance	3
3.2. Verint WFO	3
4. Verint WFO	3
5. Load Balancing Verint WFO	3
5.1. Virtual Service (VIP) Requirements	3
5.2. SSL Termination	4
6. Deployment Concept	4
7. Load Balancer Deployment Methods	4
7.1. Layer 4 SNAT Mode	4
8. Loadbalancer.org Appliance – the Basics	5
8.1. Virtual Appliance	5
8.2. Initial Network Configuration	6
8.3. Accessing the Appliance WebUI	6
Main Menu Options	7
8.4. Appliance Software Update	8
Determining the Current Software Version	8
Checking for Updates using Online Update	8
Using Offline Update	9
8.5. Ports Used by the Appliance	9
8.6. HA Clustered Pair Configuration	10
9. Appliance Configuration for Verint WFO	10
9.1. VIP1 - VERINT-WFO	10
Configuring the Virtual Service (VIP)	10
Defining the Real Servers (RIPs)	11
10. Testing & Verification	12
10.1. Using System Overview	12
10.2. Login to the Application	12
11. Technical Support	12
12. Further Documentation	12
13. Appendix	13
13.1. Configuring HA - Adding a Secondary Appliance	13
Non-Replicated Settings	13
Configuring the HA Clustered Pair	14
14. Document Revision History	16

1. About this Brief

This brief outlines the steps required to configure a load balanced Verint WorkForce Optimization (WFO) environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Verint WFO configuration changes that are required to enable load balancing.

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

2. Loadbalancer.org Appliances Supported

All our products can be used with Verint WFO. For full specifications of available models please refer to <https://www.loadbalancer.org/products>.

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

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.9.1 and later

Note

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

3.2. Verint WFO

- V15.2 and later

4. Verint WFO

Verint WFO is a unified suite of cloud solutions for capturing interactions and managing the performance of employees across the enterprise. This includes workforce management, call recording, automated quality management, performance management, speech, text, and desktop analytics and much more.

5. Load Balancing Verint WFO

Note

It's highly recommended that you have a working Verint WFO environment first before implementing the load balancer.

5.1. Virtual Service (VIP) Requirements

To provide load balancing and HA for Verint WFO, the following VIP is required:

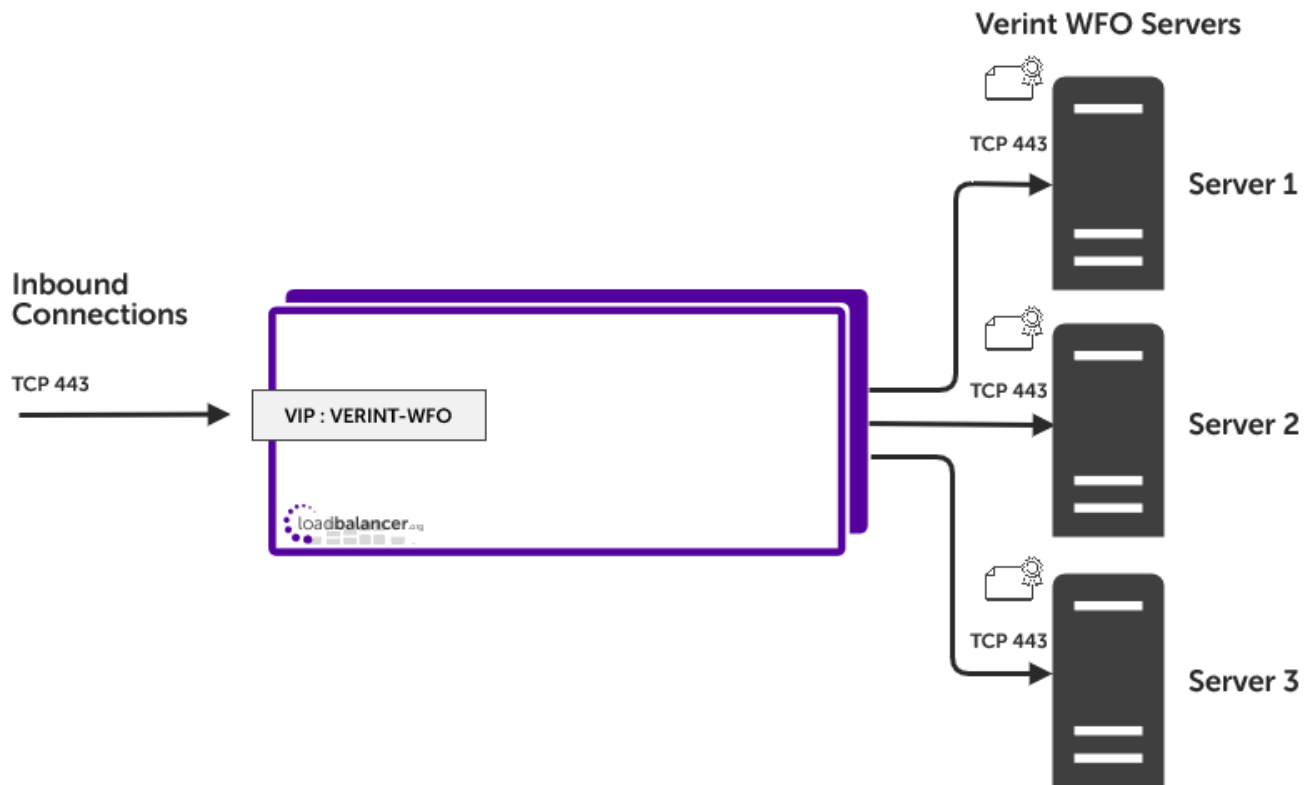


Reference	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP1	VERINT-WFO	L4 SNAT	443	Source IP	HTTPS (GET)

5.2. SSL Termination

SSL is terminated on the Verint WFO servers.

6. Deployment Concept



VIP = **V**irtual **I**P Address

Note

The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to the section [Configuring HA - Adding a Secondary Appliance](#) in the appendix for more details on configuring a clustered pair.

7. Load Balancer Deployment Methods

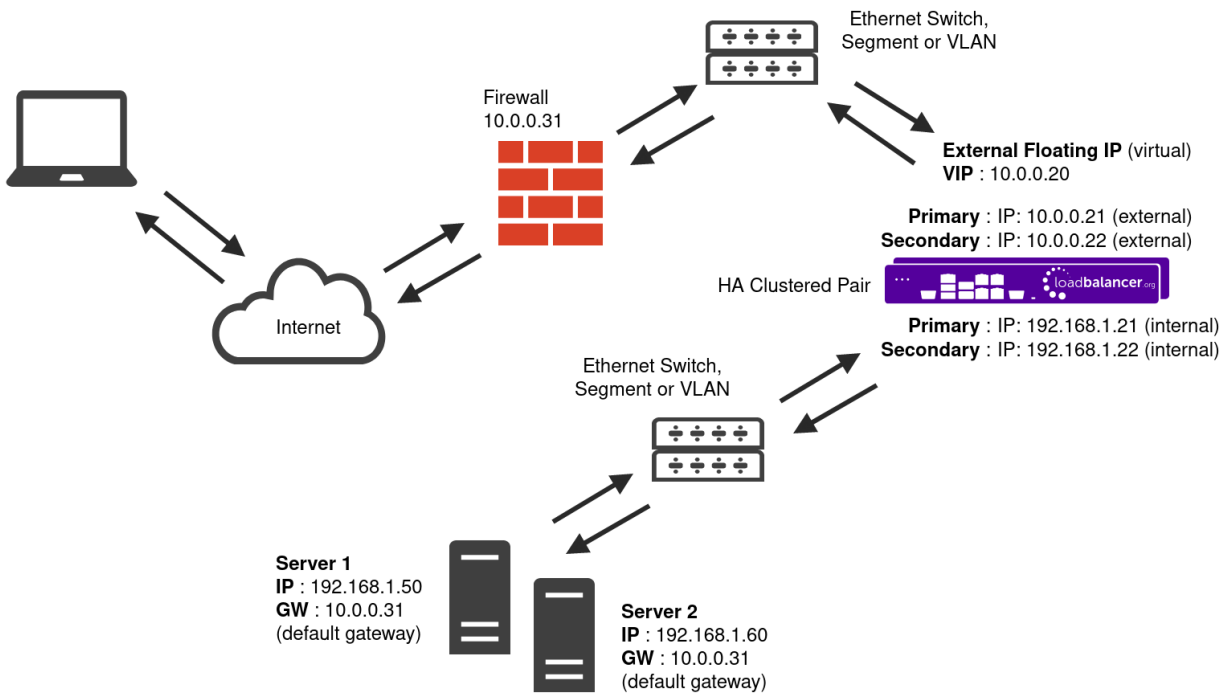
The load balancer can be deployed in 4 fundamental ways: *Layer 4 DR mode*, *Layer 4 NAT mode*, *Layer 4 SNAT mode*, and *Layer 7 SNAT mode*.

For Verint WFO, layer 4 SNAT mode is recommended. This mode is described below and is used for the configurations presented in this guide.

7.1. Layer 4 SNAT Mode



Layer 4 SNAT mode is a high performance solution, although not as fast as Layer 4 NAT mode or Layer 4 DR mode.



- The load balancer translates all requests from the external Virtual Service to the internal Real Servers in the same way as NAT mode - please refer to [Layer 4 NAT Mode](#) for more information.
- Layer 4 SNAT mode is not transparent, an iptables SNAT rule translates the source IP address to be the load balancer rather than the original client IP address.
- Layer 4 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, **eth0** is normally used for the internal network and **eth1** is used for the external network although this is not mandatory.
- If the Real Servers require Internet access, **Auto-NAT** should be enabled using the WebUI option: **Cluster Configuration > Layer 4 - Advanced Configuration**, the external interface should be selected.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 4 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 4 SNAT mode VIPs and layer 7 SNAT mode VIPs because the required firewall rules conflict.

8. Loadbalancer.org Appliance – the Basics

8.1. 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.

8.2. Initial Network Configuration

After boot up, follow the instructions on the appliance console to configure the management IP address, subnet mask, default gateway, DNS servers and other network and administrative settings.

 **Important**

Be sure to set a secure password for the load balancer, when prompted during the setup routine.

8.3. Accessing the Appliance WebUI

The WebUI is accessed using a web browser. By default, users are authenticated using Apache authentication. Users can also be authenticated against LDAP, LDAPS, Active Directory or Radius - for more information, please refer to [External Authentication](#).

 **Note**

There are certain differences when accessing the WebUI for the cloud appliances. For details, please refer to the relevant [Quick Start / Configuration Guide](#).

1. Using a browser, navigate to the following URL:

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

 **Note**

You'll receive a warning about the WebUI's SSL certificate. This is due to the default self signed certificate that is used. If preferred, you can upload your own certificate - for more information, please refer to [Appliance Security Features](#).

 **Note**

If you need to change the port, IP address or protocol that the WebUI listens on, please refer to [Service Socket Addresses](#).

2. Log in to the WebUI using the following credentials:

Username: loadbalancer

Password: <configured-during-network-setup-wizard>





Note

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

Once logged in, the WebUI will be displayed as shown below:

LOADBALANCER

Enterprise VA Max

Primary | Secondary Active | Passive Link 15 Seconds ↻

- 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](#)

System Overview ? 2024-03-15 16:27:21 UTC

Would you like to run the Setup Wizard?

VIRTUAL SERVICE ⌵ IP ⌵ PORTS ⌵ CONNS ⌵ PROTOCOL ⌵ METHOD ⌵ MODE ⌵

No Virtual Services configured.

Network Bandwidth

Bytes/s

Legend: RX (light blue), TX (dark blue)

RX	3k Min,	4k Avg,	32675k Total,
TX	6k Min,	7k Avg,	56693k Total,

System Load Average

System Load

1m average	0.00 Min,	0.12 Avg,	0.60 Max
5m average	0.00 Min,	0.06 Avg,	0.21 Max
15m average	0.00 Min,	0.02 Avg,	0.08 Max

Memory Usage

- You'll be asked if you want to run the Setup Wizard which can be used to configure layer 7 services. Click **Dismiss** if you're following a guide or want to configure the appliance manually or click **Accept** to start the wizard.

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

8.4. Appliance Software Update

To ensure that the appliance(s) are running the latest software version, we recommend a software update check is performed.

Determining the Current Software Version

The software version is displayed at the bottom of the WebUI as shown in the example below:

Copyright © Loadbalancer.org Inc. 2002 – 2024
ENTERPRISE VA Max - v8.11.1

English ▼

Checking for Updates using Online Update

Note

By default, the appliance periodically contacts the Loadbalancer.org update server and checks for updates. An update check can also be manually triggered as detailed below.

1. Using the WebUI, navigate to: **Maintenance > Software Update**.
2. Select **Online Update**.
3. If the latest version is already installed, a message similar to the following will be displayed:

Information: Version v8.11.1 is the current release. No updates are available

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

Note

Do not navigate away whilst the update is ongoing, this may cause the update to fail.

6. Once complete (the update can take several minutes depending on download speed and upgrade version) the following message will be displayed:

Information: Update completed successfully.

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



Using Offline Update

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

Note

Please contact support@loadbalancer.org to check if an update is available and obtain the latest offline update files.

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:

Software Update

Offline Update

The following steps will lead you through offline update.

1. Contact **Loadbalancer.org support** to obtain the offline update archive and checksum.
2. Save the archive and checksum to your local machine.
3. Select the archive and checksum files in the upload form below.
4. Click *Upload and Install* to begin the update process.

Archive: No file chosen
Checksum: No file chosen

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.

8.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22 *	SSH
TCP & UDP	53 *	DNS / GSLB
TCP & UDP	123	NTP
TCP & UDP	161 *	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000 *	Gateway service (Centralized/Portal Management)



Protocol	Port	Purpose
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

Note

The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to [Service Socket Addresses](#).

8.6. HA Clustered Pair Configuration

Loadbalancer.org recommend that load balancer appliances are deployed in pairs for high availability. In this guide a single unit is deployed first, adding a secondary unit is covered in the section [Configuring HA - Adding a Secondary Appliance](#) of the appendix.

9. Appliance Configuration for Verint WFO

9.1. VIP1 - VERINT-WFO

Configuring the Virtual Service (VIP)

- Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
- Enter the following details:

Virtual Service	
Label	<input type="text" value="VERINT-WFO"/> ?
IP Address	<input type="text" value="192.168.110.100"/> ?
Ports	<input type="text" value="443"/> ?
Protocol	
Protocol	<input type="text" value="TCP"/> ?
Forwarding	
Forwarding Method	<input type="text" value="SNAT"/> ?

- Specify the *Label* for the virtual service, e.g. **VERINT-WFO**.
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.110.100**.



5. Set the *Ports* field to **443**.
6. Set the *Protocol* to **TCP**.
7. Set the *Forwarding Method* to **SNAT**.
8. Click **Update** to create the Virtual Service.
9. Now click **Modify** next to the newly created VIP.
10. Scroll to the *Persistence* section.
 - Ensure that *Enable* is checked.
 - Set the *Timeout* to **1800**, i.e. 30 minutes.
11. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate**.
 - Set the *Protocol* to **HTTPS**.
 - Set *Request to Send* to **/wfo/control/appTest**.
 - Set *Response Expected* to **Application test has completed successfully**.
12. Click **Update**.

Defining the Real Servers (RIPs)

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

Label	<input type="text" value="VERINT-WFO-APP1"/>	?
Real Server IP Address	<input type="text" value="192.168.110.160"/>	?
Real Server Port	<input type="text" value="443"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

3. Define the *Label* for the real server as required, e.g. **VERINT-WFO-APP1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.110.160**.
5. Set the *Real Server Port* field to **443**.
6. Click **Update**.

7. Repeat these steps to add additional Real Servers as required.

10. Testing & Verification

Note









For additional guidance on diagnosing and resolving any issues you may have, please also refer to [Diagnostics & Troubleshooting](#).

10.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the Verint WFO servers) and shows the state/health of each server as well as the state of the cluster as a whole. The example below shows that the Virtual Service and all three application servers are healthy and available to accept connections:

System Overview

2024-04-03 12:48:16 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	VERINT-WFO	192.168.110.100	443	0	TCP	Layer 4	SNAT	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	VERINT-WFO-APP1	192.168.110.160	443	100	0	Drain	Halt	
	VERINT-WFO-APP2	192.168.110.161	443	100	0	Drain	Halt	
	VERINT-WFO-APP3	192.168.110.162	443	100	0	Drain	Halt	

10.2. Login to the Application

First ensure that DNS is configured so that the FQDN for the application (e.g. verint-vip.example.com) resolves to VIP1. Then verify that you're able to successfully login and access the application using the relevant URL, e.g. <https://verint-vip.example.com/wfo/control/signin>.

11. Technical Support

For more details about configuring the appliance and assistance with designing your deployment please don't hesitate to contact the support team using the following email address: support@loadbalancer.org.

12. Further Documentation

For additional information, please refer to the [Administration Manual](#).



13. Appendix

13.1. Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution. We recommend that the Primary appliance is fully configured first, then the Secondary appliance can be added to create an HA pair. Once the HA pair is configured, load balanced services must be configured and modified on the Primary appliance. The Secondary appliance will be automatically kept in sync.

 **Note**

For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the [documentation library](#)

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

Non-Replicated Settings

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

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

⚠ Important

Make sure that where any of the above have been configured on the Primary appliance, they're also configured on the Secondary.

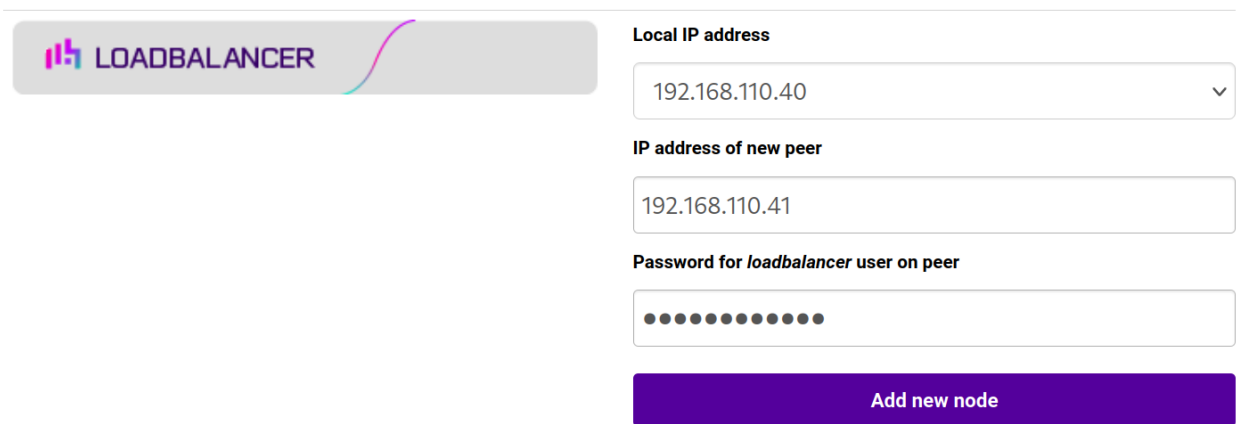
Configuring the HA Clustered Pair

📌 Note

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

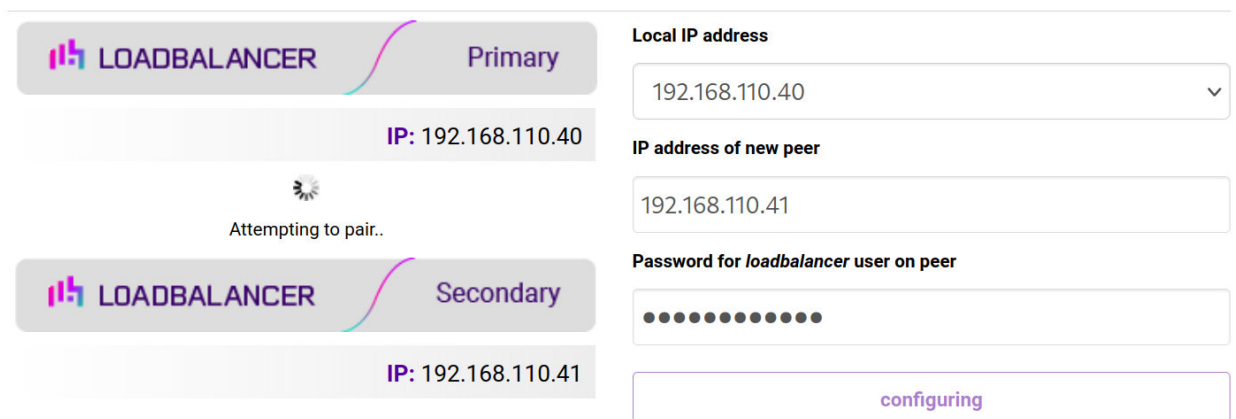
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: **Cluster Configuration > High-Availability Configuration**.

Create a Clustered Pair



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:

Create a Clustered Pair



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

High Availability Configuration - primary

The screenshot displays a configuration interface for a High Availability (HA) setup. It features two Loadbalancer appliances in a clustered pair. The Primary appliance is shown with the IP address 192.168.110.40, and the Secondary appliance is shown with the IP address 192.168.110.41. A red button labeled "Break Clustered Pair" is positioned to the right of the appliances.

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

Note

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

Note

For more details on configuring HA with 2 appliances, please refer to [Appliance Clustering for HA](#).

Note

For details on testing and verifying HA, please refer to [Clustered Pair Diagnostics](#).

14. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	04 April 2024	Initial version		JR, RJC





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](https://twitter.com/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.

