



Load Balancing IBM Watson Health iConnect Enterprise Archive by Merative

Version 1.3.0

Table of Contents

1. About this Guide	3
2. Loadbalancer.org Appliances Supported	3
3. Loadbalancer.org Software Versions Supported	3
4. IBM Watson Health iConnect Enterprise Archive by Merative Software Versions Supported	3
5. Load Balancing iConnect Enterprise Archive	3
Port Requirements	3
Deployment Concept	4
Virtual Service (VIP) Requirements	4
Deployment Mode	5
6. Loadbalancer.org Appliance – the Basics	5
Virtual Appliance	5
Initial Network Configuration	5
Accessing the WebUI	5
Main Menu Options	7
HA Clustered Pair Configuration	7
7. Appliance & iConnect Enterprise Archive Configuration	7
Appliance Configuration	7
Network Configuration	7
Floating IP Configuration (For The Cluster’s Default Gateway)	8
Configuring VIP1 – All VNA Services	8
iConnect Enterprise Archive Server Configuration	10
8. Testing & Verification	10
Automatic Failover	10
Manual Failover	10
Client Connection Tests	11
9. Technical Support	11
10. Additional Documentation	11
11. Conclusion	11
12. Appendix	12
Configuring HA - Adding a Secondary Appliance	12
Non-Replicated Settings	12
13. Document Revision History	15

1. About this Guide

This guide details the steps required to configure a highly available IBM Watson Health iConnect Enterprise Archive by Merative environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any iConnect Enterprise Archive configuration changes that are required.

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 IBM Watson Health iConnect Enterprise Archive by Merative. For full specifications of available models please refer to <https://www.loadbalancer.org/products>. Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

3. Loadbalancer.org Software Versions Supported

- V8.3.8 and later

Note

The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If using an older software version, note that the screenshots presented here may not match the WebUI exactly.

4. IBM Watson Health iConnect Enterprise Archive by Merative Software Versions Supported

- IBM Watson Health iConnect Enterprise Archive by Merative – all versions

5. Load Balancing iConnect Enterprise Archive

For high availability, IBM Watson Health recommend that a load balancer is used to enable rapid failover to the secondary iConnect Enterprise Cluster should the Primary Cluster become unavailable.

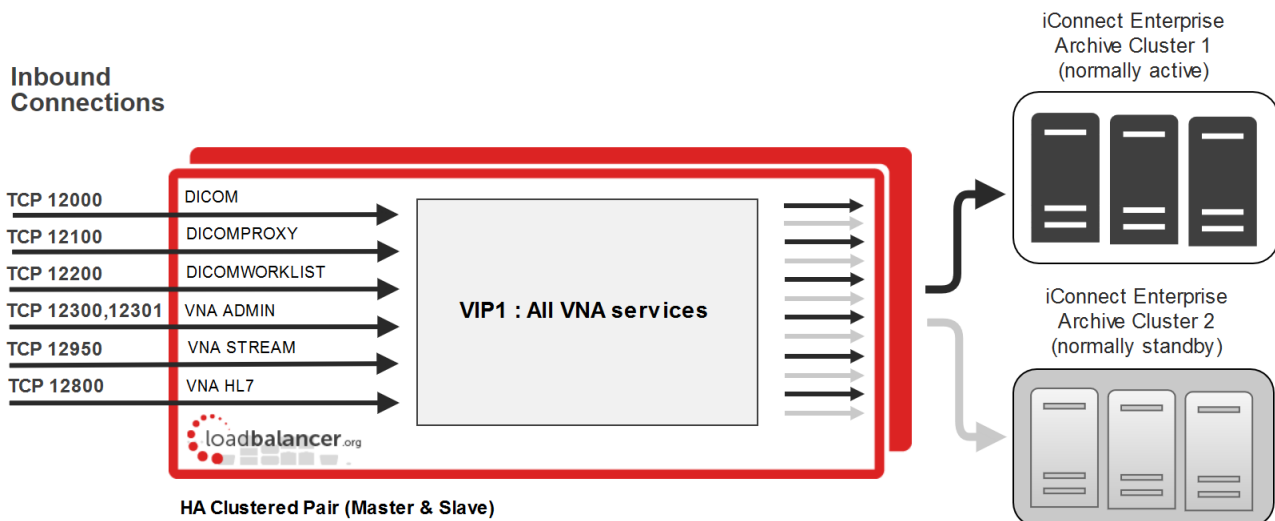
Port Requirements

The following table shows the ports used by iConnect Enterprise Archive. The load balancer must be configured to listen on the same ports.

Port	Protocols	Use
12000	TCP	DICOM
12100	TCP	DICOM Proxy
12200	TCP	DICOM Work List
12300 & 12301	TCP	VNA Admin
12950	TCP	VNA Stream
12800	TCP	VNA HL7

Deployment Concept

When iConnect Enterprise Archive is deployed with the load balancer, clients connect to the Virtual Service (VIP) on the load balancer rather than connecting directly to one of the iConnect Enterprise Archive Clusters. Under normal conditions, these connections are then forwarded to the Primary Cluster.



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 [Configuring HA - Adding a Secondary Appliance](#) for more details on configuring a clustered pair.

Should the Primary Cluster become unavailable, failover to the Secondary Cluster can be handled in either of the following ways:

- **Automatically** – In this case, health checks are configured at 30 second intervals. Should there be 10 consecutive health check failures, failover to the Secondary Cluster occurs.
- **Manually** – In this case, failover to the Secondary Cluster must be triggered manually using the 'Halt' feature in the load balancer's WebUI. Please refer to [Manual Failover](#) for more details.

Note

The way the Virtual Service's health check is configured determines which of these failover methods is used.

Virtual Service (VIP) Requirements

A single multi-port VIP is used that listens on all required ports. The VIP is configured as follows:

- Deployment mode: Layer 4 NAT (Network Address Translation) mode
- Listens on a total of 7 ports as described on the table and diagram in [Port Requirements](#)
- The health-check configuration depends on whether automatic or manual failover is required:
 - for *automatic* failover an external script is used, the script checks that *all* 7 ports are available and runs every 30 seconds, if connection to one or more of the ports fails, the health check is deemed to have failed, if there are 10 consecutive health check failures, cluster failover occurs
 - for *manual* failover the health check is set to: **No checks, always On**
- The associated Real Server is configured to be the cluster IP address of the Primary Cluster

- The fallback server is configured to be the cluster IP address of the Secondary Cluster

Deployment Mode

As mentioned above, the VIP is configured using Layer 4 NAT mode. With this mode, return traffic must pass via the load balancer. To achieve this, the default gateway of each cluster must be set to be the load balancer. For a clustered pair (our recommended configuration), an additional floating IP address must be used for this purpose. This allows the same IP address to be brought up on the Secondary appliance should an appliance failover occur.

6. Loadbalancer.org Appliance – the Basics

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.

Initial Network Configuration

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

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

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/>

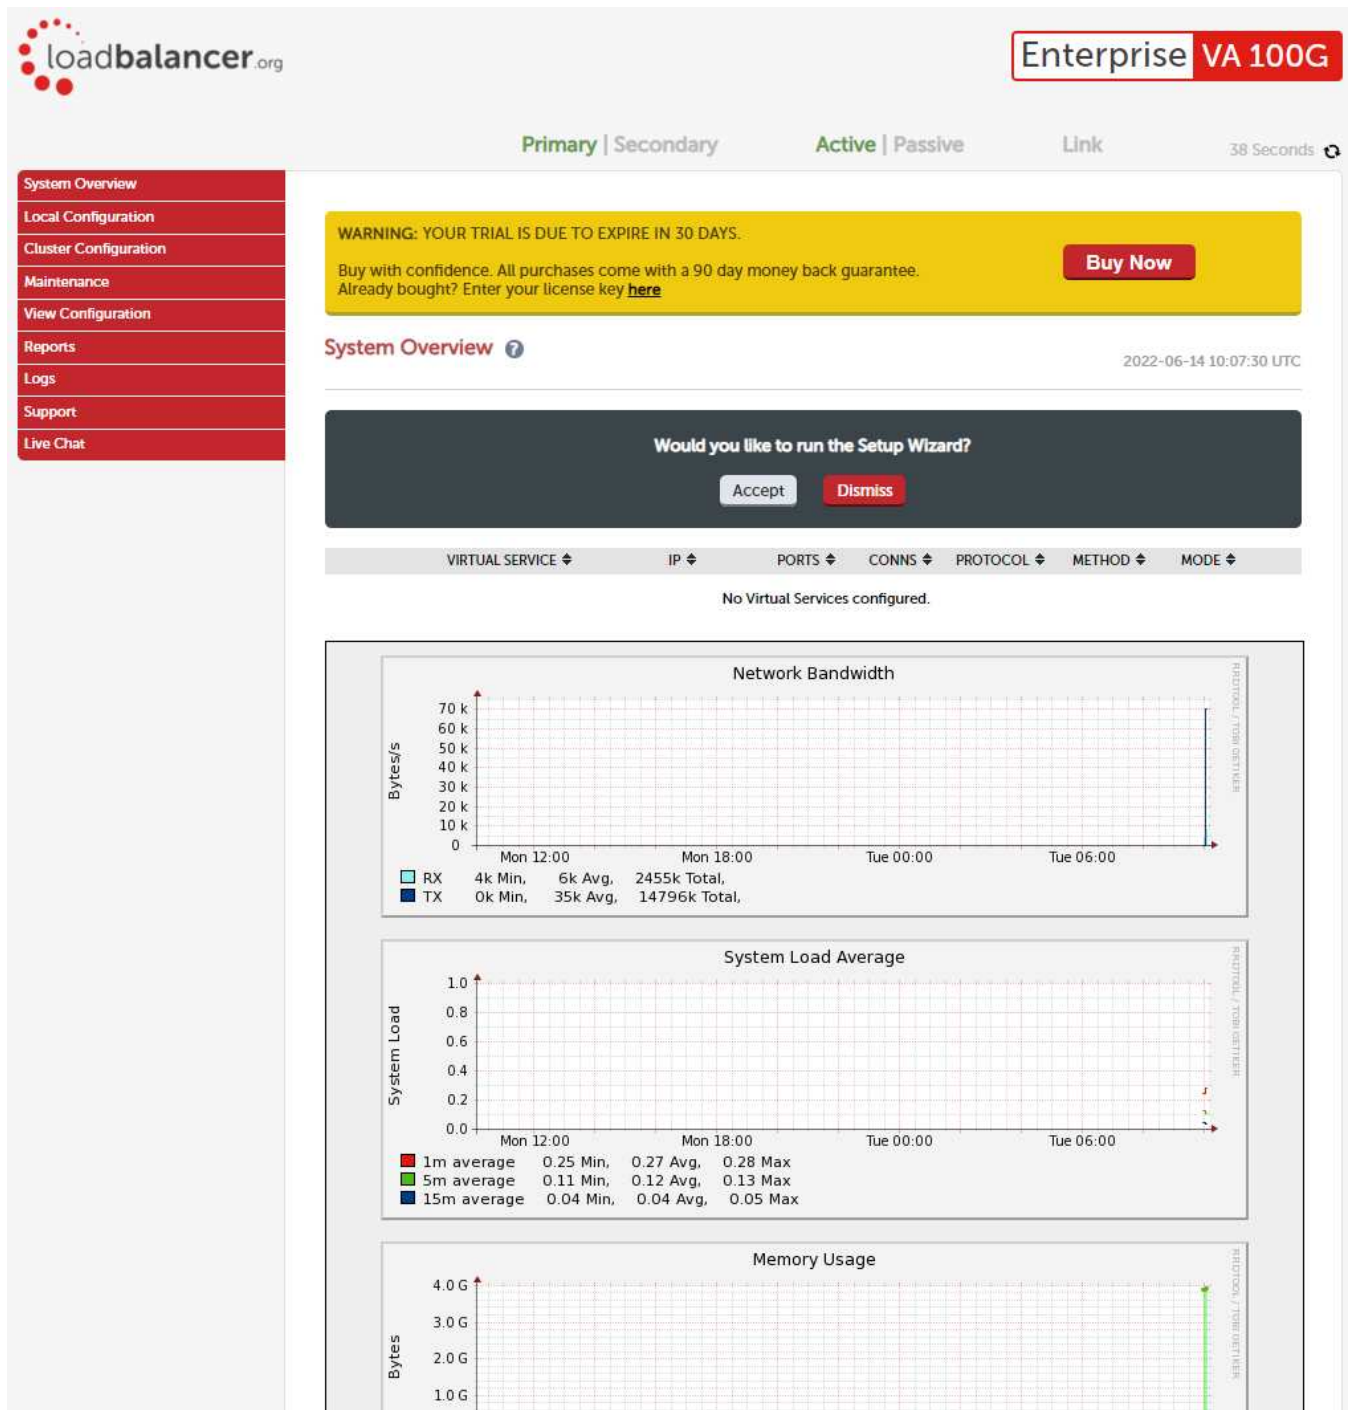
2. Log in to the WebUI:

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:



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.

3. 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

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 [Configuring HA - Adding a Secondary Appliance](#).

7. Appliance & iConnect Enterprise Archive Configuration

Appliance Configuration





Network Configuration

Layer 4 NAT mode is typically used in a 2-arm configuration where the VIP is located in one subnet and the load balanced Real Servers are located in another. This can be achieved by using two network adapters, or by creating VLAN's on a single adapter. Single arm configuration is also supported under certain conditions - for more information please refer to [Layer 4 NAT Mode](#).

To configure an additional network interface for a 2-arm configuration:

1. Using the WebUI, navigate to *Local Configuration > Network Interface Configuration*.
2. Scroll to the *IP Address Assignment* section.

IP Address Assignment

				
	eth0	eth1	eth2	eth3
	1 GB/s	1 GB/s		

eth0	<input type="text" value="192.168.100.1/24"/>	MTU <input type="text" value="1500"/> bytes
eth1	<input type="text" value="192.168.200.1/24"/>	MTU <input type="text" value="1500"/> bytes

3. Specify an appropriate IP address for **eth1** in CIDR format as shown above.

4. Click **Configure Interfaces**.

Note | There are no restrictions on which interface is used for each requirement.

Floating IP Configuration (For The Cluster's Default Gateway)

As mentioned in **Deployment Mode**, when using Layer 4 NAT mode and a clustered pair of load balancers, a floating IP address must be configured on the load balancer for use as the iConnect Enterprise Archive server's default gateway.

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IP's*.

New Floating IP

Add Floating IP

2. Enter the IP address you'd like to use as the default gateway. e.g. **192.168.100.254**.
3. Click **Add Floating IP**.

Configuring VIP1 – All VNA Services

a) Configure the External Health Check Script (used for automatic failover)

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

Health Check Details		
Name:	<input type="text" value="IBM-WHI-iConnect-Enterpris"/>	?
Type:	<input type="text" value="Virtual Service"/>	?
Template:	<input type="text" value="IBM-WHI-iConnect-Enterprise-Archive"/>	?

Primary Node Health Check Contents

2. Specify an appropriate *Name* for the health check, e.g. **IBM-WHI-iConnect-Enterprise-Archive**.
3. Set *Type* to **Virtual Service**.
4. Set *Template* to **IBM-WHI-iConnect-Enterprise-Archive**.
5. Click **Update**.

b) Setting up the Virtual Service (VIP)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Services* and click **Add a new Virtual Service**.
2. Enter the following details:

Label	<input type="text" value="VNA"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.200.100"/>	?
Ports	<input type="text" value="12000,12100,12200,12300,1"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="NAT"/>	?

3. Enter an appropriate label (name) for the VIP, e.g. **VNA**.
4. Set the *Virtual Service IP* address field to the required IP address, e.g. **192.168.200.100**.
5. Set the *Virtual Service Ports* field to **12000,12100,12200,12300,12301,12950,12800**.
6. Leave *Protocol* set to **TCP**.
7. Set the *Forwarding Method* to **NAT**.
8. Click **Update**.
9. Now click **Modify** next to the newly created VIP.
10. Scroll to the *Health Checks* section.

For *automatic* failover:

- a. Set *Check Type* to **External Script**.
- b. Set *External Script* to **IBM-WHI-iConnect-Enterprise-Archive**.

For *manual* failover:

- a. Set the *Check Type* to **No checks, Always On**.

11. Scroll to the *Fallback Server* section.
 - a. Set the *IP Address* field to the IP address of the Secondary iConnect Enterprise Archive Cluster.
 - b. Set the *Port* to **0** (numerical zero), this ensures that the fallback server (i.e. the Secondary Cluster) can receive connections on all required ports.
 - c. Enable (check) the *MASQ Fallback* checkbox.
12. Click **Update**.

c) Setting up the Real Server (RIP)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – 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="PrimaryCluster"/>	
Real Server IP Address	<input type="text" value="192.168.100.110"/>	
Real Server Port	<input type="text"/>	
Weight	<input type="text" value="100"/>	
Minimum Connections	<input type="text" value="0"/>	
Maximum Connections	<input type="text" value="0"/>	

3. Enter an appropriate label (name) for the RIP, e.g. **PrimaryCluster**.
4. Set the *Real Server IP Address* field to the IP address of the Primary iConnect Enterprise Archive Cluster.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.

iConnect Enterprise Archive Server Configuration

As mentioned in [Deployment Mode](#), when using Layer 4 NAT mode and a clustered pair of load balancers, a floating IP address must be configured for use as the default gateway. Set the default gateway of each iConnect Enterprise Archive to be this IP address.

8. Testing & Verification

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

Under normal circumstances the Primary Cluster handles all connections. Failover to the Secondary Cluster is handled automatically or manually depending on how the VIP is configured (see [Virtual Service \(VIP\) Requirements](#)).

Automatic Failover


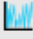

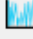
Automatic failover occurs after 5 minutes. To trigger a failover, the Primary Cluster must be continuously unavailable for this time.

Manual Failover



To trigger a failover to the Secondary Cluster, the 'Halt' option in the System Overview is used:

	REAL SERVER	IP	PORTS	WEIGHT	CONNS		
	PrimaryCluster	192.168.100.110	5.12000,12..	100	0	Drain	Halt

Once Halted, the VIP & RIP will be shown colored blue, connections will then be forwarded to the fallback server, i.e the Secondary Cluster:

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	VNA	192.168.200.100	5.12000,1..	0	TCP	Layer 4	NAT	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	PrimaryCluster	192.168.100.110	5.12000,12..	halt	0	Online (halt)		

To return to the Primary Cluster, the 'Online' option is used:

	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	PrimaryCluster	192.168.100.110	5.12000,12..	halt	0	Online (halt)		

Client Connection Tests

Ensure that clients can connect via the load balancer to the iConnect Enterprise Archive Cluster / MergePACS cluster. You'll probably need to create new DNS records or modify your existing DNS records, replacing the IP addresses of individual servers or the cluster with the IP address of the Virtual Service on the load balancer.

9. 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.

10. Additional Documentation

The Administration Manual contains much more information about configuring and deploying the appliance. It's available here: <https://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf>.

11. Conclusion

Loadbalancer.org appliances provide a very cost effective solution for highly available load balanced IBM Watson Health iConnect Enterprise Archive by Merative environments.

12. Appendix

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

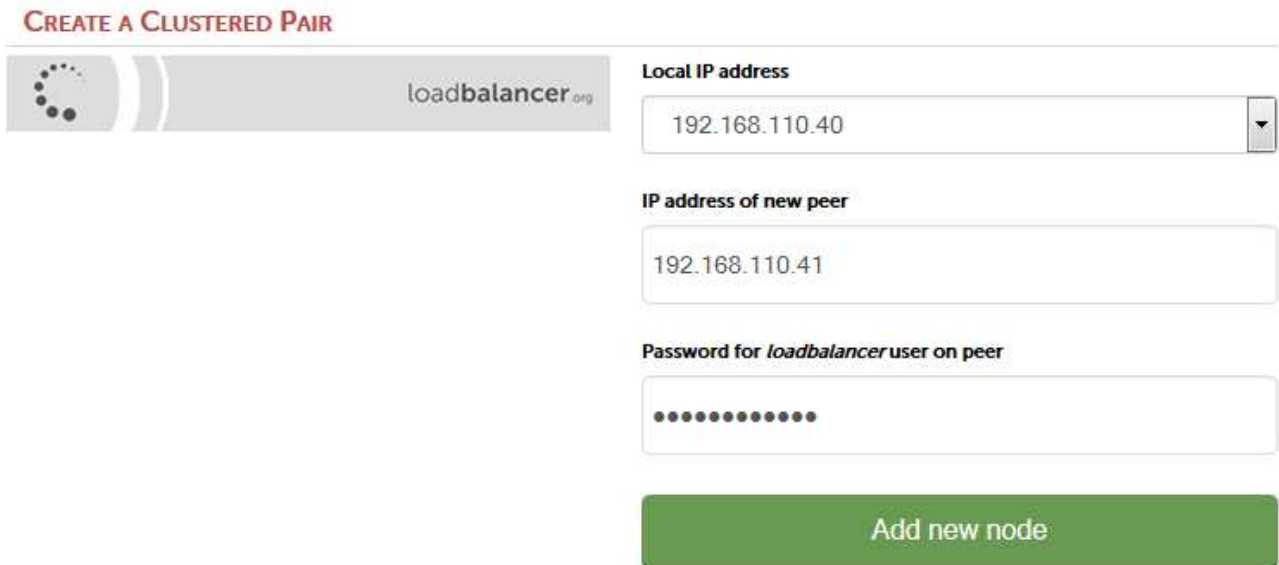
Important

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

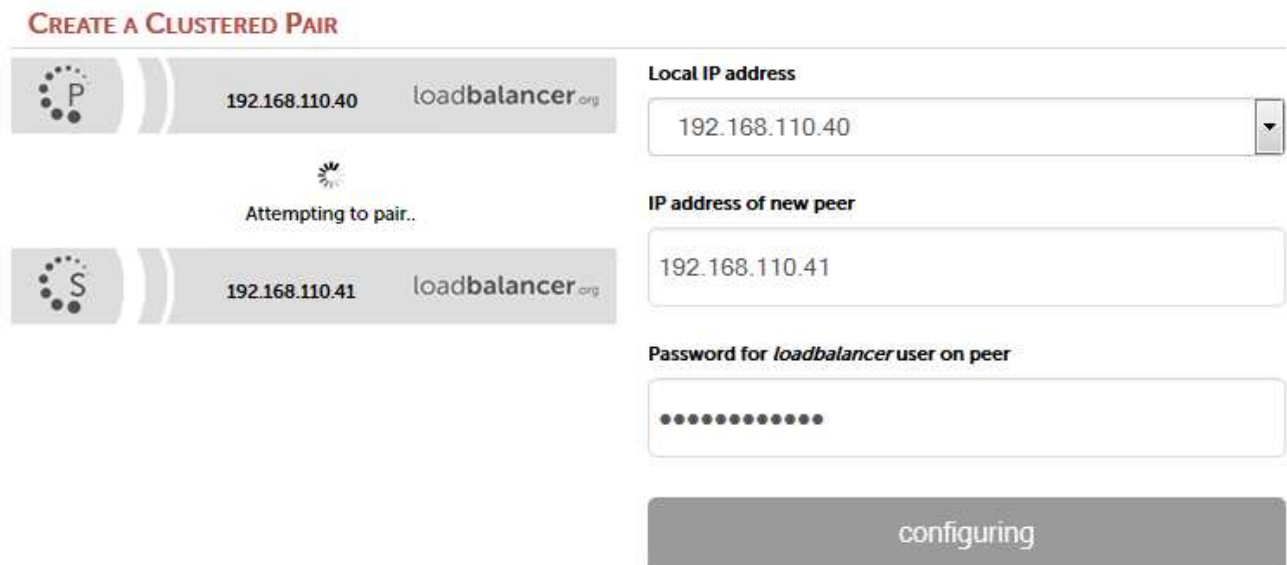
To add a Secondary node - i.e. create a highly available clustered pair:

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*.



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:



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. On the left, there are two appliance cards. The top card is labeled 'P' (Primary) and shows the IP address 192.168.110.40 and the domain loadbalancer.org. The bottom card is labeled 'S' (Secondary) and shows the IP address 192.168.110.41 and the domain loadbalancer.org. To the right of these cards is a prominent red button with the text '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](#).

13. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.1.0	2 August 2019	Styling and layout	General styling updates	RJC
1.1.1	20 August 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.2.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.2.1	11 May 2022	Updated external health check related content to reflect latest software version	New software release	RJC
1.3.0	6 September 2022	Renamed document and amended references to the product	Product acquisition	AH

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