

Load Balancing Inspirata Dynamyx

Version 1.2.0



Table of Contents

1. About this Guide	3
2. Loadbalancer.org Appliances Supported	3
3. Software Versions Supported	3
3.1. Loadbalancer.org Appliance	3
3.2. Inspirata Dynamyx	3
4. Inspirata Dynamyx	3
5. Load Balancing Overview	3
5.1. Basic Concepts	3
5.2. Load Balancer Deployment	4
5.3. Load Balancing Deployment Modes	5
Our Recommendation	5
5.4. Load Balanced Ports & Services	5
5.5. Persistence (Server Affinity)	6
5.6. Server Health Checking	6
6. Loadbalancer.org Appliance – the Basics	6
6.1. Virtual Appliance	6
6.2. Initial Network Configuration	7
6.3. Accessing the Appliance WebUI	7
Main Menu Options	8
6.4. Appliance Software Update	9
Determining the Current Software Version	9
Checking for Updates using Online Update	9
Using Offline Update	9
6.5. Ports Used by the Appliance	10
6.6. Clustered Pair Configuration	11
7. Appliance & Server Configuration	11
7.1. Load Balancing Mode	11
7.2. Health Check Configuration	11
7.3. Load Balancing Web Nodes (IIS)	11
Setting up the Virtual Service (VIP)	11
Setting up the Real Servers (RIPs)	12
7.4. Load Balancing Digital Archive	12
Setting up the Virtual Service (VIP)	12
Setting up the Real Servers (RIPs)	13
Restart HAProxy	14
8. Testing & Verification	14
8.1. Using the System Overview	14
8.2. System Logs & Reports	15
9. Technical Support	15
10. Further Documentation	15
11. Appendix	16
11.1. Configuring HA - Adding a Secondary Appliance	16
Non-Replicated Settings	16
Adding a Secondary Appliance - Create an HA Clustered Pair	17
11.2. Solving the ARP Problem	18
12. Document Revision History	19

1. About this Guide

This guide details the steps required to configure a load balanced Inspirata Dynamyx environment utilizing Loadbalancer.org appliances.

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 for load balancing Inspirata Dynamyx. 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. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.4 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.

3.2. Inspirata Dynamyx

- All versions

4. Inspirata Dynamyx

Dynamyx from Inspirata affords an 'open' architecture purposely designed to enable healthcare providers to arrive at their preferred blend of laboratory and diagnostic technologies. In providing a vendor-agnostic, end-to-end digital pathology solution deliberately architected for multi-vendor environments, Inspirata assist high-throughput clinical laboratories navigate the specific and frequent challenges of system interoperability while also allowing for increased resource flexibility and a future-proofing of investments culminating in an overall lower total cost of ownership.

5. Load Balancing Overview

5.1. Basic Concepts

To provide resilience and high availability, multiple Virtual Services (VIPs) are configured for the various protocols and systems. Clients and systems then connect to these VIPs rather than directly to the application servers. Each VIP can be configured in one of the following ways:

- **Load balanced mode**



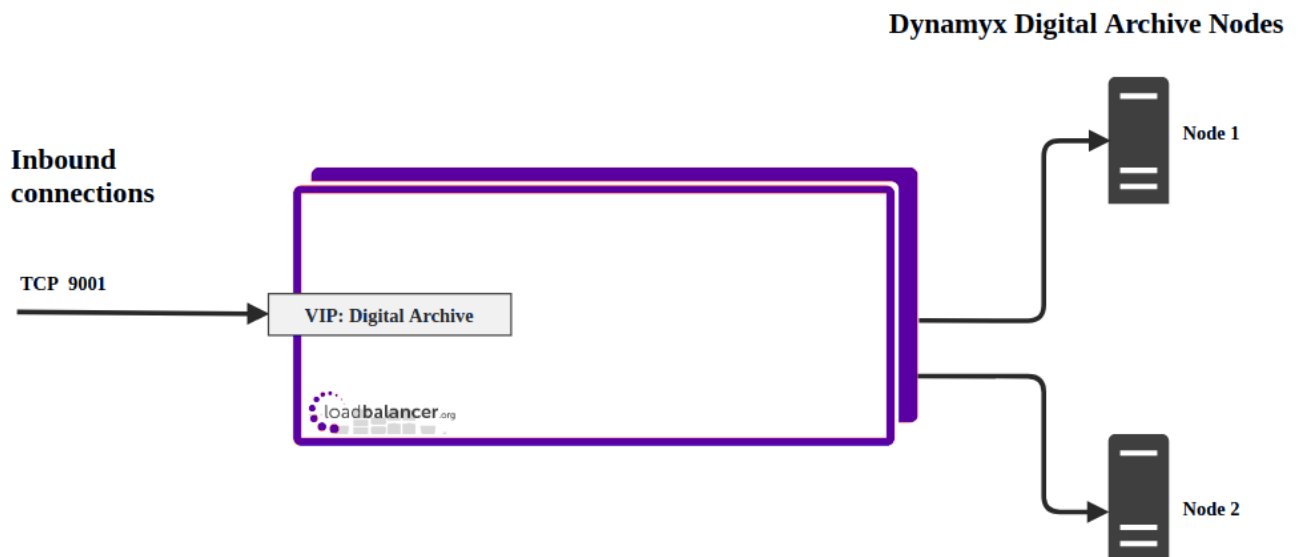
Load is distributed across all configured servers/endpoints

- **Failover mode**

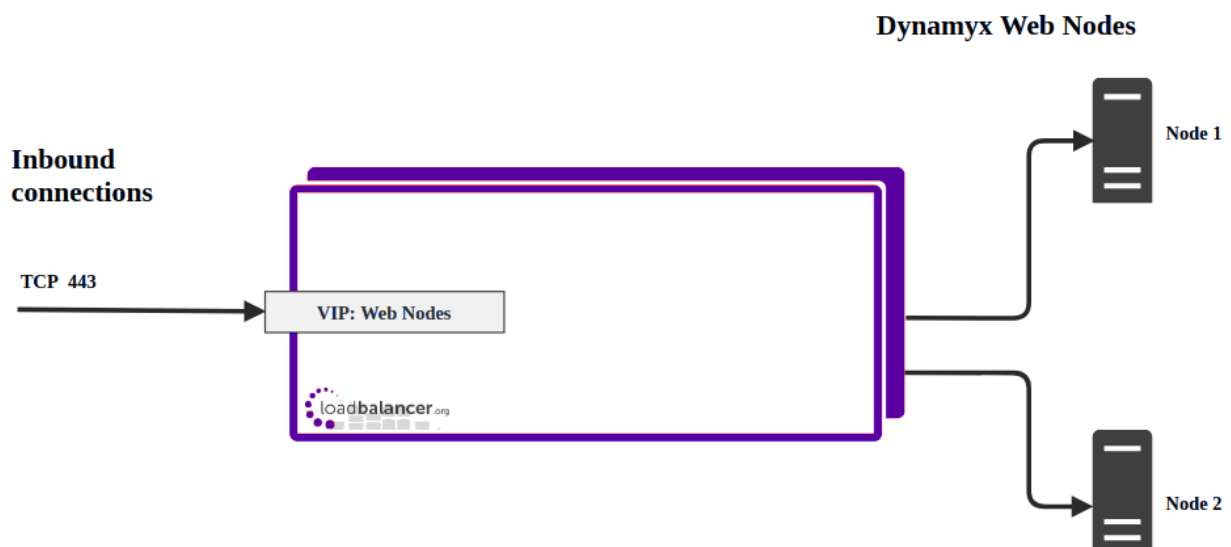
The second/backup server is used only when the first server/endpoint fails

5.2. Load Balancer Deployment

The following diagram shows a simplified view of Inspirata Dynamyx Digital Archive in load balancing mode:



The following diagram shows a simplified view of Inspirata Dynamyx Digital Archive Web Access in load balancing mode:



Notes

1. **VIP (Virtual IP)** – This is IP address presented by the load balancer. Clients and other systems connect to

this rather than directly to the back end servers/endpoints.

2. A single load balancer appliance can be used to load balance all services. More than one load balancer appliance may be required depending on throughput and physical network topology.

5.3. Load Balancing Deployment Modes

The load balancer supports the following deployment modes:

Layer 4 DR Mode – This mode offers the best performance and requires limited physical Real Server changes. The load balanced application must be able to bind to the Real Server's own IP address and the VIP at the same time. This mode requires the **ARP Problem** to be solved as described [here](#). Layer 4 DR mode is transparent, i.e. the Real Servers will see the source IP address of the client.

Layer 4 NAT Mode – This mode is also a high performance solution but not as fast as DR mode. It requires the default gateway of each Real Server to be the load balancer and supports both one-arm and two-arm configurations. Layer 4 NAT mode is transparent, i.e. the Real Servers will see the source IP address of the client.

Layer 4 SNAT Mode – This mode is also a high performance solution but not as fast as the other layer 4 modes. It does not require any changes to the Real Servers and can be deployed in one-arm or two-arm mode. This mode is ideal for example 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. Layer 4 SNAT mode is non-transparent, i.e. the Real Servers will see the source IP address of the load balancer.

Layer 7 SNAT Mode – This mode offers greater flexibility but at lower performance levels. It supports HTTP cookie insertion, RDP cookies, Connection Broker integration and works very well with either Pound or STunnel when SSL termination is required. It also enables content switching and header manipulation rules to be implemented. It does not require any changes to the Real Servers and can be deployed in one-arm or two-arm mode. HAProxy is a high performance solution, but since it operates as a full proxy it cannot perform as fast as the layer 4 solutions. Layer 7 SNAT mode is non-transparent by default, i.e. the Real Servers will see the source IP address of the load balancer. This mode can be made transparent through the use of TProxy.

Our Recommendation

When load balancing Inspirata Dynamyx, we recommend that Layer 7 SNAT mode is used. This mode offers high performance with no real server or network changes required since replies go via the same path as the ingress traffic. Using a layer 7 configuration will lose client source IP address transparency. If source IP transparency is required, i.e. if the back end servers **must** see inbound traffic as originating from the client's true source address, then it is suggested to use either a layer 4 DR or NAT mode configuration. Ultimately, the final choice depends on your specific requirements and infrastructure.

Note

If you are using Microsoft Windows Real Servers (i.e. the backend servers) make sure that Windows **NLB** (Network Load Balancing) is **completely disabled** to ensure that this does not interfere with the operation of the load balancer.

5.4. Load Balanced Ports & Services

The following table shows the typical ports/services that are load balanced.



Port	Protocols	Use
443	TCP/HTTPS	Web based access to the digital archives
9001	TCP/Digital Archive	Communication between digital archives

5.5. Persistence (Server Affinity)

Source IP address persistence is used for the HTTPS-based protocol only. This ensures that a particular client will connect to the same load balanced server/endpoint for the duration of a session.

5.6. Server Health Checking

The default health check used for new virtual services is a TCP 'connect to port' check. This verifies that a given port is open and accepting connections. However, it does not necessarily guarantee that the associated service is fully operational. Also, repeated ongoing connections to a service's port may cause multiple log entries reporting incomplete connections or other issues.

More robust service-oriented health checks can be configured for both layer 4 and layer 7 services using the negotiate option. This effectively tests and verifies the running service.

For example, the load balancer can be configured to look for specific content on an HTTP web page on the load balanced Real Server. If the page can be opened and the content can be found then the check will have passed. If not, the check will fail and the server/endpoint will be marked as down.

If the service running is not HTTP based, a custom page could be setup on the load balanced servers that simply indicates service status. The load balancer can then use this for health checking.

The page to check and the content to be verified can easily be configured for layer 4 and layer 7 VIPs using the WebUI. Select the required negotiate option and configure the required settings. For more details on configuring health-checks please refer to [Real Server Health Monitoring & Control](#).

Note

The configuration examples in this guide use a TCP 'connect to port' check (the default) to check the health of load balanced servers.

6. Loadbalancer.org Appliance – the Basics

6.1. 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.

6.2. Initial Network Configuration

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

Important

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

6.3. 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](#).

Note

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

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 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](#).

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


Username: loadbalancer

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

Note

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

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



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



Note

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

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

6.4. 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 – 2023
ENTERPRISE VA Max - v8.9.0

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 v8.9.0 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.

6.5. 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
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



6.6. 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 & Server Configuration

7.1. Load Balancing Mode

As mentioned in [Load Balancing Deployment Modes](#), Virtual Services can be configured in one of four fundamental ways: *Layer 4 DR mode*, *Layer 4 NAT mode*, *Layer 4 SNAT mode*, and *Layer 7 SNAT mode*. The following sections illustrate how to configure the Virtual Services using the recommended load balancing mode, *Layer 7 SNAT mode*. If a different load balancing mode is required for a particular VIP then please don't hesitate to contact our support team: support@loadbalancer.org.

7.2. Health Check Configuration

As mentioned in [Server Health Checking](#), health checks can be configured in several different ways. The sections below all use a TCP 'connect to port' check using the port of the service in question.

7.3. Load Balancing Web Nodes (IIS)

(Using Layer 7 SNAT Mode)

Setting up the Virtual Service (VIP)

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

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="IIS_VIP"/>	?
IP Address	<input type="text" value="192.168.86.177"/>	?
Ports	<input type="text" value="443"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **IIS_VIP**.
4. Set the *IP Address* field to the required IP address, e.g. **192.168.86.177**.



5. Set the **Ports** field to the required port(s), e.g. **443**.
6. Set **Protocol** to **TCP**.
7. Click **Update**.
8. Now click **Modify** next to the newly created Virtual Service.
9. Ensure **Persistence Mode** is set to **Source IP**.
10. Set the **Check Type** to **Connect to port**.
11. Click **Update**.

Setting up the Real Servers (RIPs)

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

Layer 7 Add a new Real Server - IIS_VIP

Label	<input type="text" value="IIS1"/>	?
Real Server IP Address	<input type="text" value="192.168.86.50"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Enter an appropriate name (Label) for the first IIS server, e.g. **IIS1**.
4. Change the **Real Server IP Address** field to the required IP address, e.g. **192.168.86.50**.
5. Click **Update**.
6. Repeat these steps to add additional server(s).

7.4. Load Balancing Digital Archive

Setting up the Virtual Service (VIP)

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

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="DA_VIP"/>	?
IP Address	<input type="text" value="192.168.86.166"/>	?
Ports	<input type="text" value="9001"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/> ▼	?

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **DA_VIP**.
4. Set the *IP Address* field to the required IP address, e.g. **192.168.86.166**.
5. Set the *Ports* field to the required port, e.g. **9001**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created Virtual Service.
9. Ensure *Persistence Mode* is set to **None**.
10. Set the *Health Checks* to **Connect to port**.
11. Scroll down to the *Other* section and enable the **Timeout** check box.
12. Set the *Client* and *Real Server Timeout* to **5m**.
13. Click **Update**.

Setting up the Real Servers (RIPs)

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

Layer 7 Add a new Real Server - DA_VIP

Label	<input type="text" value="DA1"/>	?
Real Server IP Address	<input type="text" value="192.168.86.50"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

CancelUpdate

3. Enter an appropriate name (Label) for the first Digital Archive server, e.g. **DA1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.86.50**.
5. Click **Update**.
6. Repeat these steps to add additional server(s).

Restart HAProxy

1. To apply the new settings, restart HAProxy using the WebUI option **Maintenance > Restart Services** and clicking **Restart HAProxy**.

Note

If you will be configuring additional layer 7 services, you can restart HAProxy at the end once all layer 7 Virtual Services and Real Servers have been defined.

8. Testing & Verification

Note

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

8.1. Using the System Overview

Verify that all virtual services and their associated real servers are reported as online/healthy (green) as shown below:

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
↑	IIS_VIP	192.168.86.177	443	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	IIS1	192.168.86.50	443	100	0	Drain	Halt	
↑	IIS2	192.168.86.51	443	100	0	Drain	Halt	
↑	DA_VIP	192.168.86.166	9001	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	DA1	192.168.86.50	9001	100	0	Drain	Halt	
↑	DA2	192.168.86.51	9001	100	0	Drain	Halt	

If certain servers are down, i.e. failing their health checks, they will show up as red, as shown below:

↓	DA_VIP	192.168.86.166	9001	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↓	DA1	192.168.86.50	9001	100	0	Drain	Halt	
↓	DA2	192.168.86.51	9001	100	0	Drain	Halt	

8.2. System Logs & Reports

Various system logs & reports can be used to help diagnose problems and help solve appliance issues. Logs can be accessed using the WebUI options: *Logs & Reports*.

9. Technical Support

If you have any questions regarding the appliance or would like assistance designing your deployment, please don't hesitate to contact our support team: support@loadbalancer.org.

10. Further Documentation

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

11. Appendix

11.1. 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 configured first and 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.


Adding a Secondary Appliance - Create an 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

 **LOADBALANCER**

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 in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:


Create a Clustered Pair

 **LOADBALANCER**

Primary

IP: 192.168.110.40

Attempting to pair..

 **LOADBALANCER**

Secondary

IP: 192.168.110.41

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

 LOADBALANCER  Primary	Break Clustered Pair
IP: 192.168.110.40	
 LOADBALANCER  Secondary	Make Active
IP: 192.168.110.41	

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](#).

11.2. Solving the ARP Problem

Layer 4 DR mode works by changing the MAC address of the inbound packets to match the Real Server selected by the load balancing algorithm. To enable DR mode to operate:

- Each Real Server must be configured to accept packets destined for both the VIP address and the Real Servers IP address (RIP). This is because in DR mode the destination address of load balanced packets is the VIP address, whilst for other traffic such as health checks, administration traffic etc. it's the Real Server's own IP address (the RIP). The service/process (e.g. IIS, httpd) must respond to both addresses.
- Each Real Server must be configured so that it does not respond to ARP requests for the VIP address – only the load balancer should do this.

Configuring the Real Servers in this way is referred to as '**Solving the ARP problem**'. The steps required depend on the particular OS being used.

For detailed steps on solving the ARP problem for Linux, Windows and various other operating systems, please refer to [DR Mode Considerations](#).

12. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	4 March 2020	Initial document creation		IBG
1.0.1	1 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.1.0	1 December 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.1.1	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.1.2	5 January 2023	Combined software version information into one section Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section	Housekeeping across all documentation	AH
1.1.3	2 February 2023	Updated screenshots	Branding update	AH
1.1.4	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.2.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH



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.

