

Load Balancing Medical Imaging & Information System Protocols

Version 1.4.0



Table of Contents

1. About this Guide	4
2. Loadbalancer.org Appliances Supported	4
3. Software Versions Supported	4
3.1. Loadbalancer.org Appliance	4
3.2. Medical Systems Supported	4
4. Medical Imaging and Information Systems & Components	4
4.1. Picture Archiving and Communication System (PACS)	4
4.2. Vendor Neutral Archive (VNA)	4
4.3. Imaging Modalities	5
4.4. Health Care Administration Systems	5
4.5. Workstations/Viewers	5
5. Medical Information System Standards & Protocols	5
5.1. DICOM	5
5.2. HL7	6
5.3. IHE XDS	6
6. Load Balancing Overview	6
6.1. Basic Concepts	6
6.2. Load Balancer Deployment	7
6.3. Load Balancing Deployment Modes	8
6.3.1. Our Recommendation	9
6.4. Load Balanced Ports & Services	9
6.5. Persistence (Server Affinity)	10
6.6. Server Health Checking	10
7. Loadbalancer.org Appliance – the Basics	10
7.1. Virtual Appliance	10
7.2. Initial Network Configuration	11
7.3. Accessing the Appliance WebUI	11
7.3.1. Main Menu Options	12
7.4. Appliance Software Update	13
7.4.1. Online Update	13
7.4.2. Offline Update	13
7.5. Ports Used by the Appliance	14
7.6. Clustered Pair Configuration	15
8. Appliance & Server Configuration	15
8.1. Load Balancing Mode	15
8.2. Health-Check Configuration	15
8.3. Load Balancing DICOM	15
8.3.1. Setting up the Virtual Service (VIP)	15
8.3.2. Setting up the Real Servers (RIPs)	16
8.3.3. Configuring the load balanced DICOM servers	16
8.4. Load Balancing HL7	16
8.4.1. Setting up the Virtual Service (VIP)	16
8.4.2. Setting up the Real Servers (RIPs)	17
8.5. Load Balancing XDS (Registry & Repository)	18
8.5.1. Setting up the Virtual Service (VIP)	18
8.5.2. Setting up the Real Servers (RIPs)	19
8.5.3. Configuring the load balanced XDS servers	19
8.6. Load Balancing HTTPS	19

8.6.1. Setting up the Virtual Service (VIP)	19
8.6.2. Setting up the Real Servers (RIPs)	20
8.7. Finalizing Appliance Settings	21
8.7.1. Configure HAProxy Timeouts	21
8.7.2. Restart HAProxy	21
9. Testing & Verification	21
9.1. Using the System Overview	21
9.2. System Logs & Reports	22
10. Technical Support	22
11. Further Documentation	22
12. Appendix	23
12.1. Configuring HA - Adding a Secondary Appliance	23
12.1.1. Non-Replicated Settings	23
12.1.2. Configuring the HA Clustered Pair	24
12.2. Solving the ARP Problem	25
13. Document Revision History	26

1. About this Guide

This guide details the steps required to configure a load balanced Medical Imaging and Information System environment utilizing Loadbalancer.org appliances. It includes details on load balancing DICOM, HL7 & IHE XDS.

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 Medical Imaging and Information Systems. For full specifications of available models please refer to <https://www.loadbalancer.org/products/enterprise>.

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. Medical Systems Supported

- Any systems that utilizes medical system standards & protocols such as DICOM, HL7, XDS, XDS-1

4. Medical Imaging and Information Systems & Components

4.1. Picture Archiving and Communication System (PACS)

A picture archiving and communication system (PACS) is a medical imaging technology which provides economical storage and convenient access to images from multiple imaging modalities. Electronic images and reports are transmitted digitally via PACS; this eliminates the need to manually file, retrieve, or transport film jackets. The universal format for PACS image storage and transfer is DICOM (Digital Imaging and Communications in Medicine). Non-image data, such as scanned documents, may be incorporated using consumer industry standard formats like PDF (Portable Document Format), once encapsulated in DICOM.

4.2. Vendor Neutral Archive (VNA)

A VNA is an archival system that can be used to store virtually any type of digital data irrespective of the original



source of the data. The VNA will also serve that data to any requesting system (with proper authentication and authorization) without regard to the vendor of the system requesting the data. It is the independence from the vendors that provide the source data or the data request that renders it "vendor neutral." VNAs are also sometimes referred to as a PACS Neutral Archive.

VNAs are distinguished from picture archiving and communication systems by functioning more as a central store for images from many sources and diverse vendors. PACS are proprietary systems that share little, if at all, and are typically scattered around a health-care system.

4.3. Imaging Modalities

These are the various sources of medical images and include equipment such as:

- CT (Computed Tomography) scanners
- MRI (Magnetic Resonance Imaging) scanners
- PET (Positron Emission Tomography) scanners
- X-RAY scanners
- Ultrasound scanners

4.4. Health Care Administration Systems

Various health-care systems are used within hospitals and ideally are interfaced to share data using protocols such as HL7, these include:

- HIS – Hospital Information System
- RIS – Radiology Information System
- PAS – Patient Administration System
- ADT – Admission, Discharge and Transfer System

4.5. Workstations/Viewers

To enable access to stored images and associated data, DICOM workstations are used. These connect directly to the DICOM source. Viewer servers are also used which enable client PCs to view DICOM images using a web browser via HTTPS.

5. Medical Information System Standards & Protocols

5.1. DICOM

The Digital Imaging and Communications in Medicine (DICOM) Standard describes the means of formatting, storing and exchanging medical images and image related information to facilitate the connectivity of medical devices and systems. The DICOM Standard endorsed by the National Electrical Manufacturers Association (NEMA) is a result of joint efforts of users and manufacturers of medical imaging and health-care information technology.

Today, virtually all imaging devices (Modalities) that are used in radiology, such as CT, MRI, Ultrasound, RF, and



other digital rooms, supports the DICOM standard for the exchange of images and related information.

5.2. HL7

Health Level Seven (HL7) is an American National Standards Institute accredited Standards Developing Organization (SDO) operating in the health-care arena. Since its inception, HL7 has specified standards for a large number of application areas. HL7 standards cover generic application fields such as patient administration, patient care, order entry, results reporting, document and financial management. In addition to that, HL7 addresses the departmental information system communication needs of clinical specialties like laboratory medicine and diagnostic imaging. HL7 is the language used for communication between health-care IT systems.

5.3. IHE XDS

Cross-Enterprise Document Sharing (XDS) is focused on providing a standards-based specification for managing the sharing of documents between any health-care enterprise, ranging from a private physician office to a clinic to an acute care in-patient facility and personal health record systems. This is managed through federated document repositories and a document registry to create a longitudinal record of information about a patient within a given clinical affinity domain. These are distinct entities with separate responsibilities:

A Document Repository is responsible for storing documents in a transparent, secure, reliable and persistent manner and responding to document retrieval requests.

A Document Registry is responsible for storing information about those documents so that the documents of interest for the care of a patient may be easily found, selected and retrieved irrespective of the repository where they are actually stored.

Documents are provided by one or more Document Sources.

They are then accessed by one or more Document Consumers.

XDS/XDS-I enables sharing of non-DICOM (i.e. JPEG images, scanned documents, text-based documents) information across disparate health-care systems.

6. Load Balancing Overview

6.1. 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

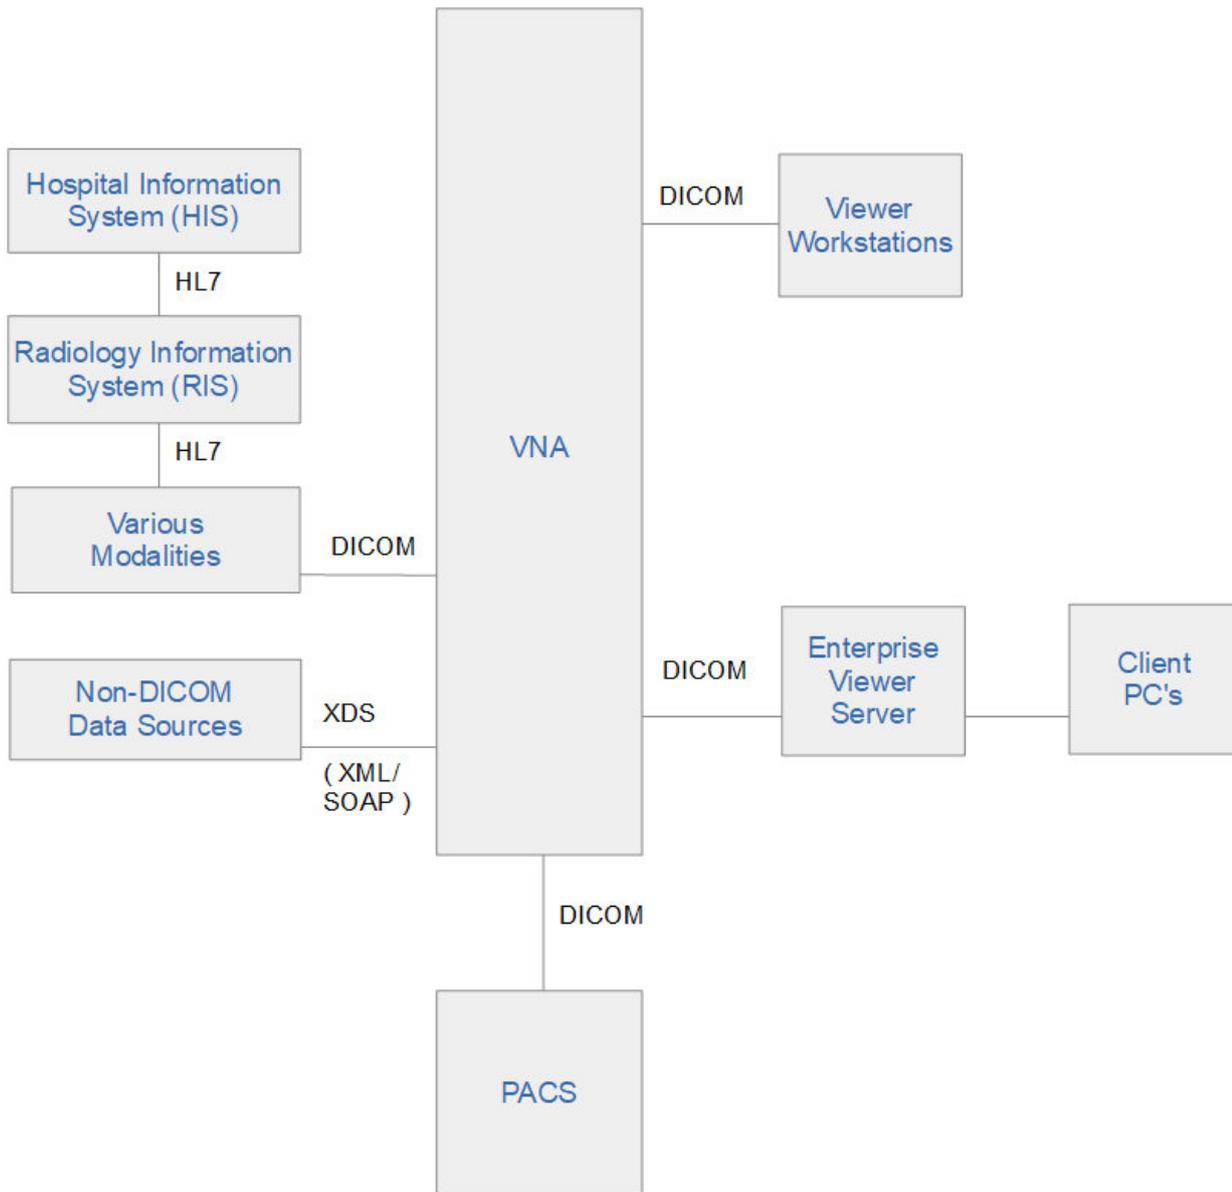
- **Failover mode**

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

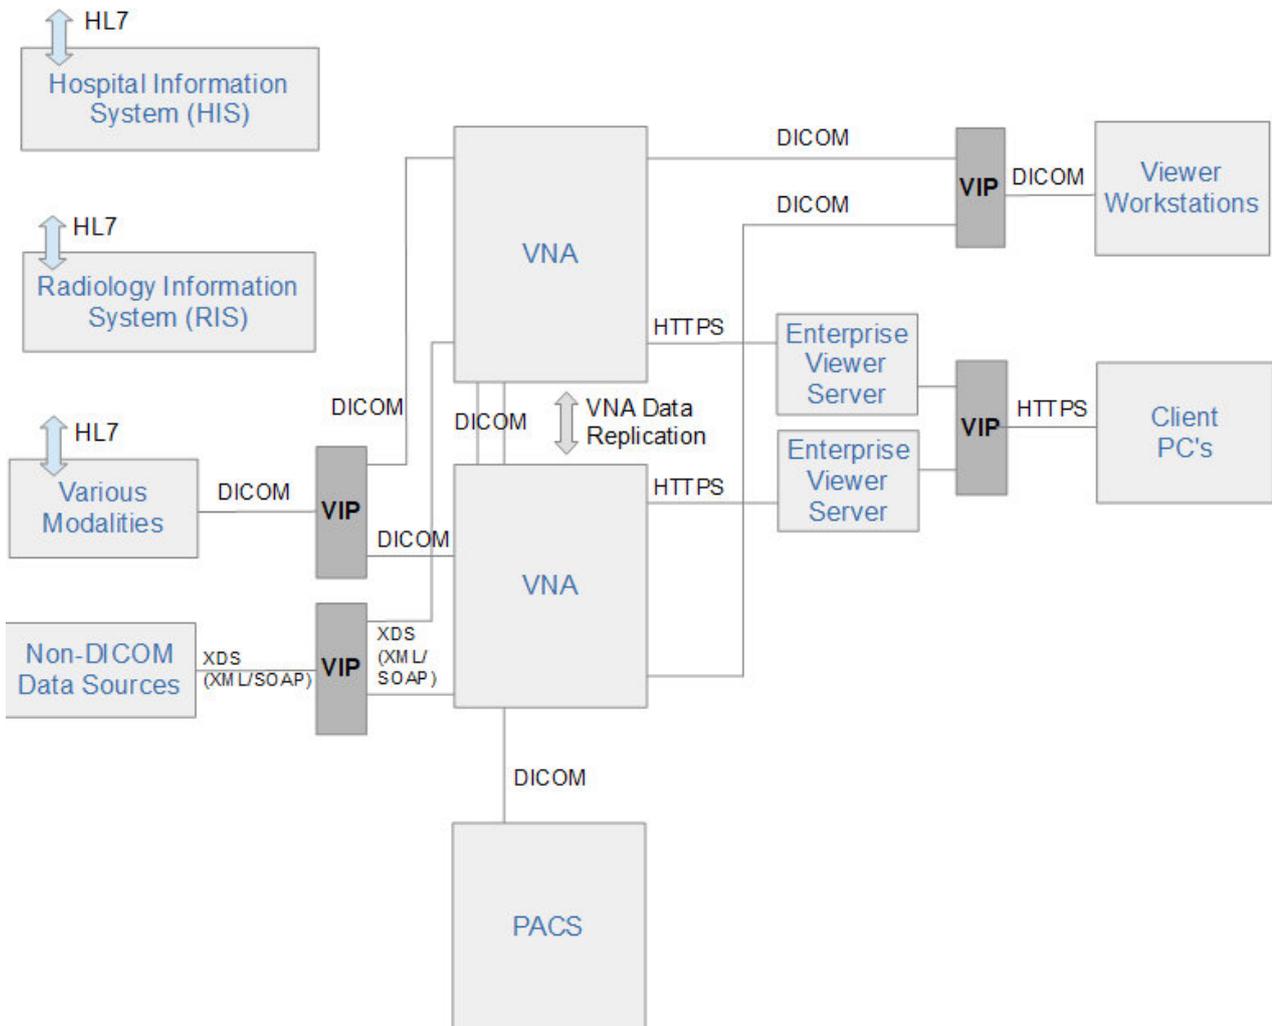
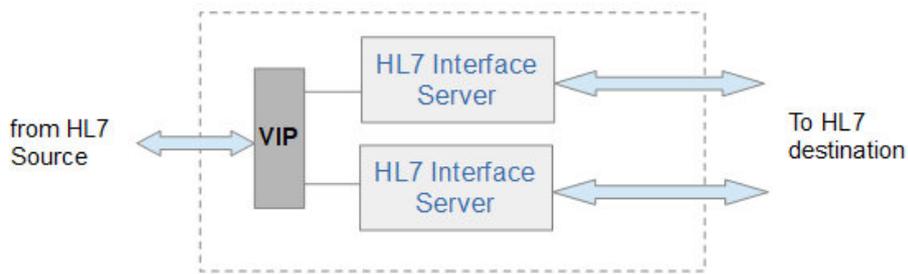


6.2. Load Balancer Deployment

The following diagram shows a simplified view of a typical Medical Imaging & Information System without load balancing:



The diagram below shows a highly available system that utilizes multiple system components and load balancing:



Notes

- **VIP (Virtual IP)** – This is the IP address presented by the load balancer. Clients and other systems connect to this rather than directly to the back end servers/endpoints.
- 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.

6.3. 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.

6.3.1. Our Recommendation

Where possible we recommend that Layer 4 Direct Routing (DR) mode is used. This mode offers the best possible performance since replies go directly from the Real Servers to the client, not via the load balancer. It's also relatively simple to implement. Ultimately, the final choice does depend 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.

6.4. Load Balanced Ports & Services

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

Port	Protocols	Use
104	TCP/DICOM	exchange of images and related information
11112	TCP/DICOM	exchange of images and related information
2575	TCP/HL7/MLLP	communication between health-care IT systems
443	TCP/HTTPS	client viewer connectivity
17035 *	TCP/XDS/SOAP/XML	XDS repository
17035 *	TCP/XDS/SOAP/XML	XDS registry

(*) There is no specific standard port for XDS data. Either use the suggested port (17035) or choose an alternative.



6.5. Persistence (Server Affinity)

Source IP address persistence is used for all protocols. This ensures that a particular client will connect to the same load balanced server/endpoint for the duration of the session.

6.6. Server Health Checking

The default health-check used for new VIPs is a TCP port connect. This verifies that the port is open and accepting connections. However, it does not necessarily guarantee that the associated service is fully operational. Also, repeated ongoing connections to the service 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, 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 port connect (the default) to check the health of load balanced servers.

7. Loadbalancer.org Appliance – the Basics

7.1. 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.

7.2. 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.

7.3. 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:



Primary | Secondary Active | Passive Link 8 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](#)

Buy Now

System Overview ? 2025-05-08 12:37:21 UTC

Would you like to run the Setup Wizard?

Accept
Dismiss

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

No Virtual Services configured.

Network Bandwidth

RX	28 Min,	2713 Avg,	27344772 Total,
TX	0 Min,	13777 Avg,	138872181 Total,

System Load Average

1m average	0.00 Min,	0.08 Avg,	0.68 Max
5m average	0.00 Min,	0.04 Avg,	0.30 Max
15m average	0.00 Min,	0.02 Avg,	0.12 Max

Memory Usage

- 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.

7.3.1. 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 creating 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

7.4. Appliance Software Update

We recommend that the appliance is kept up to date to ensure that you benefit from the latest bug fixes, security updates and feature improvements. Both online and offline update are supported.

 **Note**

For full details, please refer to [Appliance Software Update](#) in the Administration Manual.

 **Note**

Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

7.4.1. Online Update

The appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. This is the default behavior and can be disabled if preferred. If an update is found, a notification similar to the example below will be displayed at the top of the WebUI:

Information: Update 8.13.1 is now available for this appliance.

Online Update

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

 **Important**

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

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

Information: Update completed successfully. Return to **system overview**.

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

7.4.2. Offline Update

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

To check for the latest version, please refer to our product roadmap page available [here](#). To obtain the latest offline update files contact support@loadbalancer.org.

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.

7.5. 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)
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](#)



7.6. 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](#).

8. Appliance & Server Configuration

8.1. Load Balancing Mode

As mentioned in [Load Balancing Deployment Modes](#), Virtual Services can be configured in one of four fundamental ways, i.e. *Layer 4 DR mode*, *Layer 4 NAT mode*, *Layer 4 SNAT mode* or *Layer 7 SNAT mode*. The following sections illustrate how to configure the Virtual Services using various modes. If a different mode is required for a particular VIP, please refer to one of the other sections that uses that mode for guidance. Please also don't hesitate to contact our support team: support@loadbalancer.org.

8.2. 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 port connect on the service port.

8.3. Load Balancing DICOM

(Using Layer 4 DR Mode)

8.3.1. Setting up the Virtual Service (VIP)

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

Label	<input type="text" value="DICOM-Modalities"/>		
Virtual Service	IP Address	<input type="text" value="10.12.1.100"/>	
	Ports	<input type="text" value="104,11112"/>	
Protocol	<input type="text" value="TCP"/>		
Forwarding Method	<input type="text" value="Direct Routing"/>		

- Enter an appropriate name (Label) for the Virtual Service, e.g. **DICOM-Modalities**.
- Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.100**.

5. Set the *Virtual Service Ports* field to the required port(s), e.g. **104,11112**.
6. Set *Protocol* to **TCP**.
7. Set *Forwarding Method* to **Direct Routing**.
8. Click **Update**.
9. Now click **Modify** next to the newly created Virtual Service.
10. Set *Persistent Timeout* to **3600**, i.e. 1 hour.
11. Set the *Check Type* to **Connect to port** (the default).
12. Set the *Check Port* to the required port - by default this is set to the first port (**104**) of a multi-port VIP.
13. Click **Update**.

8.3.2. Setting up the Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created Virtual Service.
2. Enter the following details:

Label	<input type="text" value="DICOM1"/>	
Real Server IP Address	<input type="text" value="10.12.1.110"/>	
Weight	<input type="text" value="100"/>	
Minimum Connections	<input type="text" value="0"/>	
Maximum Connections	<input type="text" value="0"/>	

3. Enter an appropriate name (Label) for the first DICOM server, e.g. **DICOM1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.110**.
5. Click **Update**.
6. Now repeat for your remaining DICOM server(s).

8.3.3. Configuring the load balanced DICOM servers

1. As mentioned in [Load Balancing Deployment Modes](#), the ARP problem must be solved for all load balanced servers. Please refer to [Solving the ARP Problem](#).

8.4. Load Balancing HL7

(Using Layer 7 SNAT Mode)

8.4.1. Setting up the Virtual Service (VIP)



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

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="HL7"/>	?
IP Address	<input type="text" value="10.12.1.120"/>	?
Ports	<input type="text" value="2575"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

- Enter an appropriate name (Label) for the Virtual Service, e.g. **HL7**.
- Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.120**.
- Set the *Virtual Service Ports* field to the required port, e.g. **2575**.
- Set the *Layer 7 Protocol* to **TCP Mode**.
- Click **Update**.
- Now click **Modify** next to the newly created Virtual Service.
- In the *Persistence* section, click **Advanced** to show more options.
- Ensure *Persistence Mode* is set to **Source IP**.
- Set *Persistence Timeout* to **1h** (i.e. 1 hour).
- Set *Check Port* to the required port – leave blank to check the VIP port (**2575**).
- Click **Update**.

8.4.2. Setting up the Real Servers (RIPs)

- 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.
- Enter the following details:

Layer 7 Add a new Real Server

Label	<input type="text" value="HL71"/>	?
Real Server IP Address	<input type="text" value="10.12.1.130"/>	?
Real Server Port	<input type="text" value="2575"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Enter an appropriate name (Label) for the first HL7 server, e.g. **HL71**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.130**.
5. Set the *Real Server Port* field to **2575**.
6. Click **Update**.
7. Now repeat for your remaining HL7 server(s).

8.5. Load Balancing XDS (Registry & Repository)

(Using Layer 4 DR Mode)

8.5.1. 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="XDS-Registry"/>	?	
Virtual Service	IP Address	<input type="text" value="10.12.1.140"/>	?
	Ports	<input type="text" value="17035"/>	?
Protocol	<input type="text" value="TCP"/>	?	
Forwarding Method	<input type="text" value="Direct Routing"/>	?	

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **XDS-Registry**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.40**.
5. Set the *Virtual Service Ports* field to the required port, e.g. **17035**.
6. Set *Protocol* to **TCP**.

7. Set *Forwarding Method* to **Direct Routing**.
8. Click **Update**.
9. Now click **Modify** next to the newly created Virtual Service.
10. Set *Persistent Timeout* to **3600** , i.e. 1 hour.
11. Set *Check Port* to the required port – leave blank to check the VIP port (**17035**).
12. Click **Update**.

8.5.2. Setting up the Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created Virtual Service.
2. Enter the following details:

Label	<input type="text" value="XDS1"/>	
Real Server IP Address	<input type="text" value="10.12.1.150"/>	
Weight	<input type="text" value="100"/>	
Minimum Connections	<input type="text" value="0"/>	
Maximum Connections	<input type="text" value="0"/>	

3. Enter an appropriate name (Label) for the first XDS server, e.g. **XDS1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.150**.
5. Click **Update**.
6. Now repeat for your remaining XDS server(s).

8.5.3. Configuring the load balanced XDS servers

1. As mentioned in [Load Balancing Deployment Modes](#), the ARP problem must be solved for all load balanced servers. Please refer to [Solving the ARP Problem](#) for more details.

8.6. Load Balancing HTTPS

(Using Layer 7 SNAT Mode)

8.6.1. 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="HTTPS-Viewer"/>	?
IP Address	<input type="text" value="10.12.1.160"/>	?
Ports	<input type="text" value="443"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **HTTPS-Viewer**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.160**.
5. Set the *Virtual Service Ports* field to the required port, e.g. **443**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created Virtual Service.
9. In the *Persistence* section, click **Advanced** to show more options.
10. Ensure *Persistence Mode* is set to **Source IP**.
11. Set *Persistence Timeout* to **1h** (i.e. 1 hour).
12. Set *Check Port* to the required port – leave blank to check the VIP port (**443**).
13. Click **Update**.

8.6.2. 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

Label	<input type="text" value="VIEWER1"/>	?
Real Server IP Address	<input type="text" value="10.12.1.170"/>	?
Real Server Port	<input type="text" value="443"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Enter an appropriate name (Label) for the first Viewer server, e.g. **VIEWER1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.170**.
5. Set the *Real Server Port* field to **443**.
6. Click **Update**.
7. Now repeat for your remaining Viewer server(s).

8.7. Finalizing Appliance Settings

8.7.1. Configure HAProxy Timeouts

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Advanced Configuration*.

Connection Timeout	<input type="text" value="4000"/>	ms	?
Client Timeout	<input type="text" value="1h"/>	ms	?
Real Server Timeout	<input type="text" value="1h"/>	ms	?

2. Change *Client Timeout* to **1h** as shown above (i.e. 1 hour).
3. Change *Real Server Timeout* to **1h** as shown above (i.e. 1 hour).
4. Click the **Update** button to save the settings.

8.7.2. Restart HAProxy

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the "Commit changes" box at the top of the screen or by using the *Restart Services* menu option:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.

Note

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

9. Testing & Verification

Note

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

9.1. Using the System Overview

Verify that all VIPs & associated RIPS are reported as up (green) as shown below:



VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
DICOM-Modalities..	10.12.1.100	104	0	TCP	Layer 4	DR	
<i>REAL SERVER</i>							
<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>				
DICOM1	10.12.1.110	104	100	0	Drain	Halt	
DICOM2	10.12.1.111	104	100	0	Drain	Halt	
HL7	10.12.1.120	2575	0	TCP	Layer 7	Proxy	
<i>REAL SERVER</i>							
<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>				
HL71	10.12.1.130	2575	100	0	Drain	Halt	
HL72	10.12.1.131	2575	100	0	Drain	Halt	
HTTPS-Viewer	10.12.1.160	443	0	TCP	Layer 7	Proxy	
<i>REAL SERVER</i>							
<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>				
VIEWER1	10.12.1.170	443	100	0	Drain	Halt	
VIEWER2	10.12.1.171	443	100	0	Drain	Halt	

If certain servers are down, i.e. failing their health check, they will be highlighted red as shown below:

HTTPS-Viewer	10.12.1.160	443	0	TCP	Layer 7	Proxy	
<i>REAL SERVER</i>							
<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>				
VIEWER1	10.12.1.170	443	100	0	Drain	Halt	
VIEWER2	10.12.1.171	443	100	0	Drain	Halt	

9.2. 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*.

10. 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.

11. Further Documentation

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

12. Appendix

12.1. 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.

12.1.1. 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.

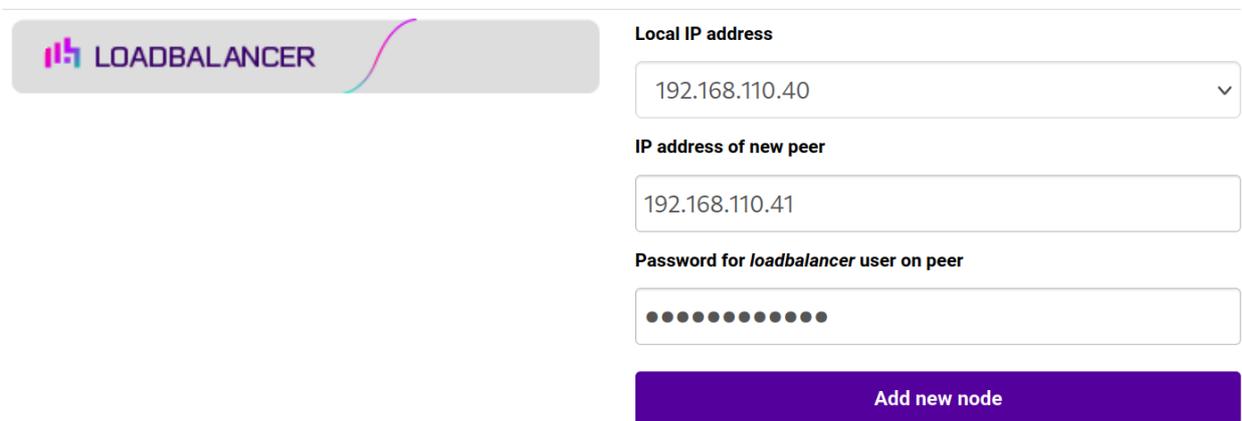
12.1.2. Configuring the HA Clustered Pair

📌 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.

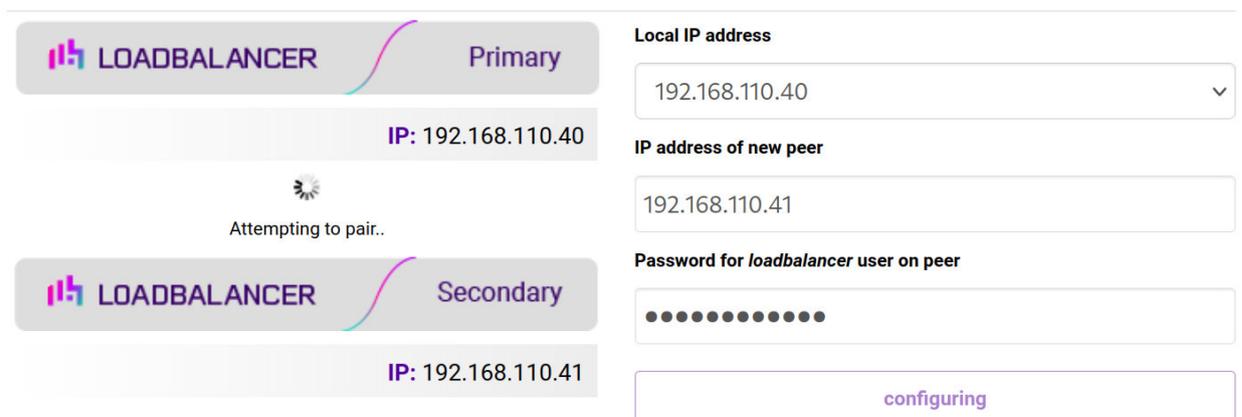
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: *Cluster Configuration > High-Availability Configuration*.

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: a Primary one with IP 192.168.110.40 and a Secondary one with IP 192.168.110.41. A prominent red button labeled 'Break Clustered Pair' is located to the right of the Primary appliance.

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](#).

12.2. 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](#).

13. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.2.0	14 August 2019	Styling and layout	General styling updates	RJC
1.2.1	24 August 2020	New title page Updated Canadian contact details Amended instructions for setting persistence timeouts	Branding update Change to Canadian contact details Changes to the appliance WebUI	AH
1.3.0	1 December 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.3.1	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.3.2	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 Removed references to the colour of certain UI elements	Housekeeping across all documentation	AH
1.3.3	2 February 2023	Updated screenshots	Branding update	AH
1.3.4	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.4.0	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.

