DEPLOYMENT GUIDE

Load Balancing Medical Imaging & Information System Protocols

Version 1.3.0

loadbalancer.org

Table of Contents

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. 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. Medical Systems Supported

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

5. Medical Imaging and Information Systems & Components

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.

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.

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

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

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.

6. Medical Information System Standards & Protocols

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.

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.

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.

7. Load Balancing Overview

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

Load Balancer Deployment

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

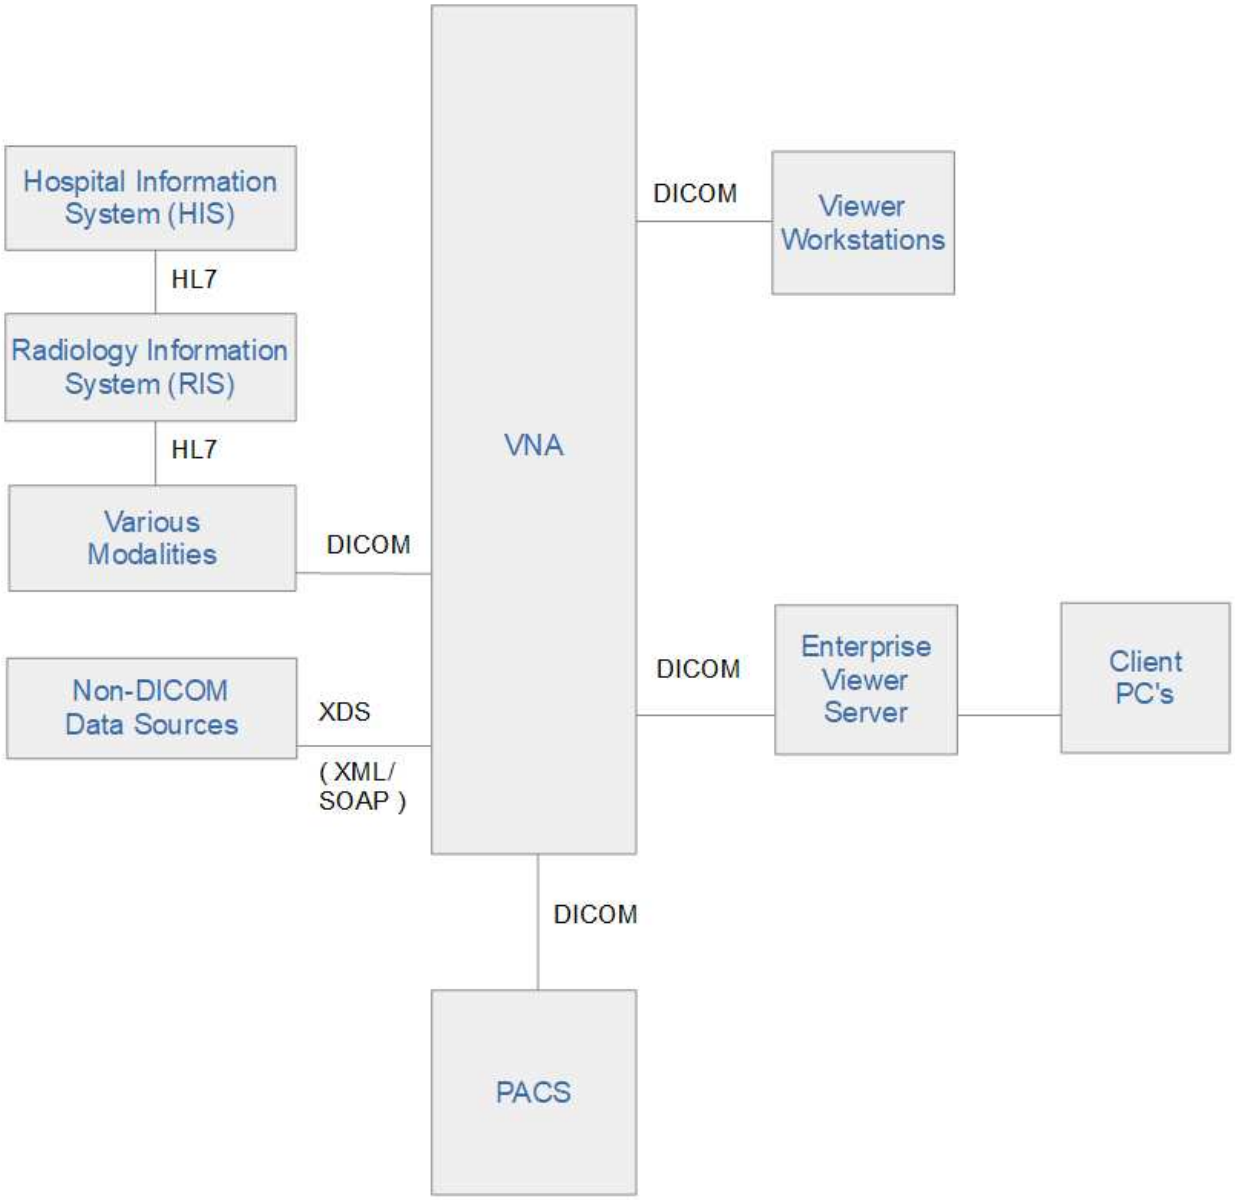

The diagram below shows a highly available system that utilizes multiple system components and load balancing:

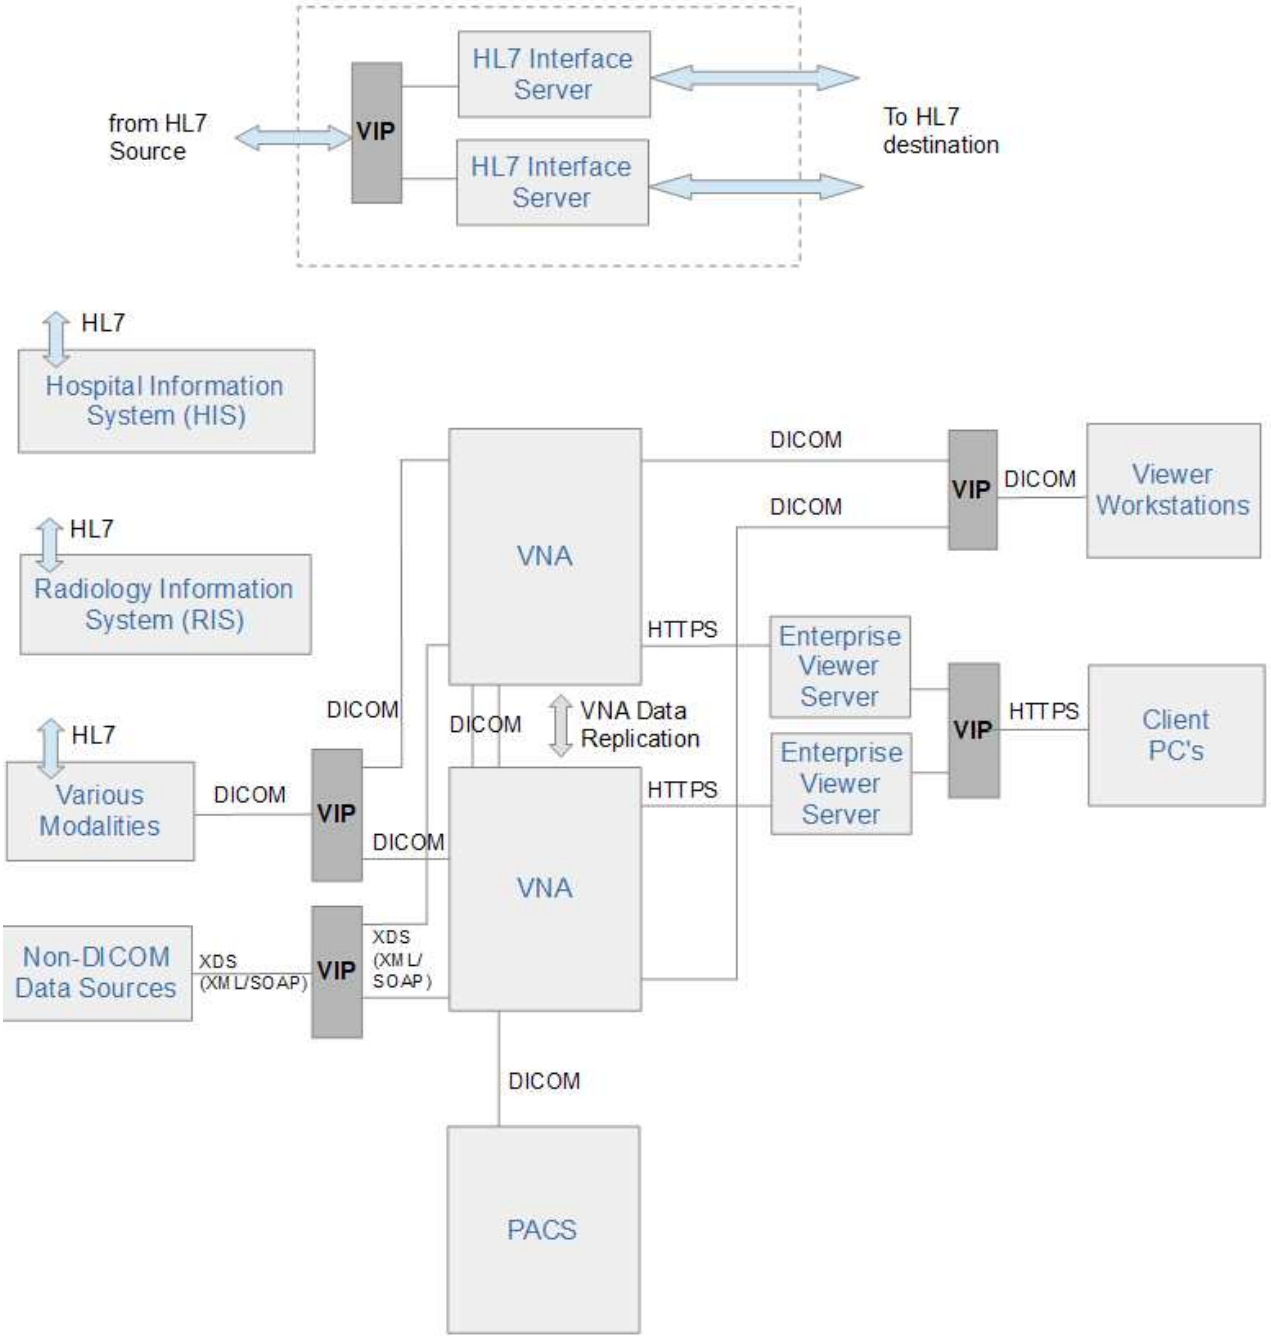

Notes

- **VIP** (**V**irtual **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 that one load balancer appliance may be required depending on throughput and physical network topology.

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

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.

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.

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.

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.

8. 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 The Virtual Appliance - Hypervisor Deployment and the ReadMe.txt text file included in the VA download for more detailed information on deploying the VA using various Hypervisors.
Note	For the VA, 4 NICs are included but only eth0 is connected by default at power up. If the other NICs are required, these should be connected using the network configuration screen within the Hypervisor.

Initial Network Configuration

After boot up, follow the instructions on the console to configure the IP address, subnet mask, default gateway, DNS 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:

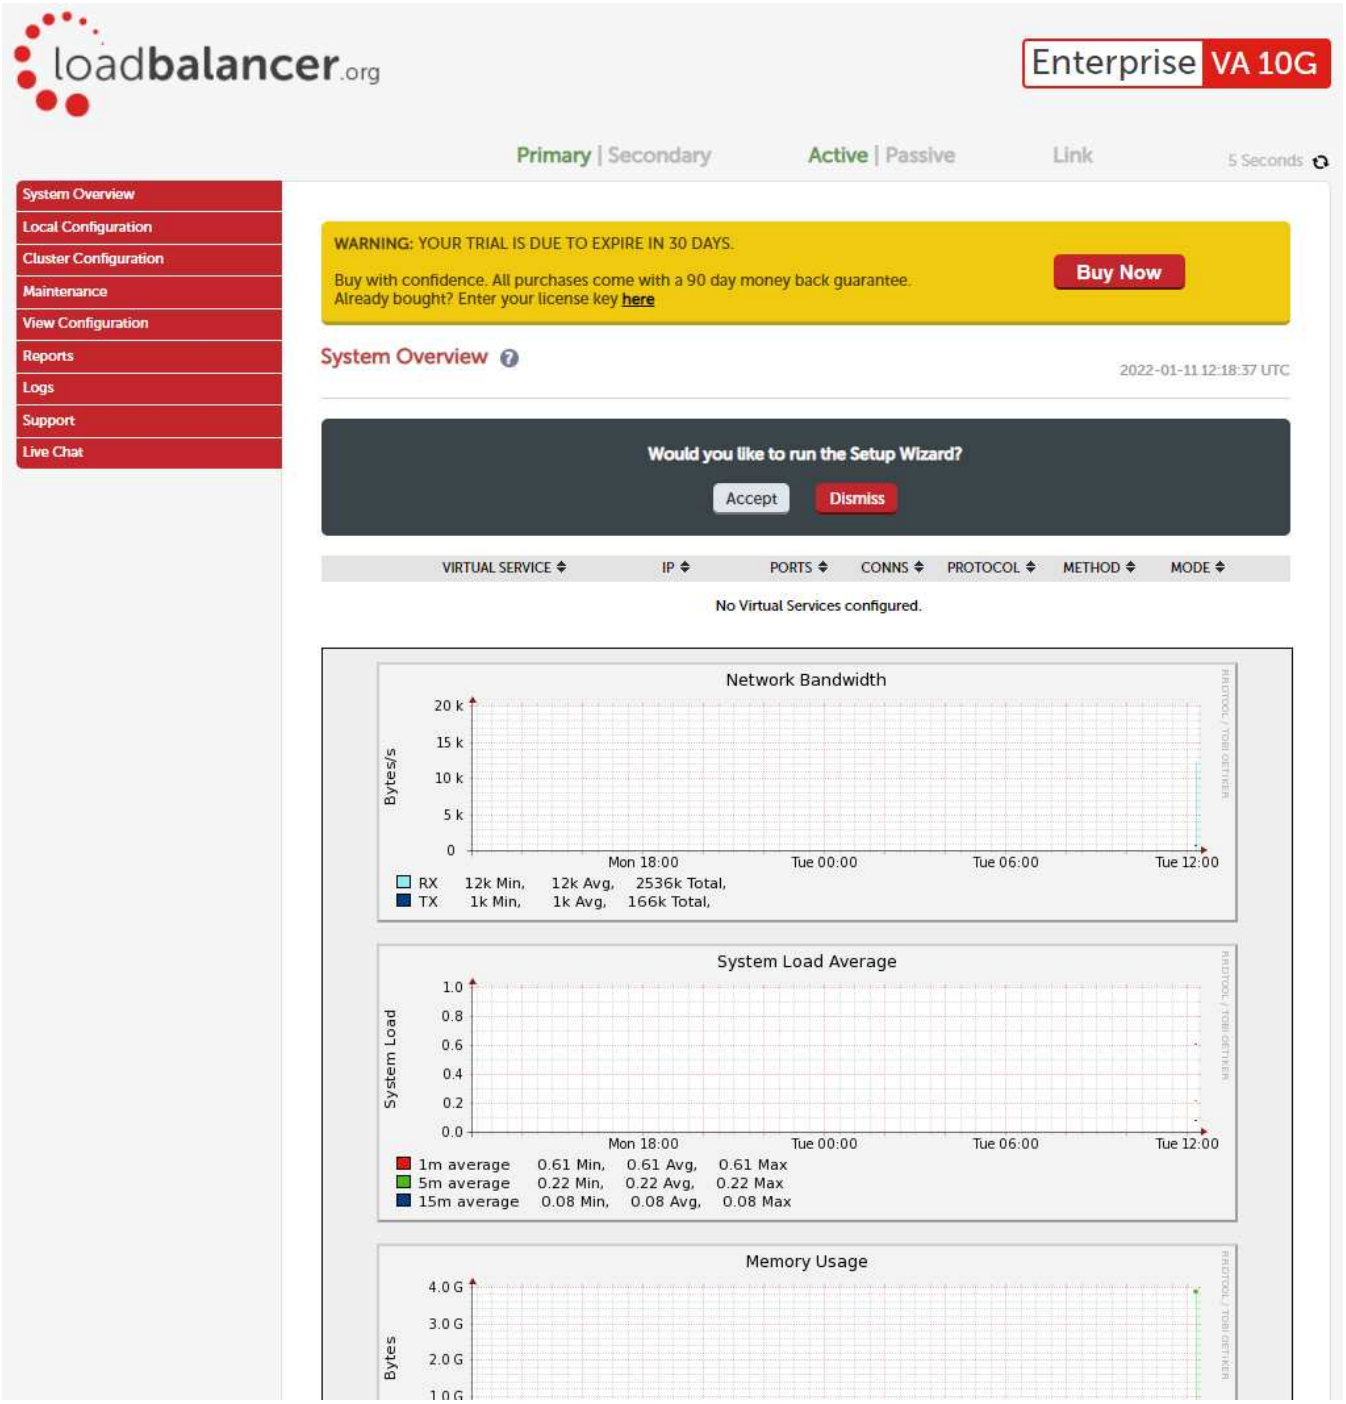

Note	The WebUI for the VA is shown, the hardware and cloud appliances are very similar. The yellow licensing related message is platform & model dependent.

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, simple 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

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.

9. Appliance & Server Configuration

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.

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.

Load Balancing DICOM

(Using Layer 4 DR Mode)

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:

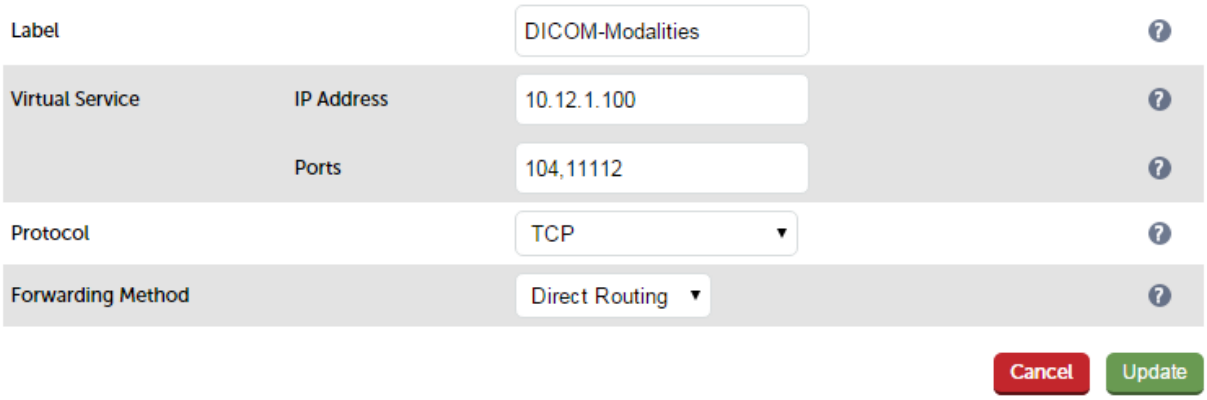

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **DICOM-Modalities**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.100**.
5. Set the *Virtual Service Ports* field to the required port(s), e.g. **104,11112**.
6. Set *Protocol* to **TCP**.

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**.

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:

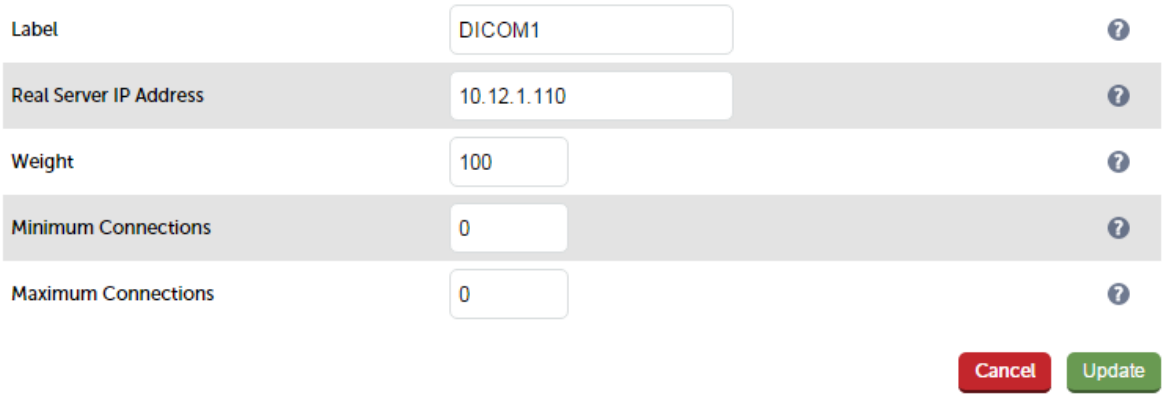

3. Enter an appropriate name (Label) for the first DICOM server, e.g. **DICOM1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.110**.
5. Click **Update**.
6. Now repeat for your remaining DICOM server(s).

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.

Load Balancing HL7

(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:

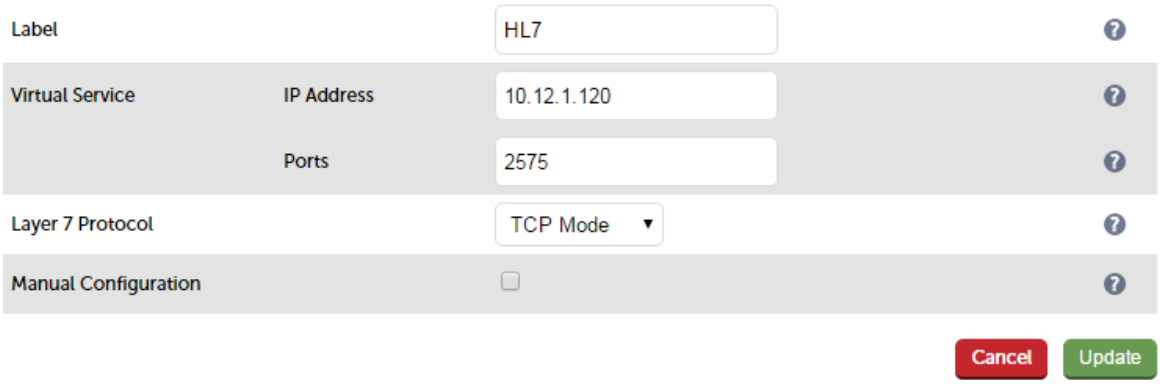

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **HL7**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.120**.
5. Set the *Virtual Service Ports* field to the required port, e.g. **2575**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **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 (**2575**).
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:

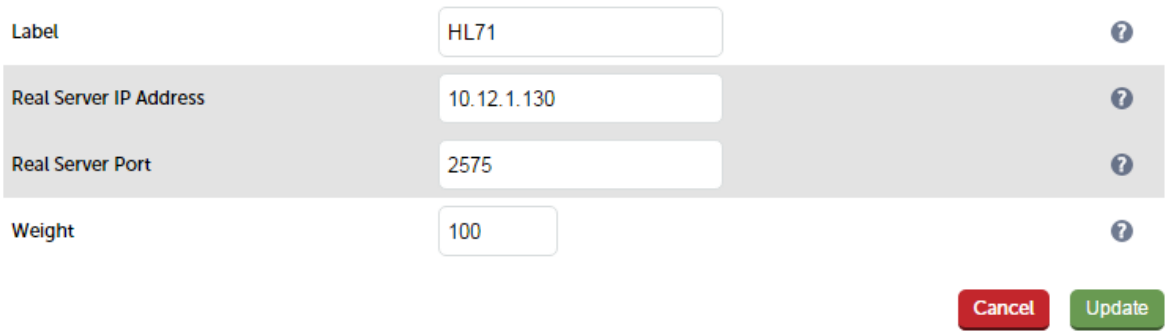

3. Enter an appropriate name (Label) for the first HL7 server, e.g. **HL71**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.130**.
5. Set the *Real Server Port* field to **2575**.
6. Click **Update**.

7. Now repeat for your remaining HL7 server(s).

Load Balancing XDS (Registry & Repository)

(Using Layer 4 DR Mode)

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:

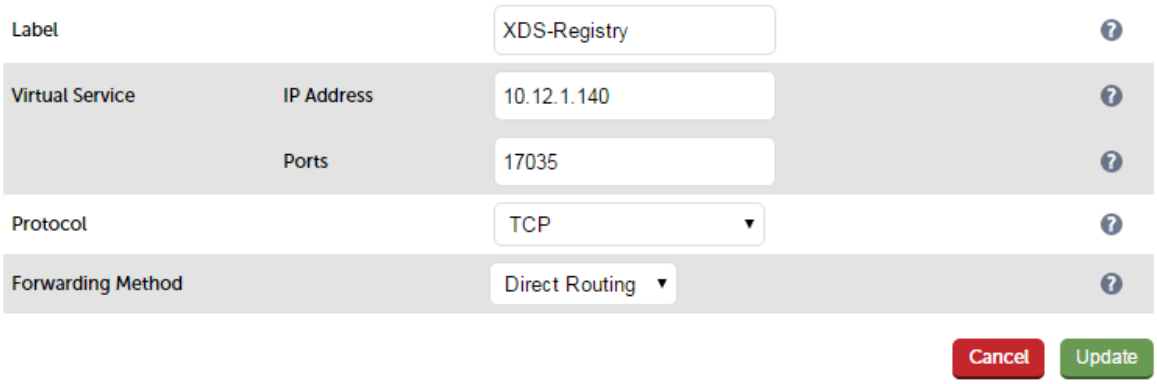

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **XDS-Registry**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.40**.
5. Set the *Virtual Service Ports* field to the required port, e.g. **17035**.
6. Set *Protocol* to **TCP**.
7. Set *Forwarding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

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:

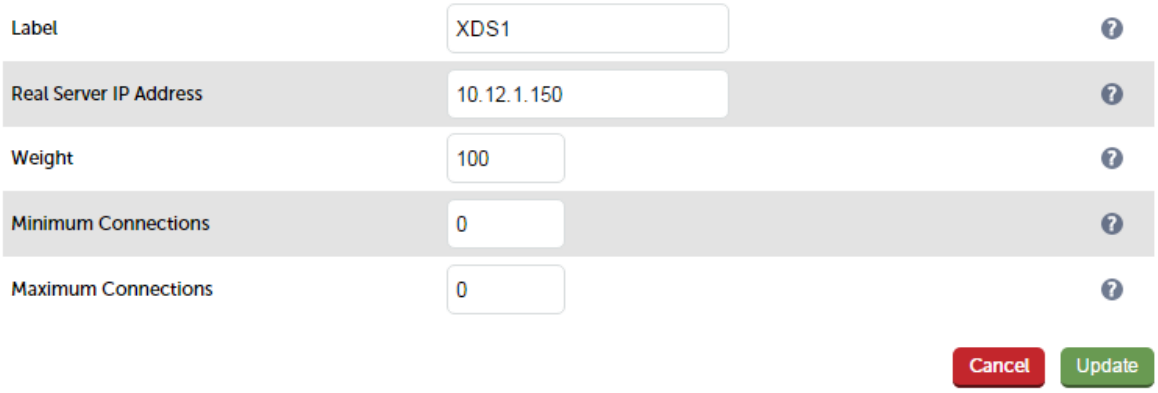

3. Enter an appropriate name (Label) for the first XDS server, e.g. **XDS1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.150**.
5. Click **Update**.
6. Now repeat for your remaining XDS server(s).

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.

Load Balancing HTTPS

(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:

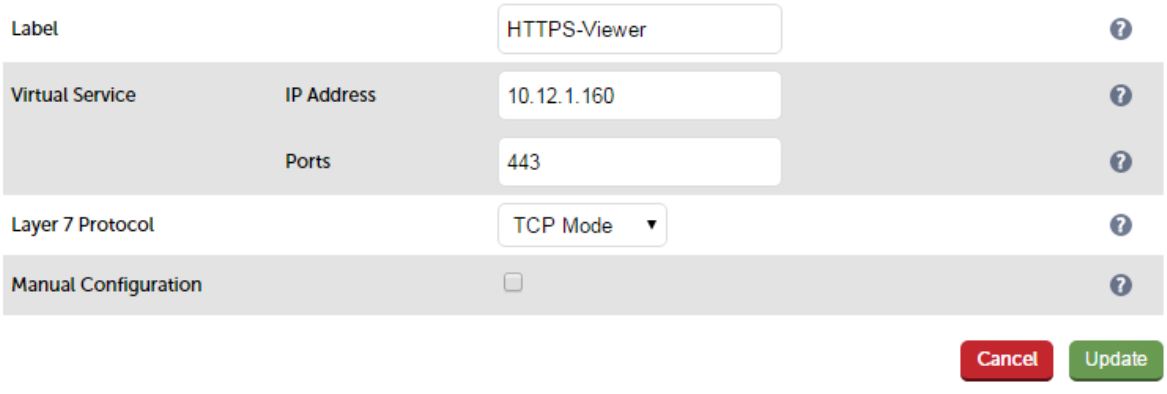

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **HTTPS-Viewer**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.12.1.160**.
5. Set the *Virtual Service Ports* field to the required port, e.g. **443**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
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**.

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:

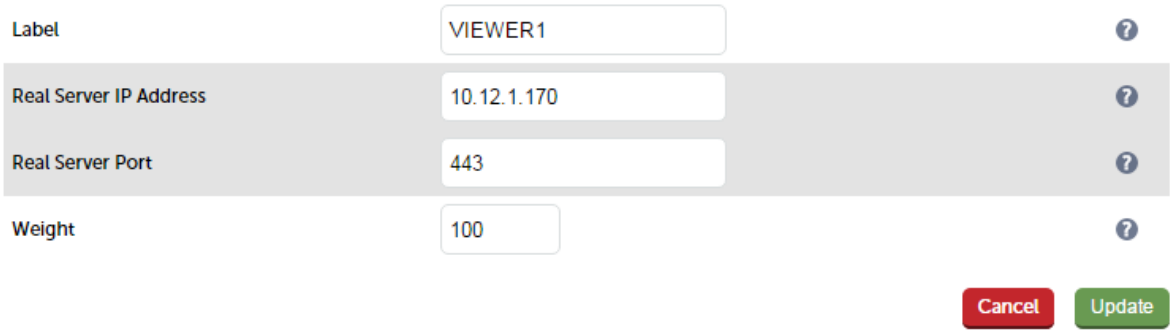

3. Enter an appropriate name (Label) for the first Viewer server, e.g. **VIEWER1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.12.1.170**.
5. Set the *Real Server Port* field to **443**.
6. Click **Update**.
7. Now repeat for your remaining Viewer server(s).

Finalizing Appliance Settings

Configure HAProxy Timeouts

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

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.

Restart HAProxy

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the blue 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.

10. Testing & Verification

Note	For additional general guidance please also refer to Testing Load Balanced Services.

Using the System Overview

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

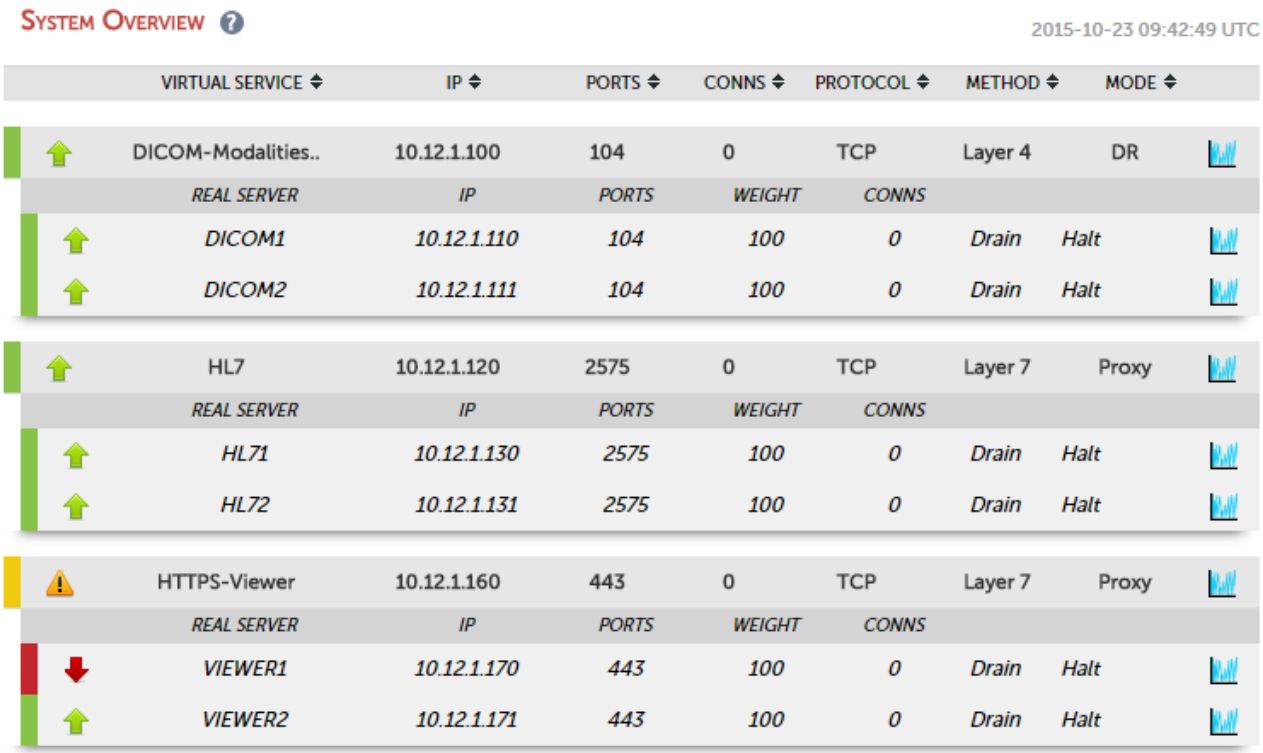

If certain servers are down, i.e. failing their health check, they will be highlighted red as shown below:

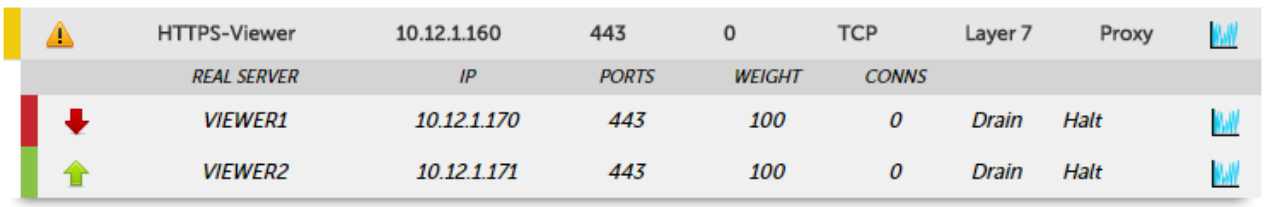

System Logs & Reports

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*.

11. 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.

12. Further Documentation

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

13. Conclusion

Loadbalancer.org appliances provide a very cost effective and flexible solution for highly available load balanced Medical Imaging Systems environments.

14. 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.

Note

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

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

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*.

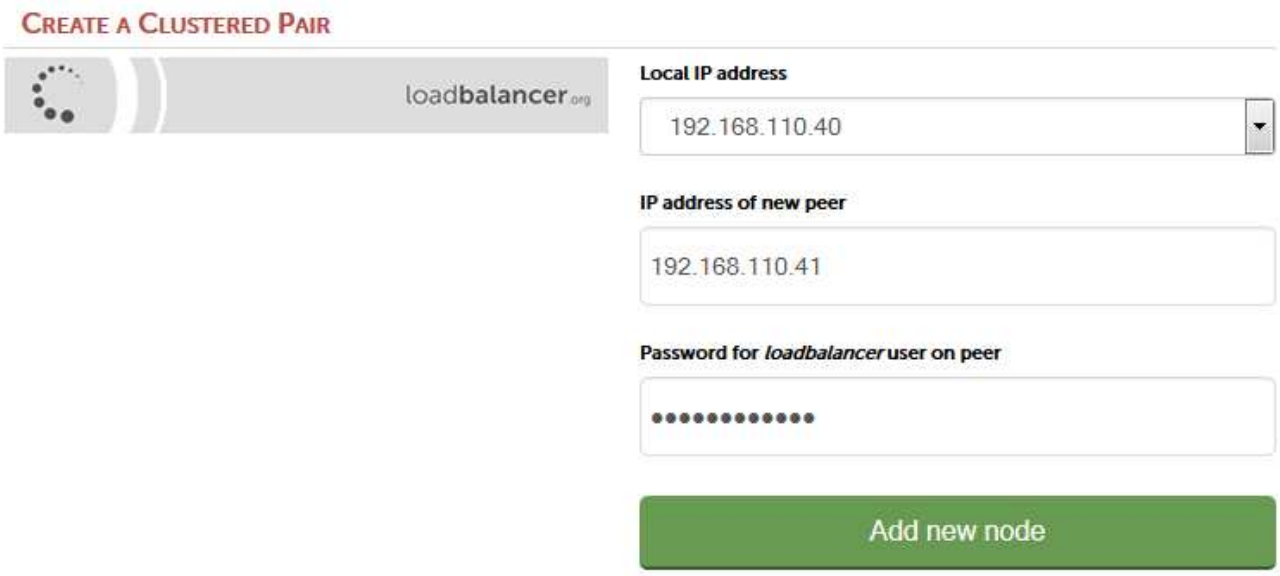

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

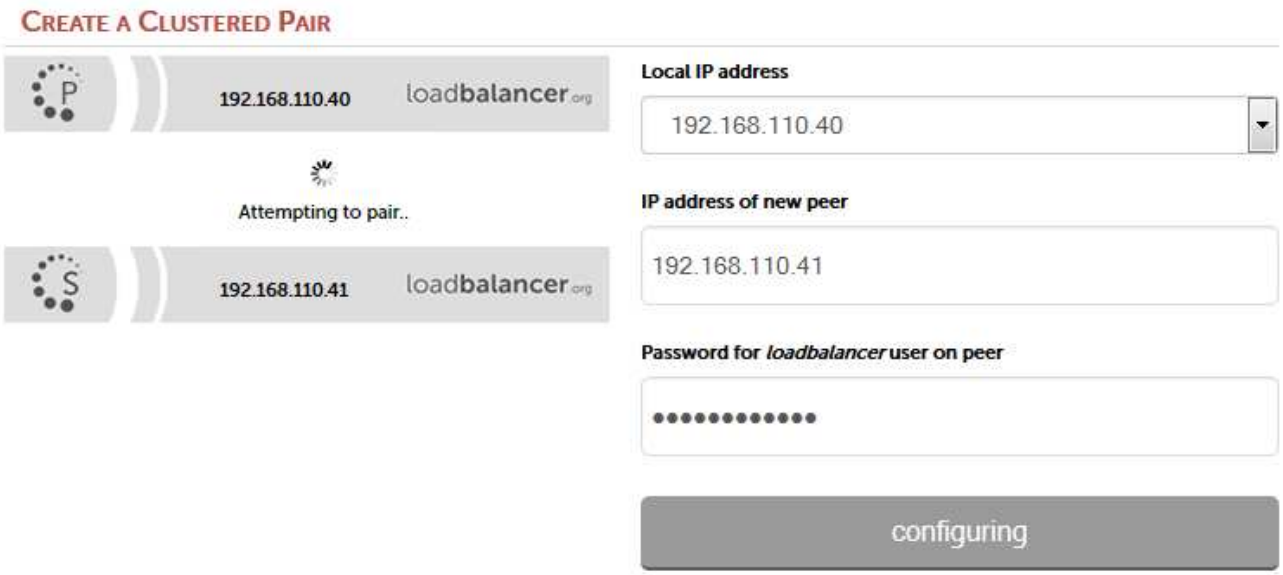

6. Once complete, the following will be displayed on the Primary appliance:

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.

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.

15. 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

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