



# Load Balancing Medical Imaging & Information System Protocols

Deployment Guide  
**v1.1.2**

---

## Table of Contents

1. About this Guide.....	3
2. Loadbalancer.org Appliances Supported.....	3
3. Loadbalancer.org Software Versions Supported.....	3
4. Medical Systems Supported.....	3
5. Medical Imaging and Information Systems & Components.....	4
Picture Archiving and Communication System (PACS).....	4
Vendor Neutral Archive (VNA).....	4
Imaging Modalities.....	4
Health Care Administration Systems.....	4
Workstations/Viewers.....	4
6. Medical Information System Standards & Protocols.....	5
DICOM.....	5
HL7.....	5
IHE XDS.....	5
7. Load Balancing Overview.....	5
Basic Concepts.....	5
Load Balancer Deployment.....	6
Supported Load Balancing Methods.....	8
Our Recommendation.....	8
Load Balanced Ports & Services.....	8
Persistence (Server Affinity).....	9
Server Health Checking.....	9
8. Loadbalancer.org Appliance – the Basics.....	9
Virtual Appliance Download & Deployment.....	9
Initial Network Configuration.....	10
Accessing the Web User Interface (WebUI).....	10
Clustered Pair Configuration.....	11
9. Appliance & Server Configuration.....	12
Load Balancing DICOM.....	12
Load Balancing HL7.....	13
Load Balancing XDS (Registry & Repository).....	15
Load Balancing HTTPS.....	16
10. Testing & Verification.....	18
Using the System Overview.....	18
System Logs & Reports.....	19
11. Technical Support.....	19
12. Further Documentation.....	19
13. Conclusion.....	19
14. Appendix.....	20
1 - Clustered Pair Configuration – Adding a Slave Unit.....	20
2 – Solving the ARP Problem.....	22
3 - Company Contact Information.....	23

## 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 relevant Administration Manual:

- [v7 Administration Manual](#)
- [v8 Administration Manual](#)

## 2. Loadbalancer.org Appliances Supported

All our products can be used with Medical Imaging and Information Systems. The complete list of models is shown below:

Discontinued Models	Current Models *
Enterprise R16	Enterprise R20
Enterprise VA R16	Enterprise MAX
Enterprise VA	Enterprise 10G
	Enterprise Ultra
	Enterprise VA R20
	Enterprise VA MAX

\* For full specifications of these models please refer to: <http://www.loadbalancer.org/products/hardware>

## 3. Loadbalancer.org Software Versions Supported

- V7.6.4 & later

## 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." VNA's are also sometimes referred to as a PACS Neutral Archive.

VNA's 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 PC's 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 a 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

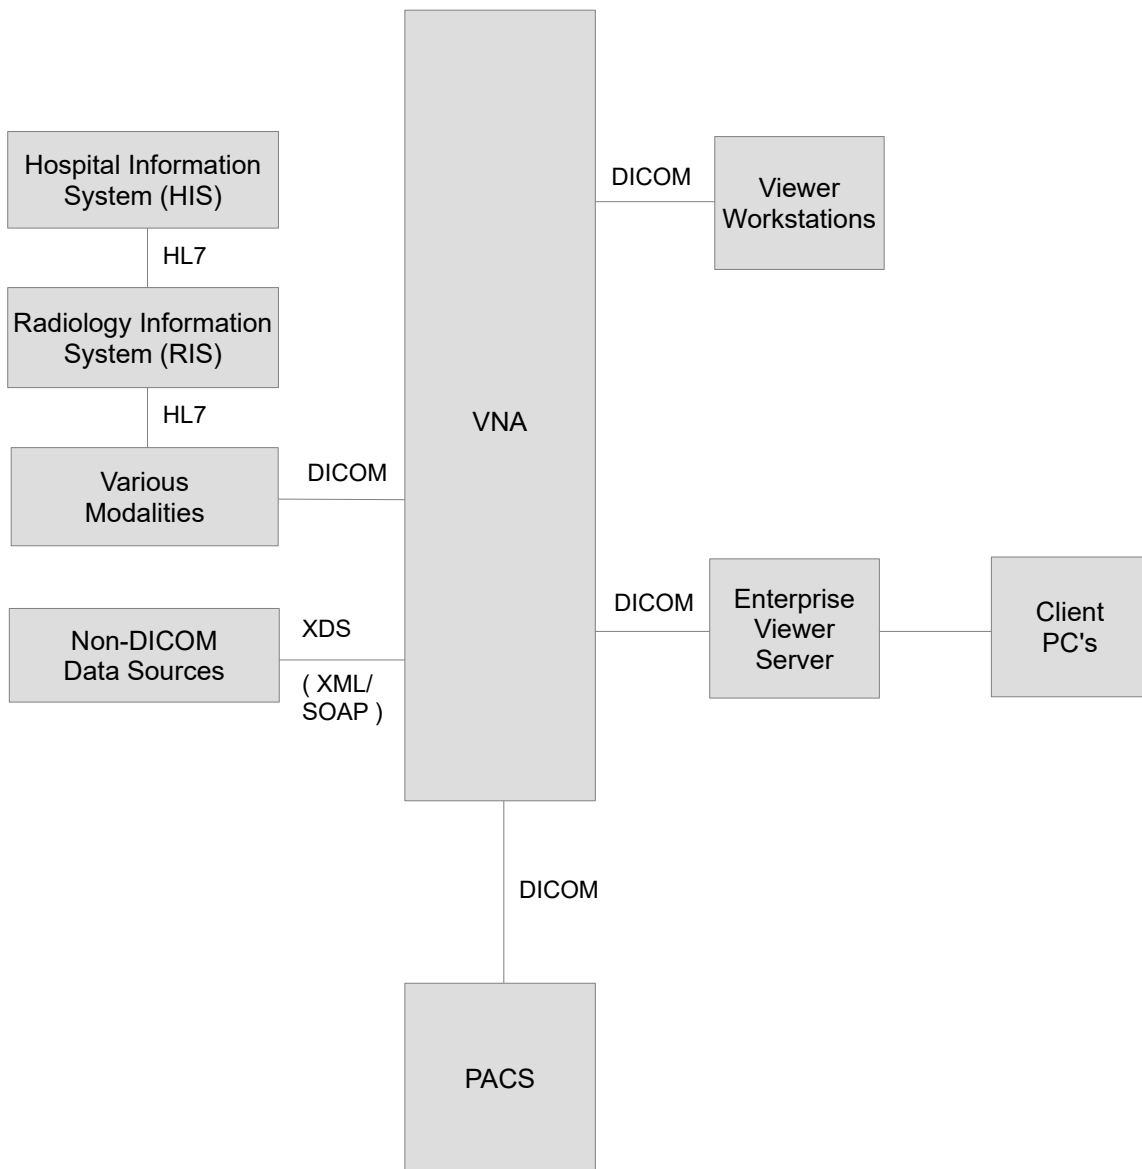
*Load is distributed across all configured servers/endpoints*

- **Failover mode**

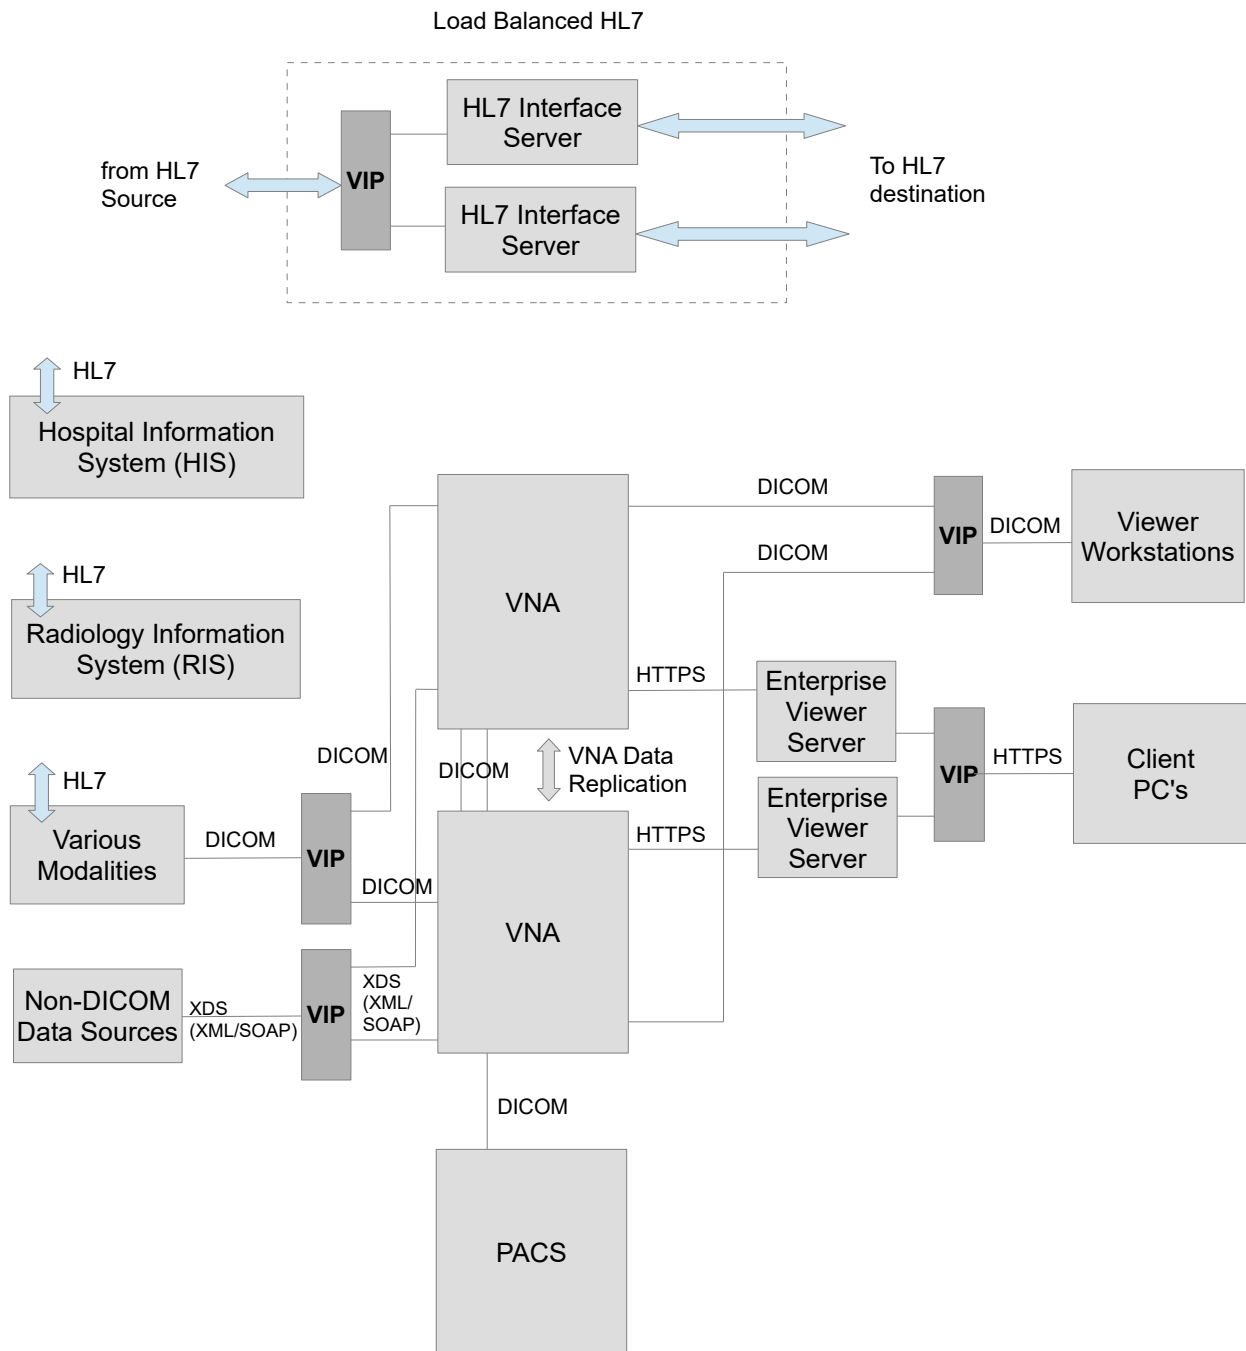
*The second server is used only when the first server/endpoint fails*

## 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 (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
- 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
- All Loadbalancer.org models support unlimited VIPs except the Enterprise R20 which supports up to 5 VIPs, each with up to 4 load balanced servers

## SUPPORTED LOAD BALANCING METHODS

**Layer 4 DR Mode** offers the best performance and requires limited physical Real Server changes. The load balanced application must be able to bind to the Real Servers own IP address and the VIP at the same time. This mode requires the "ARP Problem" to be solved as described on page [22](#).

**Layer 4 NAT Mode** is also a high performance solution but not as fast as DR mode. It requires the implementation of a two-arm infrastructure with an internal and external subnet to carry out the translation (the same way a firewall works). Also each Real Server must use the load balancer as the default gateway.

**Layer 7 SNAT 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 does not require any changes to the Real Servers and can be deployed in one-arm or two-arm mode and. HAProxy is a high performance solution, but since it operates as a full proxy, it cannot perform as fast as the layer 4 solutions.

## 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 Chapter 8 in the [Administration Manual](#).

### 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 DOWNLOAD & DEPLOYMENT

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 and XEN and has been optimized for each Hypervisor. By default, the VA is allocated 1 CPU, 2GB of RAM and has an 8GB 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 [Administration Manual](#) and the ReadMe.txt text file included in the VA download for more detailed information on deploying the VA using various Hypervisors.

## INITIAL NETWORK CONFIGURATION

The IP address, subnet mask, default gateway and DNS settings can be configured in several ways as detailed below:

### *Method 1 - Using the Network Setup Wizard at the console*

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

### *Method 2 - Using the WebUI*

Using a browser, connect to the WebUI on the default IP address/port: **http://192.168.2.21:9080**

To set the IP address & subnet mask, use: *Local Configuration > Network Interface Configuration*

To set the default gateway, use: *Local Configuration > Routing*

To configure DNS settings, use: *Local Configuration > Hostname & DNS*

### *Method 3 - Using Linux commands*

At the console, set the initial IP address using the following command:

```
ip addr add <IP address>/<mask> dev eth0
```

At the console, set the initial default gateway using the following command:

```
route add default gw <IP address> <interface>
```

At the console, set the DNS server using the following command:

```
echo nameserver <IP address> >> /etc/resolv.conf
```

#### Note:

If method 3 is used, you must also configure these settings using the WebUI, otherwise the settings will be lost after a reboot

## ACCESSING THE WEB USER INTERFACE (WEBUI)

The WebUI can be accessed via HTTP at the following URL: **http://192.168.2.21:9080/lbadmin**

\* *Note the port number → 9080*

The WebUI can be accessed via HTTPS at the following URL: **https://192.168.2.21:9443/lbadmin**

\* *Note the port number → 9443*

*(replace 192.168.2.21 with the IP address of your load balancer if it's been changed from the default)*

Login using the following credentials:

**Username:** loadbalancer

**Password:** loadbalancer

#### Note:

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

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

The screenshot displays the Loadbalancer.org Enterprise VA MAX web interface. The top navigation bar shows 'Master | Slave', 'Active | Passive', and 'Link' status, along with a refresh button and '5 Seconds' interval. The left sidebar contains a menu with items like 'System Overview', 'Local Configuration', 'Cluster Configuration', 'Maintenance', 'View Configuration', 'Reports', 'Logs', and 'Support'. The main content area is titled 'SYSTEM OVERVIEW' and includes a 'Would you like to run the Setup Wizard?' dialog with 'Accept' and 'Dismiss' buttons. Below the dialog, there are filters for 'VIRTUAL SERVICE', 'IP', 'PORTS', 'CONNS', 'PROTOCOL', 'METHOD', and 'MODE', with a message stating 'No Virtual Services configured.' Three charts are displayed: 'Network Bandwidth' showing RX and TX traffic, 'System Load Average' showing 1m, 5m, and 15m averages, and 'Memory Usage' showing Used, Page, Buffer, and Free memory.

(shows v8.2.x)

## 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 slave unit is covered in section 1 of the Appendix on page [20](#).

## 9. Appliance & Server Configuration

### Load Balancing Mode

As mentioned on page 8, Virtual Services can be configured in one of three fundamental ways, i.e. *Layer 4 DR mode*, *Layer 4 NAT mode* or *Layer 7 SNAT mode*. The flow chart on page 8 can be used to help choose which is most appropriate. Each section below shows how to configure a Virtual Service using a specific mode. If a different mode is required, please refer to one of the other sections that uses that mode for guidance.

### Health-Check Configuration

As mentioned on page 9, 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)

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

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

- 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**
- Set the *Virtual Service Ports* field to the required port(s), e.g. **104,11112**
- Set *Protocol* to **TCP**
- Set *Forwarding Method* to **Direct Routing**
- Click **Update**
- Now click **Modify** next to the newly created Virtual Service
- Set *Persistent Timeout* to **3600**, i.e. 1 hour
- Set the *Check Type* to **Connect to port** (the default)
- Set the *Check Port* to the required port - by default this is set to the first port (**104**)
- Click **Update**

## Setting up the Real Servers (RIPs)

- 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
- 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"/>	
<input type="button" value="Cancel"/> <input type="button" value="Update"/>		

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

## Configuring the load balanced DICOM servers

- As mentioned on page [8](#), the ARP problem must be solved for all load balanced servers. Please refer to section 2 in the appendix on page [22](#) for more details.

## LOAD BALANCING HL7

(Using Layer 7 SNAT Mode)

### 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:





Label	<input type="text" value="HL7"/>		
Virtual Service	IP Address	<input type="text" value="10.12.1.120"/>	
	Ports	<input type="text" value="2575"/>	
Layer 7 Protocol	<input type="text" value="TCP Mode"/>		
Manual Configuration	<input type="checkbox"/>		
<input type="button" value="Cancel"/> <input type="button" value="Update"/>			

- 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. Ensure *Persistence Mode* is set to **Source IP**
10. Set *Persistence Timeout* to **1h** (i.e. 1 hour)
11. Set *Check Port* to the required port – leave blank to check the service port (**2575**)
12. 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:

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"/>	
Weight	<input type="text" value="100"/>	
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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)

### Configure HAProxy Timeout Settings

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

## 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 once all layer 7 Virtual Services and Real Servers have been defined.

## LOAD BALANCING XDS (REGISTRY & REPOSITORY)

(Using Layer 4 NAT 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:

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="NAT"/>		
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>	

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. **10.12.1.140**
5. Set the *Virtual Service Ports* field to the required port, e.g. **17035**
6. Set *Protocol* to **TCP**
7. Set *Forwarding Method* to **NAT**
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 service 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:

Label	<input type="text" value="XDS1"/>	?
Real Server IP Address	<input type="text" value="10.12.1.150"/>	?
Real Server Port	<input type="text" value="17035"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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. Set the *Real Server Port* field to the required port, e.g. **17035**
6. Click **Update**
7. Now repeat for your remaining XDS server(s)

### Configuring the load balanced XDS servers

1. As mentioned on page [8](#), the ARP problem must be solved for all load balanced servers. Please refer to page [22](#) 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:

Label	<input type="text" value="HTTPS-Viewer"/>	?	
Virtual Service	IP Address	<input type="text" value="10.12.1.160"/>	?
	Ports	<input type="text" value="443"/>	?
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	▼	?
Manual Configuration	<input type="checkbox"/>		?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>	

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. Ensure *Persistence Mode* is set to **Source IP**
10. Set *Persistence Timeout* to **1h** (i.e. 1 hour)
11. Set *Check Port* to the required port – leave blank to check the service port (**443**)
12. 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:

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"/>	?
Weight	<input type="text" value="100"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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

*If this has not already been configured:*

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

## Restart HAProxy

- 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 once all layer 7 Virtual Services and Real Servers have been defined.

# 10. Testing & Verification

## USING THE SYSTEM OVERVIEW

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

SYSTEM OVERVIEW <span>?</span>								2015-10-23 09:42:49 UTC	
	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		

## 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

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

## 12. Further Documentation

The Administration Manual contains much more information about configuring and deploying the appliance. It's available here: <http://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

### 1 - CLUSTERED PAIR CONFIGURATION – ADDING A SLAVE UNIT

If you initially configured just the master unit and now need to add a slave - our recommended procedure, please refer to the relevant section below for more details:

**Note:**

A number of settings are not replicated as part of the master/slave pairing process and therefore must be manually configured on the slave appliance. These are listed below:

- Hostname & DNS settings
- Network settings including IP addresses, bonding configuration and VLANs
- Routing configuration including default gateways and static routes
- Date & time settings
- Physical – Advanced Configuration settings including Internet Proxy IP address & port, Firewall table size, SMTP relay and Syslog server
- SNMP settings
- Graphing settings
- Firewall Script & Firewall Lockdown Script settings
- Software updates

**Version 7:**

Please refer to **Chapter 8 – Appliance Clustering for HA** in the [v7 Administration Manual](#).

**Version 8:**

*To add a slave node – i.e. create a highly available clustered pair:*

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

**CREATE A CLUSTERED PAIR**

loadbalancer.org

Local IP address  
192.168.1.20



IP address of new peer  
192.168.1.21

Password for *loadbalancer* user on peer  
.....

Add new node



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

**CREATE A CLUSTERED PAIR**

 192.168.1.20 loadbalancer.org	Attempting to pair..	<p><b>Local IP address</b></p> <input style="width: 90%;" type="text" value="192.168.1.20"/>
 192.168.1.21 loadbalancer.org		<p><b>IP address of new peer</b></p> <input style="width: 90%;" type="text" value="192.168.1.21"/>
		<p><b>Password for <i>loadbalancer</i> user on peer</b></p> <input style="width: 90%;" type="password" value="....."/>
<div style="background-color: #ccc; padding: 5px; display: inline-block; border: 1px solid #000;">configuring</div>		

6. Once complete, the following will be displayed:

**HIGH AVAILABILITY CONFIGURATION - MASTER**

 192.168.1.20 loadbalancer.org	 192.168.1.21 loadbalancer.org	<div style="background-color: #c00; color: white; padding: 10px; display: inline-block; border: 1px solid #000;">Break Clustered Pair</div>
---	---	---

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 master appliance will also automatically restart heartbeat on the slave appliance.

**Note:**

Please refer to chapter 9 – Appliance Clustering for HA in the [Administration Manual](#) for more detailed information on configuring HA with 2 appliances.

## 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 the [Administration Manual](#) and search for "*DR Mode Considerations*".

### 3 - COMPANY CONTACT INFORMATION

<b>Website</b>	URL: <a href="http://www.loadbalancer.org">www.loadbalancer.org</a>
<b>North America (US)</b>	<p>Loadbalancer.org, Inc.  4250 Lancaster Pike, Suite 120  Wilmington  DE 19805  USA</p> <p>Tel: +1 888.867.9504  Fax: +1 302.213.0122  Email (sales): <a href="mailto:sales@loadbalancer.org">sales@loadbalancer.org</a>  Email (support): <a href="mailto:support@loadbalancer.org">support@loadbalancer.org</a></p>
<b>North America (Canada)</b>	<p>Loadbalancer.org Ltd  300-422 Richards Street  Vancouver, BC  V6B 2Z4  Canada</p> <p>Tel: +1 866.998.0508  Fax: +1 302.213.0122  Email (sales): <a href="mailto:sales@loadbalancer.org">sales@loadbalancer.org</a>  Email (support): <a href="mailto:support@loadbalancer.org">support@loadbalancer.org</a></p>
<b>Europe (UK)</b>	<p>Loadbalancer.org Ltd.  Compass House  North Harbour Business Park  Portsmouth, PO6 4PS  UK</p> <p>Tel: +44 (0)330 3801064  Fax: +44 (0)870 4327672  Email (sales): <a href="mailto:sales@loadbalancer.org">sales@loadbalancer.org</a>  Email (support): <a href="mailto:support@loadbalancer.org">support@loadbalancer.org</a></p>
<b>Europe (Germany)</b>	<p>Loadbalancer.org GmbH  Tengstraße 27  D-80798  München  Germany</p> <p>Tel: +49 (0)89 2000 2179  Fax: +49 (0)30 920 383 6495  Email (sales): <a href="mailto:vertrieb@loadbalancer.org">vertrieb@loadbalancer.org</a>  Email (support): <a href="mailto:support@loadbalancer.org">support@loadbalancer.org</a></p>