

Load Balancing Philips Radiology Information System (RIS) / Workflow Information Management (WIM)

Version 1.3



Table of Contents

1. About this Guide	5
1.1. Acronyms & Terminology Used in the Guide	5
2. Prerequisites	5
3. Software Versions Supported	5
3.1. Loadbalancer.org Appliance	6
3.2. Philips RIS/WIM	6
4. Load Balancing RIS	6
4.1. Virtual Services (VIP) Requirements	6
4.2. SSL Termination	7
5. Ports Used by the Appliance	7
6. Deployment Concept	7
7. Load Balancer Deployment Methods	8
7.1. Layer 7 SNAT Mode	8
7.2. Layer 4 DR Mode	9
8. Configuring RIS for Load Balancing	10
8.1. Layer 7 SNAT Mode	10
8.2. Layer 4 DR Mode	10
8.2.1. Windows Server 2019	10
8.2.1.1. Step 1 of 3: Install the Microsoft Loopback Adapter	11
8.2.1.2. Step 2 of 3: Configure the Loopback Adapter	11
8.2.1.3. Step 3 of 3: Configure the strong/weak host behavior	13
8.2.1.4. Verify the Strong/Weak Host Configuration	13
9. Load Balancer Appliance Installation & Configuration for RIS	15
9.1. Overview	15
9.2. Hardware Appliance Installation	15
9.3. Virtual Appliance Installation	16
9.3.1. Download & Extract the Appliance	16
9.3.2. Virtual Hardware Resource Allocations	16
9.3.2.1. Virtual Hardware Resource Requirements	16
9.3.3. VMware vSphere Client	16
9.3.3.1. Upgrading to the latest Hardware Version	16
9.3.3.2. Installing the Appliance using vSphere Client	17
9.3.3.3. Configure Network Adapters	20
9.3.3.4. Start the Appliance	21
9.3.4. Microsoft Hyper-V	21
9.3.4.1. Installing the Appliance using Hyper-V Manager	21
9.3.4.2. Configure Network Adapters	24
9.3.4.3. Start the Appliance	24
9.4. Configuring Initial Network Settings	24
9.5. Accessing the Appliance WebUI	29
9.5.1. Main Menu Options	30
9.6. Installing the License Key	30
9.7. Appliance Software Update	31
9.7.1. Online Update	31
9.7.2. Offline Update	32
9.8. Appliance Security Mode Configuration	33
9.9. Appliance Network Configuration	33
9.9.1. Verify Network Connections	33

9.9.2. Configuring Hostname & DNS	34
9.9.3. Configuring NTP	34
9.10. Configuring Layer 4 Default Health Check Settings	34
9.11. Configuring Load Balanced RIS Services	35
9.11.1. VIP 1 - RIS	35
9.11.1.1. Virtual Service (VIP) Configuration	35
9.11.1.2. Define the Associated Real Servers (RIPs)	36
9.11.1.3. Upload the SSL Certificate	37
9.11.1.4. Configuring SSL Termination	37
9.11.2. VIP 1-B1 - RIS_EM	38
9.11.2.1. Virtual Service (VIP) Configuration	38
9.11.2.2. Define the Associated Real Servers (RIPs)	39
9.11.3. VIP 1-B2 - RIS_IM	40
9.11.3.1. Virtual Service (VIP) Configuration	40
9.11.3.2. Define the Associated Real Servers (RIPs)	41
9.11.4. VIP 1-B3 - RIS_ServicesWS	42
9.11.4.1. Virtual Service (VIP) Configuration	42
9.11.4.2. Define the Associated Real Servers (RIPs)	43
9.11.5. VIP 1-B4 - RIS_Auth	43
9.11.5.1. Virtual Service (VIP) Configuration	43
9.11.5.2. Define the Associated Real Servers (RIPs)	44
9.11.6. VIP 1-B5 - RIS_WS	45
9.11.6.1. Virtual Service (VIP) Configuration	45
9.11.6.2. Define the Associated Real Servers (RIPs)	46
9.11.7. VIP 1-B6 - RIS_WEB	47
9.11.7.1. Virtual Service (VIP) Configuration	47
9.11.7.2. Define the Associated Real Servers (RIPs)	48
9.11.8. VIP 1-B7 - RIS_IWEB_sms	48
9.11.8.1. Virtual Service (VIP) Configuration	48
9.11.8.2. Define the Associated Real Servers (RIPs)	49
9.11.9. VIP 1-B8 - RIS_IWEB_gbad	50
9.11.9.1. Virtual Service (VIP) Configuration	50
9.11.9.2. Define the Associated Real Servers (RIPs)	51
9.11.10. VIP 1-B9 - RIS_IWEBP_ws	51
9.11.10.1. Virtual Service (VIP) Configuration	51
9.11.10.2. Define the Associated Real Servers (RIPs)	52
9.11.11. VIP 2 - Patient_Assistance	53
9.11.11.1. Virtual Service (VIP) Configuration	53
9.11.11.2. Define the Associated Real Servers (RIPs)	55
9.11.11.3. Upload the SSL Certificate	55
9.11.11.4. Configuring SSL Termination	56
9.11.12. VIP 3 - R2I	57
9.11.12.1. Create the Custom Health Check	57
9.11.12.2. Virtual Service (VIP) Configuration	58
9.11.12.3. Define the Associated Real Servers (RIPs)	59
9.11.13. VIP 4 - EIS	59
9.11.13.1. Virtual Service (VIP) Configuration	59
9.11.13.2. Define the Associated Real Servers (RIPs)	60
9.11.14. VIP 5 - External_SMS	61
9.11.14.1. Virtual Service (VIP) Configuration	61
9.11.14.2. Define the Associated Real Servers (RIPs)	62

9.11.14.3. Upload the SSL Certificate	62
9.11.14.4. Configuring SSL Termination	63
9.11.15. Finalizing the Configuration	64
10. Configuring RIS Components to use the Load Balancer	64
11. Testing & Verification	66
12. Configuring HA - Adding a Secondary Appliance	68
12.1. Non-Replicated Settings	68
12.2. Configuring the HA Clustered Pair	69
13. Optional Appliance Configuration	70
13.1. SNMP Configuration	70
13.2. Configuring Email Alerts for Virtual Services	72
13.2.1. Layer 4	72
13.2.1.1. Global Layer 4 Email Settings	72
13.2.1.2. VIP Level Settings	72
13.2.2. Layer 7	73
13.3. Configuring Email Alerts for Heartbeat	74
13.4. Configuring a Smart Host (SMTP relay)	74
14. Technical Support	75
15. Further Documentation	75
16. Appendix	76
16.1. Generating a CSR on the Load Balancer	76
17. Document Revision History	78

1. About this Guide

This guide details the steps required to configure a load balanced Philips Healthcare Radiology Information System (RIS) version 11.6 and Workflow Information Management (WIM) version 15.0 environments using Enterprise Flex (Loadbalancer.org, formerly called R20) appliances. It covers the configuration of the load balancers and any RIS/WIM configuration changes that are required to enable load balancing. This configuration is also used for all RIS/WIM modules, such as the Extensible Integration System (EIS), and R2I for example.

Note

In the context of this document, the product name Radiology Information System may be referred to as RIS, and Workflow Information Management may be referred to as WIM.

1.1. Acronyms & Terminology Used in the Guide

Acronym	Description
RIS	Radiology Informatics System
WIM	Workflow Information Management
R2I	Radiology to Interface
EIS	Extensible Integration Software
VIP	Virtual IP address - the IP address of the load balanced cluster of RIPs, the address presented to connecting clients. Also refers to the logical load balancer configuration and is used as an acronym for Virtual Service.
RIP	The Real IP address - the IP address of a backend server in the cluster. Also refers to the logical load balancer configuration and is used as an acronym for Real Server.
Virtual Service	The main building block used to define load balanced services. It defines the IP address clients connect to, which Real Servers are load balanced and other settings such as health check options, persistence options and timeout settings.
Real Server	The actual backend server being load balanced. Multiple Real Servers are associated with a Virtual Service.

2. Prerequisites

1. Ensure that firewalls and other network devices are configured to allow management and other required access to the appliance - for details of all ports used refer to [Ports Used by the Appliance](#).
2. Ensure that firewalls and other network devices are configured to allow client/test access to all Virtual Services (VIPs).
3. Ensure that firewalls and other network devices are configured to allow load balancer access to all RIS servers.
4. Have IP addresses for the appliance and all required Virtual Services.

3. Software Versions Supported



3.1. Loadbalancer.org Appliance

- V8.8.1 & later

3.2. Philips RIS/WIM

- Philips RIS V11.6 & later
- WIM V15.0 & later

Note

Throughout this document the term RIS is interchangeable with the term WIM.

4. Load Balancing RIS

Note

It's highly recommended that you have a working RIS environment first before implementing the load balancer.

4.1. Virtual Services (VIP) Requirements

To provide load balancing and HA for RIS, the following VIPs are required:

Ref.	VIP Name	Mode/Type	Port(s)	Persistence	Health Check
VIP 1	RIS	L7 SNAT	80	None	No Checks, Always On
VIP 1-B1	RIS_EM	Backend Only	-	HTTP Cookie	HTTPS (GET)
VIP 1-B2	RIS_IM	Backend Only	-	HTTP Cookie	HTTPS (GET)
VIP 1-B3	RIS_ServicesWS	Backend Only	-	HTTP Cookie	HTTPS (GET)
VIP 1-B4	RIS_Auth	Backend Only	-	None	HTTPS (GET)
VIP 1-B5	RIS_WS	Backend Only	-	None	HTTPS (GET)
VIP 1-B6	RIS_WEB	Backend Only	-	None	HTTPS (GET)
VIP 1-B7	RIS_IWEB_sms	Backend Only	-	HTTP Cookie	HTTPS (GET)
VIP 1-B8	RIS_IWEB_gbad	Backend Only	-	HTTP Cookie	HTTPS (GET)
VIP 1-B9	RIS_IWEBP_ws	Backend Only	-	HTTP Cookie	ICMP Ping
VIP 2	Patient_Assistance	L7 SNAT	80	None	HTTPS (GET)
VIP 3	R2I	L4 DR	* (all ports)	None	External Script
VIP 4	EIS	L4 DR	* (all ports)	Source IP	HTTPS (GET)
VIP 5	External_SMS	L7 SNAT	80	HTTP Cookie	Connect to port

Note

VIPs with names in the format **VIP 1-B<number>** are **Backend Only** VIPs. These are used to define a pool of Real Servers. ACL's are then used by the 'parent' VIP (VIP 1) to determine which **Backend Only** VIP should be selected based on the requested URL.



Important

Ensure that the required DNS records are created that point to the appropriate VIP. If DNS records already exist, ensure that they are modified to point to the VIP rather than individual servers.

4.2. SSL Termination

SSL Termination is configured on the load balancer for the following VIPs:

- VIP 1 - **RIS**
- VIP 2 - **Patient_Assistance**
- VIP 5 - **External_SMS**

This provides a corresponding HTTPS Virtual Service for each VIP. Certificates in PEM or PFX format can be uploaded or if required a CSR can be generated on the load balancer to request a new certificate.

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 Addresses](#).

6. Deployment Concept






7. Load Balancer Deployment Methods

For RIS, both layer 4 DR mode and layer 7 SNAT mode are used. These modes are described below and are used for the configurations presented in this guide.

7.1. Layer 7 SNAT Mode

Layer 7 SNAT mode uses a proxy (HAProxy) at the application layer. Inbound requests are terminated on the load balancer and HAProxy generates a new corresponding request to the chosen Real Server. As a result, Layer 7 is typically not as fast as the Layer 4 methods. Layer 7 is typically chosen when either enhanced options such as SSL termination, cookie based persistence, URL rewriting, header insertion/deletion etc. are required, or when the network topology prohibits the use of the layer 4 methods. The image below shows an example network diagram for this mode.



- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network including across the Internet or WAN.


- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's own IP address by default, or any other local appliance IP address if preferred (e.g. the VIP address). This can be configured per layer 7 VIP. If required, the load balancer can be configured to provide the actual client IP address to the Real Servers in 2 ways. Either by inserting a header that contains the client's source IP address, or by modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. For more information on these methods please refer to [Transparency at Layer 7](#).
- Layer 7 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, **eth1** is typically used for client side connections and **eth0** is used for Real Server connections, although this is not mandatory since any interface can be used for any purpose.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

7.2. Layer 4 DR Mode

Layer 4 DR (Direct Routing) mode is a very high performance solution that requires little change to your existing infrastructure. The image below shows an example network diagram for this mode.

Note

Kemp, Brocade, Barracuda & A10 Networks call this *Direct Server Return* and F5 call it *nPath*.



- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that each Real Server (and the load balanced application) must respond to both the Real Server's own IP address and the VIP.

- The Real Server should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Server in this way is referred to as "Solving the ARP Problem". For more information please refer to [DR Mode Considerations](#).
- On average, DR mode is 8 times quicker than NAT mode for HTTP and much faster for other applications such as Remote Desktop Services, streaming media and FTP.
- The load balancer must have an interface in the same subnet as the Real Servers to ensure layer 2 connectivity which is required for DR mode to operate.
- The VIP can be brought up on the same subnet as the Real Servers or on a different subnet provided that the load balancer has an interface in that subnet.
- Port translation is not possible with DR mode, e.g. VIP:80 → RIP:8080 is not supported.
- DR mode is transparent, i.e. the Real Server will see the source IP address of the client.

8. Configuring RIS for Load Balancing

8.1. Layer 7 SNAT Mode

Layer 7 SNAT mode VIPs do not require any mode specific configuration changes to the load balanced Real Servers (RIS Servers).

8.2. Layer 4 DR Mode

Layer 4 DR mode VIPs require the "ARP problem" to be solved on each load balanced Real Server (RIS Server). This enables DR mode to work correctly.

The "ARP problem" must be solved on each Real Server associated with the following VIPs:

- VIP 3 - **R2I**
- VIP 4 - **EIS**

Detailed steps on solving the "ARP problem" for Windows 2019 are presented below.

Important

If EIS v15.0 or later is installed on the same servers as R2I, it's not possible to add 2 Loopback adapters. Instead, configure 2 IP addresses on the same Loopback adapter - the first that corresponds to the R2I VIP address and the second that corresponds to the EIS VIP address.

8.2.1. Windows Server 2019


Windows Server 2019 supports Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter that must be installed and configured on each load balanced (Real) Server. The IP address configured on the Loopback Adapter must be the same as the Virtual Service (VIP) address. This enables the server to receive packets that have their destination set as the VIP address. If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

In addition, steps must be taken to set the strong/weak host behavior on each Real Server. This is used to either prevent or allow interfaces to receive packets destined for a different interface on the same server.

Important The following 3 steps must be completed on **all** WIM R2I and EIS servers.

8.2.1.1. Step 1 of 3: Install the Microsoft Loopback Adapter


1. Click **Start**, then run **hdwwiz** to start the Hardware Installation Wizard.
2. Once the Wizard has started, click **Next**.
3. Select **Install the hardware that I manually select from a list (Advanced)**, click **Next**.
4. Select **Network adapters**, click **Next**.




5. Select **Microsoft & Microsoft KM-Test Loopback Adapter**, click **Next**.
6. Click **Next** to start the installation, when complete click **Finish**.

8.2.1.2. Step 2 of 3: Configure the Loopback Adapter

1. Open Control Panel and click **Network and Sharing Center**.
2. Click **Change adapter settings**.
3. Right-click the new Loopback Adapter and select **Properties**.
4. Uncheck all items except **Internet Protocol Version 4 (TCP/IPv4)** as shown below:



5. Ensure that **Internet Protocol Version (TCP/IPv4)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service address (VIP) with a subnet mask of **255.255.255.255**, e.g. **192.168.2.20/255.255.255.255** as shown below:



Note **192.168.2.20** is an example, make sure you specify the correct VIP address.

Note If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

6. Click **OK** then click **Close** to save and apply the new settings.

8.2.1.3. Step 3 of 3: Configure the strong/weak host behavior

The strong/weak host behavior can be configured using either of the following 2 methods:

- Option 1 - Using network shell (netsh) commands
- Option 2 - Using PowerShell cmdlets

The commands in this section assume that the LAN Adapter is named "**net**" and the Loopback Adapter is named "**loopback**" as shown in the example below:



Important

Either adjust the commands to use the names allocated to your LAN and loopback adapters, or rename the adapters before running the commands. Names are case sensitive so make sure that the interface names used in the commands match the adapter names exactly.

Option 1 - Using Network Shell (netsh) Commands

To configure the correct strong/weak host behavior run the following commands:

```
netsh interface ipv4 set interface "net" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostsend=enabled
```

Option 2 - Using PowerShell Cmdlets

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled
-DadTransmits 0 -AddressFamily IPv4
```

```
Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled -AddressFamily IPv4
```

Repeat steps 1 to 3 for all WIM R2I and EIS servers.


8.2.1.4. Verify the Strong/Weak Host Configuration

The following PowerShell Cmdlets can be used to verify the settings:

The "net" interface:



```
netsh interface ipv4 show interface net
```



Administrator: Windows PowerShell

```
PS C:\Users\Administrator> netsh interface ipv4 show interface net
```


Interface net Parameters

IfLuid	: ethernet_1
IfIndex	: 6
State	: connected
Metric	: 25
Link MTU	: 1500 bytes
Reachable Time	: 40500 ms
Base Reachable Time	: 30000 ms
Retransmission Interval	: 1000 ms
DAD Transmits	: 3
Site Prefix Length	: 64
Site Id	: 1
Forwarding	: disabled
Advertising	: disabled
Neighbor Discovery	: enabled
Neighbor Unreachability Detection	: enabled
Router Discovery	: dhcp
Managed Address Configuration	: enabled
Other Stateful Configuration	: enabled
Weak Host Sends	: disabled
Weak Host Receives	: enabled
Use Automatic Metric	: enabled
Ignore Default Routes	: disabled
Advertised Router Lifetime	: 1800 seconds
Advertise Default Route	: disabled
Current Hop Limit	: 0
Force ARPND Wake up patterns	: disabled
Directed MAC Wake up patterns	: disabled
ECN capability	: application

```
PS C:\Users\Administrator>
```

The "loopback" interface:

```
netsh interface ipv4 show interface loopback
```



Administrator: Windows PowerShell

```
PS C:\Users\Administrator> netsh interface ipv4 show interface loopback
```

Interface loopback Parameters

IfLuid	: ethernet_3
IfIndex	: 3
State	: connected
Metric	: 25
Link MTU	: 1500 bytes
Reachable Time	: 28000 ms
Base Reachable Time	: 30000 ms
Retransmission Interval	: 1000 ms
DAD Transmits	: 3
Site Prefix Length	: 64
Site Id	: 1
Forwarding	: disabled
Advertising	: disabled
Neighbor Discovery	: enabled
Neighbor Unreachability Detection	: enabled
Router Discovery	: dhcp
Managed Address Configuration	: enabled
Other Stateful Configuration	: enabled
Weak Host Sends	: enabled
Weak Host Receives	: enabled
Use Automatic Metric	: enabled
Ignore Default Routes	: disabled
Advertised Router Lifetime	: 1800 seconds
Advertise Default Route	: disabled
Current Hop Limit	: 0
Force ARPND Wake up patterns	: disabled
Directed MAC Wake up patterns	: disabled
ECN capability	: application

```
PS C:\Users\Administrator>
```

9. Load Balancer Appliance Installation & Configuration for RIS

9.1. Overview

For RIS deployments, 2 load balancer appliances must be installed and configured and then paired to create an active/passive HA clustered pair.

The following is an overview of the installation and configuration process:

1. Deploy 2 appliances - either hardware or virtual (VMware or Hyper-V) - refer to [Section 9.2](#) or [Section 9.3](#)
2. Configure the management IP address and other basic settings on **both** appliances - refer to [Section 9.4](#)
3. Run a software update check on **both** appliances - refer to [Section 9.7](#)
4. Configure the appliance security mode on **both** appliances - refer to [Section 9.8](#)
5. Verify network connections and configure any additional settings on **both** appliances - refer to [Section 9.9](#)
6. Configure the required load balanced services on the **Primary** appliance - refer to [Section 9.11](#)
7. Restart services on the **Primary** appliance - refer to [Section 9.11.15](#)
8. Verify that everything is working as expected on the **Primary** appliance - refer to [Section 11](#)
9. Configure the HA Pair on the **Primary** appliance - this will replicate all load balanced services to the Secondary appliance, once configured the Secondary appliance will be kept in-sync automatically - refer to [Section 12](#)
10. Configure any required optional settings on **both** appliances - refer to [Section 13](#)

9.2. Hardware Appliance Installation

Follow the steps below to install 2 appliances, one as the Primary the other as the Secondary.

1. Remove all packaging and rack mount the appliance if required.
2. Connect the power lead.




Note

The power supply is an auto-sensing unit (100v to 240v).

3. Connect network cables from all 4 interfaces (eth0 to eth3) to the relevant switch. All interfaces are configured as a single bond in the following section ([Section 9.4](#)).
4. Attach a monitor & keyboard to the appliance.





Note

The above image shows the Enterprise Loadbalancer on the Dell R350. Other supported Philips models include the Enterprise R20, Enterprise 1G and Dell R240.

Once both appliances are installed, connect a serial (heartbeat) cable between them.

Check that mains power is on and power up both appliances. The fans should start & the front panel LEDs should light.

9.3. Virtual Appliance Installation

Follow the relevant sections below to download and deploy 2 virtual appliances (either VMware or Hyper-V), one as the Primary the other as the Secondary.

9.3.1. Download & Extract the Appliance

1. Download the required Virtual Appliance.
2. Unzip the contents of the file to your chosen location.

9.3.2. Virtual Hardware Resource Allocations

By default the appliance is allocated the following resources:

- 2 vCPUs
- 4GB RAM
- 20GB disk

9.3.2.1. Virtual Hardware Resource Requirements

Configure the appliance's vCPU and RAM resources according to the following table:

Required Throughput	Processors	Memory
1Gbps	2 vCPU (default)	4GB (default)
4Gbps	4 vCPU	8GB

9.3.3. VMware vSphere Client

The steps below apply to VMware ESX/ESXi & vSphere Client v6.7 and later.

9.3.3.1. Upgrading to the latest Hardware Version



When the appliance is deployed, the virtual hardware version is set to 11. This enables compatibility with ESX version 6.0 and later. You can upgrade to a later hardware version if required.

 Note

Create a snapshot or backup of the virtual machine first before upgrading.

9.3.3.2. Installing the Appliance using vSphere Client

1. Right-click the inventory object where the appliance is to be located and select **Deploy OVF Template**.
2. In the **Select an OVF Template** screen, select the **Local File** option, click **Browse**, navigate to the download location, select both the **.ovf** and **.vmdk** files and click **Next**.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

☐ URL

☒ Local file

http | https://remoteserver-address/filetodeploy.ovf | .ova

Choose Files 2 files

CANCEL

BACK

NEXT

3. In the **Select a name and folder** screen, type a suitable name for the appliance - this can be up to 80 characters in length.
4. Select the required location for the appliance - by default this will be the location of the inventory object from where the wizard was started and click **Next**.

Deploy OVF Template

- ✓ 1 Select an OVF template
- 2 Select a name and folder**
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name:

Select a location for the virtual machine.

▼

vcs.loadbalancer.org

>

Portsmouth

CANCEL

BACK

NEXT

5. In the **Select a compute resource** screen, select the required compute resource for the appliance - by default this will be the inventory object from where the wizard was started and click **Next**.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- 3 Select a compute resource**
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select a compute resource

Select the destination compute resource for this operation

▼

Portsmouth

▼

Compass House

host1.loadbalancer.org

host2.loadbalancer.org

Accounts

>

Dev

>

Solutions

>

Support

>

Test Lab - 192.168.120.x

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

6. In the **Review details** screen, verify the template details and click **Next**.



Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Review details**
- 5 Select storage
- 6 Select networks
- 7 Ready to complete

Review details

Verify the template details.

Publisher	No certificate present
Description	Loadbalancer.org VA - Traffic Management and Load Balancing Appliance from www.loadbalancer.org
Download size	437.9 MB
Size on disk	1.3 GB (thin provisioned)
	20.0 GB (thick provisioned)

CANCEL

BACK

NEXT

7. In the **Select Storage** screen, first select the required storage location for the appliance.

8. Now select the required disk format and click **Next**.

Note

Loadbalancer.org recommends selecting a thick provision format. By default the appliance disk is 20GB.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 Select storage**
- 6 Select networks
- 7 Ready to complete

Select storage

Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format:

Thick Provision Lazy Zeroed

VM Storage Policy:

Datastore Default

☐ Disable Storage DRS for this virtual machine

Name	Capacity	Provisioned	Free	Type	Cluster
Portsmouth Datastore	65.49 TB	25.65 TB	39.83 TB		
ISO Store	179.99 GB	86.77 GB	93.22 GB	NFS v3	
Linux Templates	196.98 GB	59.67 GB	137.32 GB	NFS v3	
Loadbalancer Appliance...	196.98 GB	109.67 GB	87.31 GB	NFS v3	
Windows Template Store	296.98 GB	184.39 GB	112.59 GB	NFS v3	

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

9. In the **Select Networks** screen, select the required destination network using the drop-down next to **VM Network** and click **Next**.



Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Select storage
- 6 Select networks**
- 7 Ready to complete

Select networks

Select a destination network for each source network.

Source Network	Destination Network
VM Network	Office Port Group
1 items	

IP Allocation Settings

IP allocation:

Static - Manual

IP protocol:

IPv4

CANCEL

BACK

NEXT

10. In the **Ready to complete** screen, review the settings and click **Finish** to create the virtual appliance. To change a setting, use the **Back** button to navigate back through the screens as required.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Select storage
- ✓ 6 Select networks
- 7 Ready to complete**

Ready to complete

Click Finish to start creation.

Provisioning type	Deploy from template
Name	Loadbalancer.org Enterprise VA
Template name	Loadbalancer.org Enterprise VA
Download size	437.9 MB
Size on disk	20.0 GB
Folder	Portsmouth
Resource	Solutions
Storage mapping	1
All disks	Datastore: Portsmouth Datastore; Format: Thick provision lazy zeroed
Network mapping	1
VM Network	VM Network
IP allocation settings	
IP protocol	IPv4
IP allocation	Static - Manual

CANCEL

BACK

FINISH

9.3.3.3. Configure Network Adapters

The appliance has 4 network adapters. By default only the first adapter is connected. Philips Healthcare deployments require all 4 network adapters to be connected and configured as a single bond as described in [Section 9.4](#). To connect additional network adapters:



1. Right-click the appliance, select **Edit Settings**.
2. Network adapter 1 should already be connected. This will be **eth0** when viewed in the appliance WebUI.
3. Select the required **Network** for **Network Adapter 2** and tick (check) the **Connected** check-box and click **OK**. This will be **eth1** when viewed in the appliance WebUI.
4. Select the required **Network** for **Network Adapter 3** and tick (check) the **Connected** check-box and click **OK**. This will be **eth2** when viewed in the appliance WebUI.
5. Select the required **Network** for **Network Adapter 4** and tick (check) the **Connected** check-box and click **OK**. This will be **eth3** when viewed in the appliance WebUI.

9.3.3.4. Start the Appliance


Now power up the appliance.

9.3.4. Microsoft Hyper-V


The steps below apply to Windows Hyper-V 2012 & later.

9.3.4.1. Installing the Appliance using Hyper-V Manager


1. Start Hyper-V Manager, then using the right-click menu or the Actions pane select **Import Virtual Machine** and click **Next**.
2. In the **Locate folder** screen, browse to the location of the extracted download and select the **Loadbalancer.org Ent. VA** folder.




3. Click **Next** until you reach the **Choose Import Type** screen, select the option **Copy the virtual machine (create a new unique ID)** and click **Next**.




4. In the **Choose Folder For Virtual Machine Files** screen, tick (check) the checkbox **Store the Virtual Machine in different location**, then select a suitable location for the virtual machines files and click **Next**.



5. In the **Choose Folder to store Virtual Hard Disks** screen, select a suitable location for the virtual hard disk files and click **Next**.



6. In the **Completing Import Wizard** screen, verify that all settings are correct and click **Finish** to complete the import process. To change a setting, use the **Previous** button to navigate back through the screens as required.



Once complete, the load balancer will appear in the **Virtual Machines** list.

Note

For a clustered pair, make sure that you select a different folder location in steps 4 & 5.

9.3.4.2. Configure Network Adapters

The appliance has 4 network adapters, these remain disconnected once deployment completes. Philips Healthcare deployments require all 4 network adapters to be connected and configured as a single bond as described in [Section 9.4](#). To connect the network adapters:

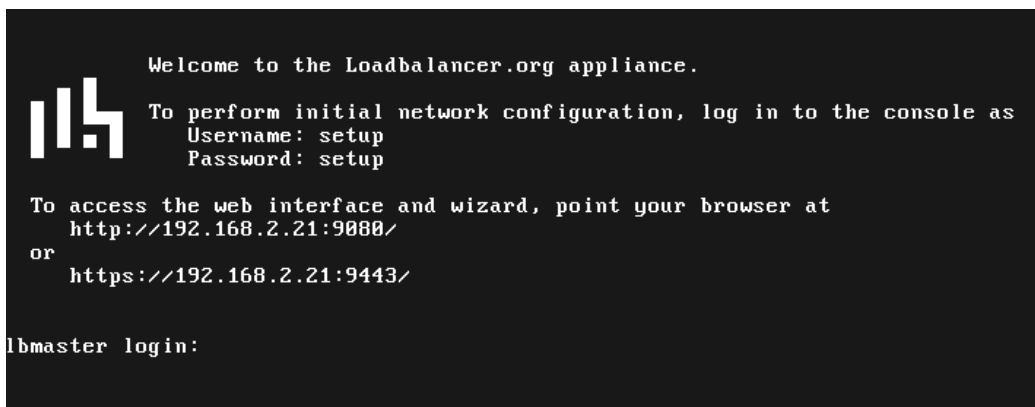
1. Right-click the appliance, select **Settings**.
2. Select the first network adapter and set the required virtual switch that the adapter should be connected to. This will be **eth0** when viewed in the appliance WebUI.
3. Select the second network adapter and set the required virtual switch that the adapter should be connected to. This will be **eth1** when viewed in the appliance WebUI.
4. Select the third network adapter and set the required virtual that the adapter should be connected to. This will be **eth2** when viewed in the appliance WebUI.
5. Select the forth network adapter and set the required virtual that the adapter should be connected to. This will be **eth3** when viewed in the appliance WebUI.
6. Click **Apply**.

9.3.4.3. Start the Appliance

Now power up the appliance.

9.4. Configuring Initial Network Settings

After power up, the following startup message is displayed on the appliance console:



```

Welcome to the Loadbalancer.org appliance.

To perform initial network configuration, log in to the console as
Username: setup
Password: setup

To access the web interface and wizard, point your browser at
http://192.168.2.21:9080/
or
https://192.168.2.21:9443/

lbmaster login:
```

As mentioned in the text, to perform initial network configuration, login as the "setup" user at the appliance console.

Once logged in, the Network Setup Wizard will start automatically. This will enable you to configure the management IP address and other network settings for the appliance.

login to the console:


Username: setup

Password: setup

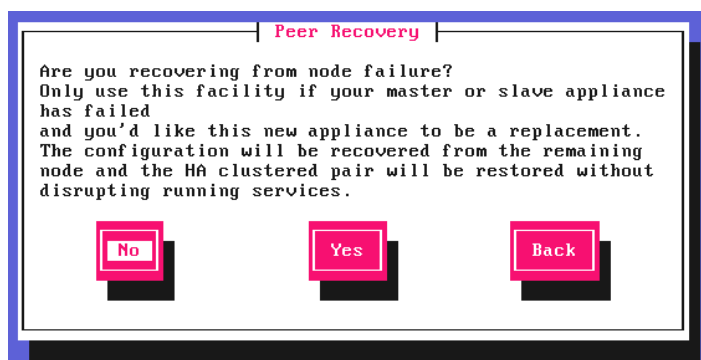
A series of screens will be displayed that allow network settings to be configured:



In the **Configure Management IP** screen, leave **Yes** selected and hit <ENTER> to continue.




In the **Peer Recovery** screen, leave **No** selected and hit <ENTER> to continue.



 Note

For more details on node recovery using this option please refer to [Disaster Recovery After Node \(Primary or Secondary\) Failure](#).


In the **Centralized Management** screen, if you have been provided with Management Server details select **Yes**, otherwise leave **No** selected, then hit <ENTER> to continue.




 Note

For information on how to modify Centralized Management settings via the WebUI, please refer to [Portal Management & Appliance Adoption](#).


In the **Available Interfaces** screen, a list of available interfaces will be displayed, hit <ENTER> to continue.




In the **Configure Bonding** screen, select **Yes** and hit <ENTER> to continue.




The **Select Interfaces** screen will be displayed. Using the space bar, select all 4 interfaces for the bond, select **Create** and hit <ENTER> to continue.



In the **Set Bond Mode** screen, select **Mode 4 - LACP 802.3ad**, select the **Select** button and hit <ENTER> to continue.



In the **Configure a VLAN** screen, leave **No** selected, then hit <ENTER> to continue.




In the **Set IP address** screen, enter the required *Static IP Address* & *CIDR Prefix* and select **Done** and hit <ENTER> to continue.




 Note

A subnet mask such as 255.255.255.0 is not valid, in this case enter 24 instead.

In the **Configure Default Gateway** screen, enter the required *Default Gateway IP Address*, select **Done** and hit <ENTER> to continue.




In the **Configure DNS Servers** screen, configure the required DNS server(s), select **Done** and hit <ENTER> to continue.




In the **Set Password** screen, hit <ENTER> to continue.



Enter the **Password** you'd like to use for the **loadbalancer** WebUI user account and the **root** Linux user account. Repeat the password, select **Done** and hit <ENTER> to continue.




If you selected **Yes** when asked if you want to enroll for Centralized Management, you'll now be prompted for the details. Default values for the **Host** and **Port** are set and can be changed if required. Enter the **Username** and **Password** for the management server account you'd like the appliance to be associated with, select **Done** and hit <ENTER> to continue.



In the **Summary** screen, check all settings. If everything is correct, leave **Configure** selected and hit <ENTER> to continue. All settings will be applied. If you need to change a setting, use the **Back** button.

Note

For v8.13.2 and later, once the settings have been applied the appliance will check if a software update is available. If an update is found, it will be installed automatically.



Once the configuration has been written, the **Configuration Complete** screen and message will be displayed. Click **OK** to exit the wizard and return to the command prompt.



9.5. 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](#).

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:



9.5.1. 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

9.6. Installing the License Key

The appliance can be used completely unrestricted for 30 days without installing a license key. After 30 days, the appliance continues to work but it's no longer possible to make configuration changes.

For an unlicensed VA, the following message is displayed:

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](#)

For an unlicensed hardware appliance, the following message is displayed:

WARNING: This appliance is unregistered. **Please enter your license key** within 30 days to activate your appliance.
If you do not have your license key please **Contact Us**

To install the license key:

1. Using the WebUI, navigate to: *Local Configuration > License Key*.

Install License Key

This unit is in evaluation mode. Please enter your license key to remove this restriction.

If you do not have a license key, please contact sales@loadbalancer.org.

No file chosen

[Install License Key](#)

2. Click **Choose File** then browse to and select the license file provided when the appliance was purchased.
3. Click **Install License Key**.

Note

Once the license is applied, these warning messages will no longer be displayed.

9.7. 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

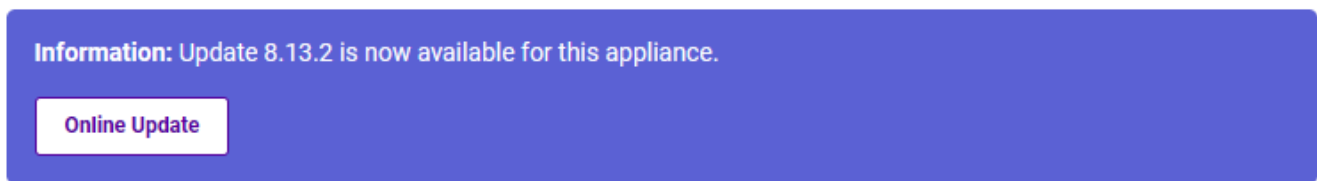
Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

9.7.1. Online Update


The appliance periodically contacts the Loadbalancer.org update server (update.loadbalancer.org) and checks



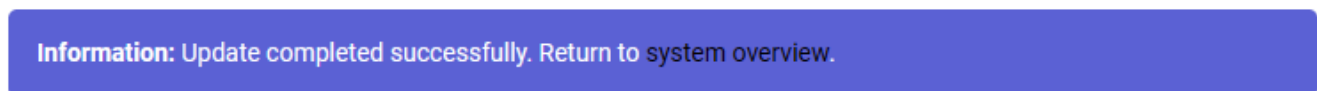
for updates. This is the default behavior and can be disabled if preferred. If an update is found, a notification similar to the example below will be displayed at the top of the WebUI:



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.



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



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

9.7.2. 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:

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.

9.8. Appliance Security Mode Configuration

To enable shell commands to be run from the WebUI, the appliance Security Mode must be configured:

1. Using the WebUI, navigate to: *Local Configuration > Security*.
2. Set *Appliance Security Mode* to **Custom**.
3. Click **Update**.


9.9. Appliance Network Configuration

The standard RIS network configuration uses all 4 network adapters configured as an LACP 802.3ad bonded interface.

9.9.1. Verify Network Connections

1. Verify that all 4 network adapters are connected to the appropriate switch.
2. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.

IP Address Assignment



The diagram shows a network configuration interface. At the top, a box labeled 'bond0' contains four network interface icons labeled 'eth0', 'eth1', 'eth2', and 'eth3', each with '10 GB/s' below it. Below this, a large text input field for 'bond0' contains the IP address '10.224.94.45/24'. To the right of this field is an 'MTU' label and a numeric input field set to '1500' with 'bytes' below it. A purple button labeled 'Configure Interfaces' is at the bottom right.

3. Verify that the network is configured as required.

Note

The IP address/CIDR prefix for **bond0** was set during the Network Setup Wizard and will be shown here, e.g. **10.224.94.45/24**.

9.9.2. Configuring Hostname & DNS

1. Using the WebUI, navigate to: *Local Configuration > Hostname & DNS*.
2. Set the required *Hostname* and *Domain Name*.
3. Configure additional DNS servers if required.
4. Click **Update**.

9.9.3. Configuring NTP

1. Using the WebUI, navigate to: *Local Configuration > System Date & Time*.
2. Select the required *System Timezone*.
3. Navigate to the first field in the *NTP Servers* section, specify the IP address of the RIS Database server.
4. Click **Set Timezone & NTP**.

9.10. Configuring Layer 4 Default Health Check Settings

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Advanced Configuration*.
2. Set the *Check Interval* to **10** (seconds).
3. Set the *Check Timeout* to **30** (seconds).
4. Set the *Failure Count* to **3** (seconds).
5. Click **Update**.



9.11. Configuring Load Balanced RIS Services

Note

All Virtual Services are listed in the table in [Virtual Services \(VIP\) Requirements](#).

9.11.1. VIP 1 - RIS

9.11.1.1. Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	RIS	?
IP Address	10.224.94.100	?
Ports	80	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel **Update**

- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS**.
 - Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.224.94.100**.
 - Set the *Ports* field to **80**.
 - Set the *Layer 7 Protocol* to **HTTP Mode**.
3. Click **Update** to create the Virtual Service.
 4. Now click **Modify** next to the newly created VIP.
 5. Scroll to the *Persistence* section.
 - Ensure that the *Persistence Mode* is set to **None**.
 6. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **No checks, Always On**.
 7. Scroll to the *ACL Rules* section, and click the **Add Rule** button.
 8. Set the **Type** to *Freetype* and copy and paste the following ACL rules into the *Freetype* field ensuring that the "#" is removed.

Important

When text is copied from a pdf document that extends beyond 1 line, additional unwanted line feed formatting characters are included. Before pasting the text into the *Freetype* field, first copy the text to an editor such as Windows Notepad, then edit the text so that each ACL rule is on a single line (ACL rules start with "http-request" or "use_backend"), making

sure there's a space between text that was on different lines. Once complete, copy and paste this text into the **Freetype** field. If this is not done, syntax checking will fail when the HAProxy service is re-started.

```
use_backend RIS_Auth if { path -m reg -i ^/[^/]*/authenticationserver/(oauth|api)/.+ } !{ path -m reg -i ^/[^/]+/authenticationserver/(oauth/token|api/changepassword) }
use_backend RIS_WS if { path_sub -i /WS/ }
use_backend RIS_ServicesWS if { path_sub -i /RIS_ServicesWS/ }
use_backend RIS_IM if { path_sub -i /InstantMessageWS/ }
use_backend RIS_EM if { path_sub -i /EnterpriseManagerWS/ }
use_backend RIS_Auth if { path_sub -i /AuthenticationServer/ }
use_backend RIS_WEB if { path_sub -i /RISWebServer/ }
use_backend RIS_WEB if { path -m reg -i ^/patientportal.? } || { path -m reg -i ^/web/(globals.bundle.js|index.html) } || { path -m reg -i ^/web/\d+/[^.]+\.[^.]*\.[^.]*\w+ } || { path -m reg -i ^/web/app/light\?shs=.+ }
use_backend RIS_WS if { path -m reg -i ^/[^/]+/riswebserver/api/.+ }
```

9. Click **OK** to save and close the ACL.
10. Scroll to the **SSL** section.
 - Enable (check) the **Enable Backend Encryption** checkbox.
11. Leave all other settings at their default value.
12. Click **Update**.

9.11.1.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

Cancel

Update

- Specify an appropriate **Label** for the RIP, e.g. **WS2019DC2**.
- Set the **Real Server IP Address** field to the required IP address, e.g. **10.224.94.42**.
- Set the **Real Server Port** field to **443**.




3. Leave all other settings at their default value.
4. Click **Update**.
5. Repeat these steps to add additional Real Server(s).

9.11.1.3. Upload the SSL Certificate

Note

If the production certificate is not currently available, continue to the next section (configuring SSL Termination) and leave **SSL Certificate** set to **Default Self Signed Certificate**. Once you have the production certificate, follow the steps in this section to upload the certificate and then navigate to *Cluster Configuration > SSL Termination*, click **Modify** next to the SSL termination, change **SSL Certificate** to the new certificate and click **Update**.

1. Using the WebUI, navigate to *Cluster Configuration > SSL Certificate* and click **Add a new SSL Certificate**.
2. Select the option **Upload prepared PEM/PFX file**.
3. Enter the following details:



- Specify an appropriate *Label*, e.g. **ris-cert**.
 - Click **Choose File**.
 - Browse to and select the relevant PEM or PFX file.
 - For PFX files specify the password if required.
4. Click **Upload Certificate**.

Note

If you don't have a certificate and want to create a Certificate Signing Request (CSR) on the load balancer, please refer to [Generating a CSR on the Load Balancer](#).

9.11.1.4. Configuring SSL Termination

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

Label	SSL-RIS	?
Associated Virtual Service	RIS	?
Virtual Service Port	443	?
SSL Operation Mode	High Security	
SSL Certificate	ris-cert	?
Source IP Address		?
Enable Proxy Protocol	☒	?
Bind Proxy Protocol to L7 VIP	RIS	?

Cancel
Update

- Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **RIS**.

Note

Once the VIP is selected, the *Label* field will be auto-populated with **SSL-RIS**. This can be changed if preferred.

- Ensure that the *Virtual Service Port* is set to **443**.
 - Leave *SSL Operation Mode* set to **High Security**.
 - Select the required *SSL Certificate*.
3. Leave all other settings at their default value.
 4. Click **Update**.

9.11.2. VIP 1-B1 - RIS_EM

9.11.2.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_EM	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel Update

- Select (Check) **Create Backend Only**.
- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_EM**.
- Leave the *Layer 7 Protocol* set to **HTTP Mode**.

4. Click **Update** to create the Virtual Service.

5. Now click **Modify** next to the newly created VIP.

6. Scroll to the *Persistence* section.

- Ensure that the *Persistence Mode* is set to **HTTP Cookie**.
- Set the *HTTP Cookie Name* to **RIS_EM**.

7. Scroll to the *Health Checks* section.

- Set the *Check Type* to **Negotiate HTTPS (GET)**.
- Set the *Request to Send* to **/CSHRIS-SGH11_6R20/EnterpriseManagerWS/api/HealthCheck/HealthCheck**.

Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.

8. Scroll to the *SSL* section.

- Enable (check) the *Enable Backend Encryption* checkbox.

9. Click **Update**.

9.11.2.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

Cancel
Update

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
 - Set the *Real Server Port* field to **443**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.3. VIP 1-B2 - RIS_IM

9.11.3.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_IM	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel
Update

- Select (Check) **Create Backend Only**.
- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_IM**.
- Leave the *Layer 7 Protocol* set to **HTTP Mode**.

4. Click **Update** to create the Virtual Service.
5. Now click **Modify** next to the newly created VIP.
6. Scroll to the *Persistence* section.
 - Ensure that the *Persistence Mode* is set to **HTTP Cookie**.
 - Set the *HTTP Cookie Name* to **RIS_IM**.
7. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate HTTPS (GET)**.
 - Set the *Request to Send* to **/CSHRIS-SGH11_6R20/InstantMessageWS/api/HealthCheck/HealthCheck**.



Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.
8. Scroll to the *SSL* section.
 - Enable (check) the *Enable Backend Encryption* checkbox.
 9. Leave all other settings at their default value.
 10. Click **Update**.

9.11.3.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
- Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
- Set the *Real Server Port* field to **443**.



3. Leave all other settings at their default value.
4. Click **Update**.
5. Repeat these steps to add additional Real Server(s).

9.11.4. VIP 1-B3 - RIS_ServicesWS

9.11.4.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_ServicesWS	?
Protocol		[Advanced +]
Layer 7 Protocol	HTTP Mode	?

Cancel Update

- Select (Check) **Create Backend Only**.
 - Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_ServicesWS**.
 - Leave the *Layer 7 Protocol* set to **HTTP Mode**.
4. Click **Update** to create the Virtual Service.
 5. Now click **Modify** next to the newly created VIP.
 6. Scroll to the *Persistence* section.
 - Ensure that the *Persistence Mode* is set to **HTTP Cookie**.
 - Set the *HTTP Cookie Name* to **RIS_ServicesWS**.
 7. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate HTTPS (GET)**.
 - Set the *Request to Send* to **/CSHRIS-SGH11_6R20/RISServicesWS/api/HealthCheck/HealthCheck**.

Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.



8. Scroll to the **SSL** section.
 - Enable (check) the **Enable Backend Encryption** checkbox.
9. Leave all other settings at their default value.
10. Click **Update**.

9.11.4.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

Cancel Update

- Specify an appropriate **Label** for the RIP, e.g. **WS2019DC2**.
 - Set the **Real Server IP Address** field to the required IP address, e.g. **10.224.94.42**.
 - Set the **Real Server Port** field to **443**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.5. VIP 1-B4 - RIS_Auth

9.11.5.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to **Cluster Configuration > Layer 7 – Virtual Services** and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_Auth	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel Update

- Select (Check) **Create Backend Only**.
- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_Auth**.
- Leave the *Layer 7 Protocol* set to **HTTP Mode**.

4. Click **Update** to create the Virtual Service.

5. Now click **Modify** next to the newly created VIP.

6. Scroll to the *Persistence* section.

- Set the *Persistence Mode* to **None**.

7. Scroll to the *Health Checks* section.

- Set the *Check Type* to **Negotiate HTTPS (GET)**.
- Set the *Request to Send* to **/CSHRIS-SGH11_6R20/AuthenticationServer/api/HealthCheck/HealthCheck**.

Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.

8. Scroll to the *SSL* section.

- Enable (check) the *Enable Backend Encryption* checkbox.

9. Leave all other settings at their default value.

10. Click **Update**.

9.11.5.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

Cancel
Update

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
- Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
- Set the *Real Server Port* field to **443**.

3. Leave all other settings at their default value.

4. Click **Update**.

5. Repeat these steps to add additional Real Server(s).

9.11.6. VIP 1-B5 - RIS_WS

9.11.6.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_WS	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel
Update

- Select (Check) **Create Backend Only**.
- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_WS**.
- Leave the *Layer 7 Protocol* set to **HTTP Mode**.

4. Click **Update** to create the Virtual Service.
5. Now click **Modify** next to the newly created VIP.
6. Scroll to the *Persistence* section.
 - Set the *Persistence Mode* to **None**.
7. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate HTTPS (GET)**.
 - Set the *Request to Send* to **/CSHRIS-SGH11_6R20/WS/Service.asmx/HealthCheck**.

 Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.
8. Scroll to the *SSL* section.
 - Enable (check) the *Enable Backend Encryption* checkbox.
 9. Leave all other settings at their default value.
 10. Click **Update**.

9.11.6.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	
Real Server IP Address	10.224.94.42	
Real Server Port	443	
Re-Encrypt to Backend	☒	
Enable Redirect	☐	
Weight	100	

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
 - Set the *Real Server Port* field to **443**.
3. Leave all other settings at their default value.

4. Click **Update**.
5. Repeat these steps to add additional Real Server(s).

9.11.7. VIP 1-B6 - RIS_WEB

9.11.7.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_WEB	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel Update

- Select (Check) **Create Backend Only**.
 - Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_WEB**.
 - Leave the *Layer 7 Protocol* set to **HTTP Mode**.
4. Click **Update** to create the Virtual Service.
 5. Now click **Modify** next to the newly created VIP.
 6. Scroll to the *Persistence* section.
 - Set the *Persistence Mode* to **None**.
 7. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate HTTPS (GET)**.
 - Set the *Request to Send* to **/cshris-sgh11_6R20/RISWebServer/api/HealthCheck/HealthCheck**.



Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.
8. Scroll to the *SSL* section.
 - Enable (check) the *Enable Backend Encryption* checkbox.



9. Leave all other settings at their default value.
10. Click **Update**.

9.11.7.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
 - Set the *Real Server Port* field to **443**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.8. VIP 1-B7 - RIS_IWEB_sms

9.11.8.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_IWEB_sms	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel Update

- Select (Check) **Create Backend Only**.
- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_IWEB_sms**.
- Leave the *Layer 7 Protocol* set to **HTTP Mode**.

4. Click **Update** to create the Virtual Service.

5. Now click **Modify** next to the newly created VIP.

6. Scroll to the *Health Checks* section.

- Set the *Check Type* to **Negotiate HTTPS (GET)**.
- Set the *Request to Send* to **/ec/0**.
- Set *Response Expected* to **Equals** and set the value to **200**.

7. Scroll to the *SSL* section.

- Enable (check) the *Enable Backend Encryption* checkbox.

8. Leave all other settings at their default value.

9. Click **Update**.

9.11.8.2. Define the Associated 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 VIP.

2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
- Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
- Set the *Real Server Port* field to **443**.

3. Leave all other settings at their default value.

4. Click **Update**.

5. Repeat these steps to add additional Real Server(s).

9.11.9. VIP 1-B8 - RIS_IWEB_gbad

9.11.9.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_IWEB_gbad	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

- Select (Check) **Create Backend Only**.
- Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_IWEB_gbad**.
- Leave the *Layer 7 Protocol* set to **HTTP Mode**.

4. Click **Update** to create the Virtual Service.
5. Now click **Modify** next to the newly created VIP.
6. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate HTTPS (GET)**.
 - Set the *Request to Send* to **/ec/gbad/0**.
 - Set *Response Expected* to **Equals** and set the value to **1**.
7. Scroll to the *SSL* section.
 - Enable (check) the *Enable Backend Encryption* checkbox.
8. Leave all other settings at their default value.
9. Click **Update**.

9.11.9.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
 - Set the *Real Server Port* field to **443**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.10. VIP 1-B9 - RIS_IWEBP_ws

9.11.10.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.

2. Click **[Advanced]**.
3. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	☐	?
Create Backend Only	☒	?
Label	RIS_IWEBP_ws	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel Update

- Select (Check) **Create Backend Only**.
 - Specify an appropriate *Label* for the Virtual Service, e.g. **RIS_IWEBP_ws**.
 - Leave the *Layer 7 Protocol* set to **HTTP Mode**.
4. Click **Update** to create the Virtual Service.
 5. Now click **Modify** next to the newly created VIP.
 6. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **External Script**.
 - Set the *check Script* to **Ping_IPv4_or_Ping_IPv6**.
 7. Scroll to the *SSL* section.
 - Enable (check) the *Enable Backend Encryption* checkbox.
 8. Leave all other settings at their default value.
 9. Click **Update**.

9.11.10.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

Cancel
Update

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
 - Set the *Real Server Port* field to **443**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.11. VIP 2 - Patient_Assistance

9.11.11.1. Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	Patient_Assistance	?
IP Address	10.224.94.101	?
Ports	80	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel
Update

- Specify an appropriate *Label* for the Virtual Service, e.g. **Patient_Assistance**.
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.224.94.101**.
- Set the *Ports* field to **80**.

- Set the *Layer 7 Protocol* to **HTTP Mode**.
3. Click **Update** to create the Virtual Service.
 4. Now click **Modify** next to the newly created VIP.
 5. Scroll to the *Persistence* section.
 - Ensure that the *Persistence Mode* is set to **None**.
 6. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Neogtiate HTTPS (GET)**.
 - Set *Request to Send* to **/CSHRIS-SGH11_6R20/RISWebServer/api/HealthCheck/HealthCheck**.

Note

CSHRIS-SGH11_6R20 is the current RIS core application folder name on the RIS servers.

- Set *Response Expected* (case sensitive field) to **Equals** and set the value to **true**.
7. Scroll to the *ACL Rules* section, and click the **Add Rule** button.
 8. Set the **Type** to *Freetype* and copy and paste the following ACL rules into the *Freetype* field ensuring that the "#" is removed.

Important

When text is copied from a pdf document that extends beyond 1 line, additional unwanted line feed formatting characters are included. Before pasting the text into the *Freetype* field, first copy the text to an editor such as Windows Notepad, then edit the text so that each ACL rule is on a single line (ACL rules start with "acl", "http-request" or "use_backend"), making sure there's a space between text that was on different lines. Once complete, copy and paste this text into the *Freetype* field. If this is not done, syntax checking will fail when the HAProxy service is re-started.

```
# ACL rules for allowed URLs (case insensitive)
acl patient path,url_dec -i -m beg /patientportal
acl webEndingGlobals path,url_dec -i /web/globals.bundle.js
acl webEndingIndex path,url_dec -i /web/index.html
acl webNumber path,url_dec -i -m reg ^/web/[0-9]+/.+\.+$
acl risweb path,url_dec -i -m reg ^/[^/]+/riswebserver/api/.+
acl authAPI path,url_dec -i -m reg ^/[^/]+/authenticationserver/api/.+
acl authOAuth path,url_dec -i -m reg ^/[^/]+/authenticationserver/oauth/.+
acl webForms path,url_dec -i -m beg /web/forms
# ACL rules for denied URLs (case insensitive)
acl denyOAuthToken path,url_dec -i -m beg /authenticationserver/oauth/token
acl denyChangePassword path,url_dec -i -m beg /authenticationserver/api/changepassword
acl denyPatientKiosk path,url_dec -i -m beg /patientportal/patientkiosk
acl denyRiswebPatientKiosk path,url_dec -i -m beg /riswebserver/api/patientkiosk
acl denyOAuthPatientKioskAuth path,url_dec -i -m beg
/authenticationserver/oauth/patientkioskauth
# Deny requests for specific denied URLs with 403 status
http-request deny deny_status 403 if denyOAuthToken
http-request deny deny_status 403 if denyChangePassword
http-request deny deny_status 403 if denyPatientKiosk
http-request deny deny_status 403 if denyRiswebPatientKiosk
```

```
http-request deny deny_status 403 if denyOAuthPatientKioskAuth
# Reject all other requests unless they match allowed URL ACLs
http-request deny deny_status 403 unless patient || webEndingGlobals || webEndingIndex ||
webNumber || risweb || authAPI || authOAuth || webForms
# Route to backends based on allowed URLs
use_backend RIS_WEB if patient || webEndingGlobals || webEndingIndex || webNumber || webForms {
nbsrv(RIS_WEB) gt 0 }
use_backend RIS_WS if risweb { nbsrv(RIS_WS) gt 0 }
use_backend RIS_Auth if authAPI || authOAuth { nbsrv(RIS_Auth) gt 0 }
```

9. Click **OK** to save and close the ACL.
10. Scroll to the **SSL** section.
 - Enable (check) the **Enable Backend Encryption** checkbox.
11. Leave all other settings at their default value.
12. Click **Update**.

9.11.11.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

Cancel Update

- Specify an appropriate **Label** for the RIP, e.g. **WS2019DC2**.
 - Set the **Real Server IP Address** field to the required IP address, e.g. **10.224.94.42**.
 - Set the **Real Server Port** field to **443**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.11.3. Upload the SSL Certificate



Note

If the production certificate is not currently available, continue to the next section (configuring



SSL Termination) and leave **SSL Certificate** set to **Default Self Signed Certificate**. Once you have the production certificate, follow the steps in this section to upload the certificate and then navigate to **Cluster Configuration > SSL Termination**, click **Modify** next to the SSL termination, change **SSL Certificate** to the new certificate and click **Update**.

1. Using the WebUI, navigate to **Cluster Configuration > SSL Certificate** and click **Add a new SSL Certificate**.
2. Select the option **Upload prepared PEM/PFX file**.
3. Enter the following details:

I would like to:

- ☒ Upload prepared PEM/PFX file
- ☐ Create a new SSL Certificate Signing Request (CSR)
- ☐ Create a new Self-Signed SSL Certificate.

Label: patient-assistance-cert

File to upload: Choose File patient-assi...nce-cert.pem

Upload Certificate

- Specify an appropriate **Label**, e.g. **patient-assistance-cert**.
 - Click **Choose File**.
 - Browse to and select the relevant PEM or PFX file.
 - For PFX files specify the password if required.
4. Click **Upload Certificate**.

Note

If you don't have a certificate and want to create a Certificate Signing Request (CSR) on the load balancer, please refer to [Generating a CSR on the Load Balancer](#).

9.11.11.4. Configuring SSL Termination

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

Label	SSL-Patient_Assistance	?
Associated Virtual Service	Patient_Assistance ▼	?
Virtual Service Port	443	?
SSL Operation Mode	High Security ▼	
SSL Certificate	patient-assistance-cert ▼	?
Source IP Address		?
Enable Proxy Protocol	☒	?
Bind Proxy Protocol to L7 VIP	Patient_Assistance ▼	?

Cancel Update

- Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **Patient_Assistance**.

Note

Once the VIP is selected, the *Label* field will be auto-populated with **SSL-Patient_Assistance**. This can be changed if preferred.

- Ensure that the *Virtual Service Port* is set to **443**.
 - Leave *SSL Operation Mode* set to **High Security**.
 - Select the required *SSL Certificate*.
- Leave all other settings at their default value.
 - Click **Update**.

9.11.12. VIP 3 - R2I

9.11.12.1. Create the Custom Health Check

- Using the WebUI, navigate to *Cluster Configuration > Health Check Scripts* and click **Add New Health Check**.
- Enter the following details:

Health Check Details		
Name:	R21-HealthCheck	?
Type:	Virtual Service ▼	?
Template:	Example ▼	?

- Specify an appropriate *Label* for the health check, e.g. **R21-HealthCheck**.
- Ensure that *Type* is to **Virtual Service**.

- Leave *Template* set to **Example**.
- Delete all existing text in the script edit window.
- Copy the following text into the script edit window:

```
#!/bin/bash

PATH="/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin"
# Command Line Parameters
# VIRTUAL_IP=${1}
# VIRTUAL_PORT=${2}
REAL_IP=${3}
REAL_PORT="5050"

RESULT=$(echo $'\v'\<PerformSelfTest\>$'\x1c'$'\r' | nc -vw 3 ${REAL_IP} ${REAL_PORT})

if echo ${RESULT} | grep 'WS_CONNECTION valid="Y"' ;
then
    exit 0
else
    exit 1
fi
```

- Click **Update**.

9.11.12.2. Virtual Service (VIP) Configuration

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

Virtual Service		
Label	R2I	
IP Address	10.224.94.102	
Ports	*	
Protocol		
Protocol	TCP/UDP	
Forwarding		
Forwarding Method	Direct Routing	

Cancel
Update

- Specify an appropriate *Label* for the Virtual Service, e.g. **R2I**.
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.224.94.102**.

- Set the *Ports* field to * (i.e. an asterisk - this means all ports).
 - Set the *Protocol* to **TCP/UDP**.
 - Leave the *Forwarding Method* set to **Direct Routing**.
3. Click **Update** to create the Virtual Service.
 4. Now click **Modify** next to the newly created VIP.
 5. Scroll to the *Persistence* section.
 - Set the *Persistence Mode* to **None**.
 6. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **External Script**.
 - Set the *Check Script* to the health check created above, e.g. **R2I-HealthCheck**.
 7. Leave all other settings at their default value.
 8. Click **Update**.

9.11.12.3. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?

Cancel
Update

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.13. VIP 4 - EIS

9.11.13.1. Virtual Service (VIP) Configuration

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

Virtual Service		
Label	EIS	
IP Address	10.224.94.104	
Ports	*	
Protocol		
Protocol	TCP/UDP	
Forwarding		
Forwarding Method	Direct Routing	

Cancel
Update

- Specify an appropriate *Label* for the Virtual Service, e.g. **EIS**.
 - Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.224.94.104**.
 - Set the *Ports* field to ***** (i.e. an asterisk - this means all ports).
 - Set the *Protocol* to **TCP/UDP**.
 - Leave the *Forwarding Method* set to **Direct Routing**.
3. Click **Update** to create the Virtual Service.
 4. Now click **Modify** next to the newly created VIP.
 5. Scroll to the *Health Checks* section.
 - Set the *Check Type* to **Negotiate**.
 - Set the *Check Port* to **4444**.
 - Set the *Protocol* to **HTTP**.
 - Set the *Request to Send* to **HealthCheck**.
 - Set the *Response Expected* (case sensitive field) to **ok**.
 6. Leave all other settings at their default value.
 7. Click **Update**.

9.11.13.2. Define the Associated 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 VIP.
2. Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?

Cancel
Update

- Specify an appropriate *Label* for the RIP, e.g. **WS2019DC2**.
 - Set the *Real Server IP Address* field to the required IP address, e.g. **10.224.94.42**.
3. Leave all other settings at their default value.
 4. Click **Update**.
 5. Repeat these steps to add additional Real Server(s).

9.11.14. VIP 5 - External_SMS

9.11.14.1. Virtual Service (VIP) Configuration

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

Virtual Service		[Advanced +]
Label	External_SMS	?
IP Address	10.224.94.103	?
Ports	80	?
Protocol		
Layer 7 Protocol	HTTP Mode	?

Cancel
Update

- Specify an appropriate *Label* for the Virtual Service, e.g. **External_SMS**.
 - Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.224.94.103**.
 - Set the *Ports* field to **80**.
 - Set the *Layer 7 Protocol* to **HTTP Mode**.
3. Click **Update** to create the Virtual Service.

- Now click **Modify** next to the newly created VIP.
- Scroll to the **ACL Rules** section, and click the **Add Rule** button.
- Set the **Type** to **Freetype** and copy and paste the following ACL rules into the **Freetype** field ensuring that the "#" is removed.

```
acl ec path,url_dec -m reg -i ^/ec.?
http-request reject unless ec
```

- Click **OK** to save and close the ACL.
- Scroll to the **SSL** section.
 - Enable (check) the **Enable Backend Encryption** checkbox.
- Leave all other settings at their default value.
- Click **Update**.

9.11.14.2. Define the Associated 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 VIP.
- Enter the following details:

Label	WS2019DC2	?
Real Server IP Address	10.224.94.42	?
Real Server Port	443	?
Re-Encrypt to Backend	☒	?
Enable Redirect	☐	?
Weight	100	?

- Specify an appropriate **Label** for the RIP, e.g. **WS2019DC2**.
 - Set the **Real Server IP Address** field to the required IP address, e.g. **10.224.94.42**.
 - Set the **Real Server Port** field to **443**.
- Leave all other settings at their default value.
 - Click **Update**.
 - Repeat these steps to add additional Real Server(s).

9.11.14.3. Upload the SSL Certificate



Note

If the production certificate is not currently available, continue to the next section (configuring



SSL Termination) and leave **SSL Certificate** set to **Default Self Signed Certificate**. Once you have the production certificate, follow the steps in this section to upload the certificate and then navigate to **Cluster Configuration > SSL Termination**, click **Modify** next to the SSL termination, change **SSL Certificate** to the new certificate and click **Update**.

1. Using the WebUI, navigate to **Cluster Configuration > SSL Certificate** and click **Add a new SSL Certificate**.
2. Select the option **Upload prepared PEM/PFX file**.
3. Enter the following details:

I would like to:

- ☒ Upload prepared PEM/PFX file
- ☐ Create a new SSL Certificate Signing Request (CSR)
- ☐ Create a new Self-Signed SSL Certificate.

Label: external_sms-cert

File to upload: Choose File external_sms-cert.pem

Upload Certificate

- Specify an appropriate **Label**, e.g. **external_sms-cert**.
 - Click **Choose File**.
 - Browse to and select the relevant PEM or PFX file.
 - For PFX files specify the password if required.
4. Click **Upload Certificate**.

Note

If you don't have a certificate and want to create a Certificate Signing Request (CSR) on the load balancer, please refer to [Generating a CSR on the Load Balancer](#).

9.11.14.4. Configuring SSL Termination

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

Label	SSL-External_SMS	?
Associated Virtual Service	External_SMS	?
Virtual Service Port	443	?
SSL Operation Mode	High Security	
SSL Certificate	external_sms-cert	?
Source IP Address		?
Enable Proxy Protocol	☒	?
Bind Proxy Protocol to L7 VIP	External_SMS	?

Cancel
Update

- Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **External_SMS**.

Note

Once the VIP is selected, the **Label** field will be auto-populated with **SSL-External_SMS**. This can be changed if preferred.

- Ensure that the *Virtual Service Port* is set to **443**.
 - Leave *SSL Operation Mode* set to **High Security**.
 - Select the required *SSL Certificate*.
- Leave all other settings at their default value.
 - Click **Update**.

9.11.15. Finalizing the Configuration

To apply the new settings, HAProxy & STunnel 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:

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

10. Configuring RIS Components to use the Load Balancer

The RIS core application must be configured to use the VIP on the load balancer.

In the non-HLB environment, the RIS database and RIS services/applications are configured to use the server IP or machine name.

In an HLB environment, the server IP or machine name is replaced by a virtual IP.



During the HLB Reconfiguration flow a RIS environment will be repointed to use the HLB virtual IP.

The RIS system configuration details are stored in two different locations:

1. Database tables
2. Configuration files

One table must be updated to point to a HLB VIP into a RIS schema:

1. WEBSERVICE, URL field

There are several configuration files in the RIS Schema ServiceTools menu that need to be repointed to the VIP of the HLB.

Related RIS Services:


1. Application Server Philips\RisAppServerWinService.exe.config
2. Document Generator Philips\DocumentGenerationWorker.exe.config
3. RIS Agent Philips\RisAgent.exe.config
4. R2I Philips\R2Interface.exe.config

Involved RIS Components:

1. Authentication Server \AuthenticationServer\web.config
2. GUI \GUI\Configuration\r2WorkstationConfig.htm
3. GUI64 \GUI64\Configuration\r2WorkstationConfig.htm
4. EnterpriseManagerWS \EnterpriseManagerWS\web.config
5. InstantMessageWS \RISServicesWS\web.config
6. RISServicesWS \RISWebServer\web.config
7. GUIWeb \GUIWeb\globals.bundle.js (Also require update to port number and SSL status)
8. WEB Service 64 \WS\web.config

The HLB Reconfiguration script can be launched from the RIS Installer and activated on the RIS Database Server to make the required changes.





This must be performed from the RIS primary database server.

11. Testing & Verification

Note

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

The System Overview can be accessed via the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the RIS servers) and shows the state/health of each server as well as the state of each cluster as a whole. The example below shows that all RIS servers are healthy (green) and available to accept connections:

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	RIS	10.224.94.100	80	0	HTTP	Layer 7	Proxy	
	Patient_Assistan..	10.224.94.101	80	0	HTTP	Layer 7	Proxy	
	R2I	10.224.94.102	*	0	TCPUDP	Layer 4	DR	
	EIS	10.224.94.104	*	0	TCPUDP	Layer 4	DR	
	External_SMS	10.224.94.103	80	0	HTTP	Layer 7	Proxy	
	RIS_EM	-	-	0	HTTP	Layer 7	Proxy	
	RIS_IM	-	-	0	HTTP	Layer 7	Proxy	
	RIS_RISServicesW..	-	-	0	HTTP	Layer 7	Proxy	
	RIS_Auth	-	-	0	HTTP	Layer 7	Proxy	
	RIS_WS	-	-	0	HTTP	Layer 7	Proxy	
	RIS_WEB	-	-	0	HTTP	Layer 7	Proxy	
	RIS_IWEB_sms	-	-	0	HTTP	Layer 7	Proxy	
	RIS_IWEB_gbad	-	-	0	HTTP	Layer 7	Proxy	
	RIS_IWEBP_ws	-	-	0	HTTP	Layer 7	Proxy	

If one of the servers within a cluster fails its health check, that server will be colored red and the cluster will be colored yellow as shown below:

	Patient_Assistan..	10.224.94.101	80	0	HTTP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	WS2019DC2	10.224.94.42	80	100	0	Drain	Halt	
	WS2019DC4	10.224.94.38	80	100	0	Drain	Halt	

If the services are up (green) verify that clients can connect to the VIPs and access all services.

Note

Make sure that DNS points at the VIPs rather than individual servers.

Once you have completed the verification process, continue to the next section and add a Secondary appliance to

form the HA (active/passive) clustered pair.

12. 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.

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. 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	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.

12.2. 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

 **LOADBALANCER**

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41

Password for *loadbalancer* user on peer

••••••••••

Add new node

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

Create a Clustered Pair

 **LOADBALANCER**

Primary

IP: 192.168.110.40

Attempting to pair..

 **LOADBALANCER**

Secondary

IP: 192.168.110.41

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41

Password for *loadbalancer* user on peer

••••••••••

configuring

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

High Availability Configuration - primary

 **LOADBALANCER**



Primary

IP: 192.168.110.40

 **LOADBALANCER**



Secondary

IP: 192.168.110.41

Break Clustered Pair

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](#).

13. Optional Appliance Configuration

13.1. SNMP Configuration

The appliance supports SNMP v1, v2 and v3.

To configure SNMP:

1. Using the WebUI, navigate to: *Local Configuration > SNMP Configuration*.



Protocol Versions		
Enable SNMP v1 and v2	☐	?
Enable SNMP v3	☐	?
Details		
SNMP location	Unknown	?
SNMP contact	IT Dept	?
Authentication		
SNMP v1/v2 community string	public	?
USM Username		?
USM Authorization Algorithm	SHA	?
USM Authorization Passphrase		?
USM Privacy Algorithm	AES	?
USM Privacy Passphrase		?

2. Enable the required SNMP version(s).
3. Enter the required **SNMP location** and **SNMP contact**.
4. For SNMP v1 & v2:
 - Enter the required **SNMP v1/v2 community string**.
5. For SNMP v3:
 - Specify the **USM Username**.
 - Select the required **USM Authorization Algorithm**.
 - Specify the **USM Authorization Passphrase**, it should be at least 8 characters.
 - Select the required **USM Privacy Algorithm**.
 - Specify **USM Privacy Passphrase**, it should be at least 8 characters.
6. Click **Update**.
7. Restart SNMPD using the **Restart SNMPD** button at the top of the screen.

Note

Valid characters for the **Community string**, **USM Username**, **USM Authorization Passphrase** and **USM Privacy Passphrase** fields are: **a-z A-Z 0-9 [] # ~ _ * ! = - \$ % ? { } @ : ; ^**

Note

For more information about the various OIDs and associated MIBs supported by the appliance, please refer to [SNMP Reporting](#).



Note

If you need to change the port, IP address or protocol that SNMP listens on, please refer to [Service Socket Addresses](#).

13.2. Configuring Email Alerts for Virtual Services

Email alerts can be configured for layer 4 and layer 7 Virtual Services. This enables emails to be sent when one or more of the associated Real Servers fail their health check and also when they subsequently start to pass their health check.

13.2.1. Layer 4

For layer 4 Virtual Services, settings can be configured globally for all VIPs or individually per VIP.

13.2.1.1. Global Layer 4 Email Settings

Once configured, these settings apply to all layer 4 VIPs by default.

To configure global email alert settings for layer 4 services:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 Advanced Configuration*.

Email Alert Source Address	lb1@loadbalancer.org	
Email Alert Destination Address	alerts@loadbalancer.org	
Auto-NAT	off	
Multi-threaded	yes	
		Update

2. Enter an appropriate email address in the *Email Alert Source Address* field.

e.g. lb1@loadbalancer.org

3. Enter an appropriate email address in the *Email Alert Destination Address* field.

e.g. alerts@loadbalancer.org

4. Click **Update**.

13.2.1.2. VIP Level Settings



Note

VIP level settings override the global settings.

Once configured, these settings apply to the individual VIP.

To configure VIP level email alerts:



1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 Virtual Service* and click **Modify** next to the VIP to be configured.
2. Scroll down to the *Fallback Server* section.

Email Alert Destination Address

alerts@loadbalancer.org



Cancel

Update

3. Enter an appropriate email address in the *Email Alert Destination Address* field.

e.g. alerts@loadbalancer.org

4. Click **Update**.

Note

You can set the *Email Alert Source Address* field as explained above if required to configure a default source address.

13.2.2. Layer 7

For layer 7 services, email settings are configured globally for all VIPs.

To configure global email alert settings for layer 7 services:

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

eMail Alert From

lb1@loadbalancer.org



eMail Alert To

alerts@loadbalancer.org



eMail Server Address

mail.loadbalancer.org



eMail Server Port

25



2. Enter an appropriate email address in the *eMail Alert From* field.

e.g. lb1@loadbalancer.org

3. Enter an appropriate email address in the *eMail Alert To* field.

e.g. alerts@loadbalancer.org

4. Enter an appropriate IP address/FQDN in the *eMail Server Address* field.

e.g. mail.loadbalancer.org



5. Enter an appropriate port in the **eMail Server Port** field.

e.g. 25

6. Click **Update**.

13.3. Configuring Email Alerts for Heartbeat

Email alerts can be setup for heartbeat once a clustered pair has been configured. This enables alerts to be sent when the primary/secondary communication state has changed. This can occur when the secondary appliance takes over from the primary, when the primary takes over from the secondary and also when there is a communication issue between the 2 appliances.

To configure email alert settings for Heartbeat:

1. Using the WebUI, navigate to: **Cluster Configuration > Heartbeat Configuration**.
2. Scroll down to the **Email Alerts** section.

Email Alerts

Email Alert Destination Address

alerts@loadbalancer.org



Email Alert Source Address

lb1@loadbalancer.org



3. Enter an appropriate email address in the **Email Alert Destination Address** field.
4. Enter an appropriate email address in the **Email Alert Source Address** field.
5. Click **Modify Heartbeat Configuration**.

13.4. Configuring a Smart Host (SMTP relay)

For Heartbeat (and layer 4 services), email alerts are sent from the load balancer directly to the mail server defined in the destination domain's DNS MX record by default. Alternatively, a custom smart host (mail relay server) can be specified. A smart host is an email server through which approved devices can send emails. Where possible, we recommend that you use a smart host for email alerts as this often helps improve the deliverability of emails.

To configure a Smart Host:

1. Using the WebUI, navigate to: **Local Configuration > Physical - Advanced Configuration**.
2. Scroll down to the **SMTP Relay** section.
3. Specify the FQDN or IP address of the **Smart Host**.
4. Click **Update**.

Note

By default the **Smart Host** is set as the destination email domain's DNS MX record when the **Email Alert Destination Address** is configured. It must either be left at its default setting or a



custom smart host must be configured to enable email alerts to be sent.

14. Technical Support

If you require any assistance please contact support@loadbalancer.org.

15. Further Documentation

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



16. Appendix

16.1. Generating a CSR on the Load Balancer

If you have an SSL certificate in either PFX or PEM file format, this can be uploaded to the Load balancer using the certificate upload option. Alternatively, you can create a Certificate Signing Request (CSR) on the load balancer and send this to your chosen CA to create a new certificate.

To generate a CSR:

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Certificates*.
2. Click **Add a new SSL Certificate** & select *Create a New SSL Certificate (CSR)*.

Add a new SSL Certificate

I would like to:

☐ Upload prepared PEM/PFX file

☒ Create a new SSL Certificate Signing Request (CSR) ?

☐ Create a new Self-Signed SSL Certificate.

Label

ris-cert ?

Domain (CN)

domain.com ?

Subject Alternate Name

domain.com ?

Organisation (O)

Company Name ?

Organisation unit (OU)

IT Department ?

City (L)

London ?

State or Province (ST)

Greater London ?

Country code (C)

United Kingdom ?

Email address

support@domain.com ?

CSR Key Length

2048 bits ?

Create

3. Enter a suitable *Label* (name) for the certificate.
4. Populate the remaining fields according to your requirements.

 **Note** To specify multiple SANs, separate each name with a comma.

5. Once all fields are complete click **Create**.
6. To view the CSR click **Modify** next to the new certificate, then expand the Certificate Signing Request (CSR) section.
7. Copy the CSR and send this to your chosen CA.
8. Once received, copy/paste your signed certificate into the **Your Certificate** section.
9. Intermediate and root certificates can be copied/pasted into the **Intermediate Certificate** and **Root Certificate** sections as required.
10. Click **Update** to complete the process.

The new certificate will now be displayed under **Cluster Configuration > SSL Certificates**.



17. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0	25 June 2024	Initial version		RJC
1.1	8 August 2024	changed all occurrences of "RIS_RISServicesWS" to "RIS_ServicesWS"	Spelling correction	RJC
1.2	21 February 2025	Updated ACL rules for VIP 1 - RIS and VIP 2 - Patient_Assistance	Technical requirement	RJC
1.3	3 June 2025	Added VMware and Hyper-V configuration steps to cater for Virtual Appliance deployments Added section to explain appliance licensing steps	Technical requirements	RJC





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.

