

Load Balancing IBM Watson Health MergePACS by Merative

Version 1.4.0



Table of Contents

1. About this Guide	3
2. Loadbalancer.org Appliances Supported.	3
3. Software Versions Supported	
3.1. Loadbalancer.org Appliance	3
3.2. IBM Watson Health MergePACS by Merative	3
4. Load Balancing MergePACS	3
4.1. Port Requirements	3
4.2. Deployment Concept	4
4.3. Virtual Service (VIP) Requirements	4
4.4. Deployment Mode	5
5. Loadbalancer.org Appliance – the Basics	5
5.1. Virtual Appliance	
5.2. Initial Network Configuration	5
5.3. Accessing the Appliance WebUI	5
Main Menu Options	7
5.4. Appliance Software Update	8
Determining the Current Software Version	8
Checking for Updates using Online Update	8
Using Offline Update	8
5.5. Ports Used by the Appliance	9
6. Appliance & MergePACS Configuration	
6.1. Appliance Configuration	10
Configuring VIP1 – All PACS Services	10
6.2. MergePACS Server Configuration	12
Windows Server 2012 & Later	12
7. Testing & Verification	
7.1. Automatic Failover	18
7.2. Manual Failover.	
7.3. Client Connection Tests	18
8. Technical Support	
9. Additional Documentation	18
10. Appendix	
10.1. Configuring HA - Adding a Secondary Appliance	
Non-Replicated Settings	
Adding a Secondary Appliance - Create an HA Clustered Pair	20
11. Document Revision History	22

1. About this Guide

This guide details the steps required to configure a highly available IBM Watson Health MergePACS by Merative environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any MergePACS configuration changes that are required.

For more information about initial appliance deployment, network configuration and using the Web User Interface (WebUI), please also refer to the Administration Manual.

2. Loadbalancer.org Appliances Supported

All our products can be used with the IBM Watson Health MergePACS by Merative environment. For full specifications of available models please refer to https://www.loadbalancer.org/products. Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

V8.3.8 and later

8 Note

The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If using an older software version, note that the screenshots presented here may not match the WebUI exactly.

3.2. IBM Watson Health MergePACS by Merative

All versions

4. Load Balancing MergePACS

For high availability, IBM Watson Health recommend that a load balancer is used to enable rapid failover to the secondary MergePACS Cluster should the primary cluster become unavailable.

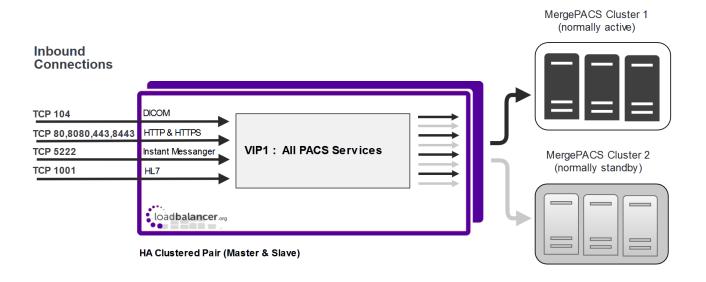
4.1. Port Requirements

The following table shows the ports used by MergePACS. The load balancer must be configured to listen on the same ports.

Port	Protocols	Use
104	TCP	DICOM
80,8080,443,8443	TCP	HTTP & HTTPS
5222	TCP	Instant Messenger
1001	TCP	HL7

4.2. Deployment Concept

When MergePACS is deployed with the load balancer, clients connect to the Virtual Service (VIP) on the load balancer rather than connecting directly to one of the MergePACS Clusters. Under normal conditions, these connections are then forwarded to the Primary Cluster.



8 Note

The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to Configuring HA - Adding a Secondary Appliance in the appendix for more details on configuring a clustered pair.

Should the Primary Cluster become unavailable, failover to the Secondary Cluster can be handled in either of the following ways:

- **Automatically** In this case, health checks are configured at 30 second intervals. Should there be 10 consecutive health check failures, failover to the Secondary Cluster occurs.
- **Manually** In this case, failover to the Secondary Cluster must be triggered manually using the 'Halt' feature in the load balancer's WebUI. Please refer to Manual Failover for more details.

8 Note

The way the Virtual Service's health check is configured determines which of these failover methods is used.

4.3. Virtual Service (VIP) Requirements

A single multi-port VIP is used that listens on all required ports. The VIP is configured as follows:

- Deployment mode: Layer 4 DR (Direct Return) mode
- Listens on a total of 7 ports as described on the table and diagram in Port Requirements
- The health-check configuration depends on whether automatic or manual failover is required:
 - for *automatic* failover an external script is used, the script checks that *all* 7 ports are available and runs every 30 seconds, if connection to one or more of the ports fails, the health check is deemed to have failed, if there are 10 consecutive health check failures, cluster failover occurs

- for manual failover the health check is set to: No checks, always On
- The associated Real Server is configured to be the cluster IP address of the Primary Cluster
- The fallback server is configured to be the cluster IP address of the Secondary Cluster

4.4. Deployment Mode

As mentioned above, the VIP is configured using Layer 4 DR (Direct Return) mode. This mode offers the best possible performance since replies go directly from the MergePACS Cluster to the client, and not via the load balancer. To use this mode, the "ARP Problem" must be solved on each MergePACS server as explained in MergePACS Server Configuration.

5. Loadbalancer.org Appliance – the Basics

5.1. Virtual Appliance

A fully featured, fully supported 30 day trial is available if you are conducting a PoC (Proof of Concept) deployment. The VA is currently available for VMware, Virtual Box, Hyper-V, KVM, XEN and Nutanix AHV and has been optimized for each Hypervisor. By default, the VA is allocated 2 vCPUs, 4GB of RAM and has a 20GB virtual disk. The Virtual Appliance can be downloaded here.

8 Note	The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.
8 Note	Please refer to Virtual Appliance Installation and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.
8 Note	The VA has 4 network adapters. For VMware only the first adapter (eth0) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

5.2. Initial Network Configuration

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

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

5.3. Accessing the Appliance WebUI

The WebUI is accessed using a web browser. By default, users are authenticated using Apache authentication. Users can also be authenticated against LDAP, LDAPS, Active Directory or Radius - for more information, please refer to External Authentication.

Note There are certain differences when accessing the WebUI for the cloud appliances. For details,

please refer to the relevant Quick Start / Configuration Guide.

8 Note

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

1. Using a browser, navigate to the following URL:

https://<IP-address-configured-during-the-network-setup-wizard>:9443/lbadmin/

8 Note

You'll receive a warning about the WebUl's certificate. This is due to the default self signed certificate that is used. If preferred, you can upload your own certificate - for more information, please refer to Appliance Security Features.

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

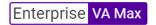
Username: loadbalancer

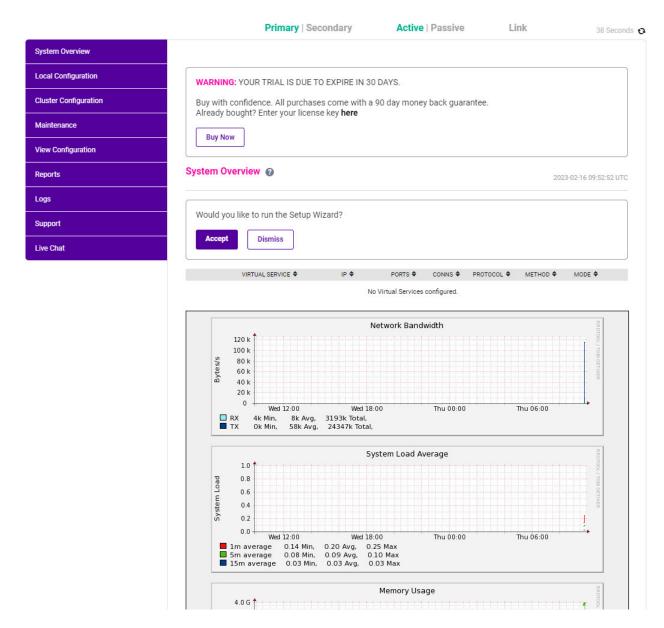
Password: <configured-during-network-setup-wizard>

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

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

IL LOADBALANCER





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

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

Main Menu Options

System Overview - Displays a graphical summary of all VIPs, RIPs and key appliance statistics

Local Configuration - Configure local host settings such as IP address, DNS, system time etc.

Cluster Configuration - Configure load balanced services such as VIPs & RIPs

Maintenance - Perform maintenance tasks such as service restarts and taking backups

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

Support - Create a support download, contact the support team & access useful links



5.4. Appliance Software Update

To ensure that the appliance(s) are running the latest software version, we recommend a software update check is performed.

Determining the Current Software Version

The software version is displayed at the bottom of the WebUI as shown in the example below:

Copyright © Loadbalancer.org Inc. 2002 - 2023 ENTERPRISE VA Max - v8.9.0



Checking for Updates using Online Update

8 Note By default, the appliance periodically contacts the Loadbalancer.org update server and checks for updates. An update check can also be manually triggered as detailed below.

- 1. Using the WebUI, navigate to: *Maintenance > Software Update*.
- 2. Select Online Update.
- 3. If the latest version is already installed, a message similar to the following will be displayed:

Information: Version v8.9.0 is the current release. No updates are available

- 4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
- 5. Click **Online Update** to start the update process.
 - Note Do not navigate away whilst the update is ongoing, this may cause the update to fail.
- 6. Once complete (the update can take several minutes depending on download speed and upgrade version) the following message will be displayed:

Information: Update completed successfully.

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

Using Offline Update

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



8 Note

Please contact support@loadbalancer.org to check if an update is available and obtain the latest offline update files.

To perform an offline update:

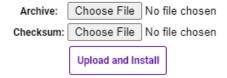
- 1. Using the WebUI, navigate to: *Maintenance > Software Update*.
- 2. Select Offline Update.
- 3. The following screen will be displayed:

Software Update

Offline Update

The following steps will lead you through offline update.

- Contact Loadbalancer.org support to obtain the offline update archive and checksum.
- 2. Save the archive and checksum to your local machine.
- Select the archive and checksum files in the upload form below.
- 4. Click Upload and Install to begin the update process.



- 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.

5.5. Ports Used by the Appliance

By default, the appliance uses the following TCP & UDP ports:

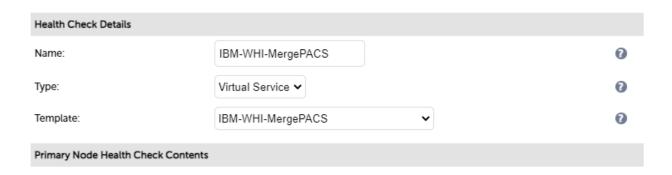
Protocol	Port	Purpose
TCP	22	SSH
TCP & UDP	53	DNS
TCP & UDP	123	NTP
TCP & UDP	161	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback page
TCP	9443	WebUI - HTTPS

6. Appliance & MergePACS Configuration

6.1. Appliance Configuration

Configuring VIP1 - All PACS Services

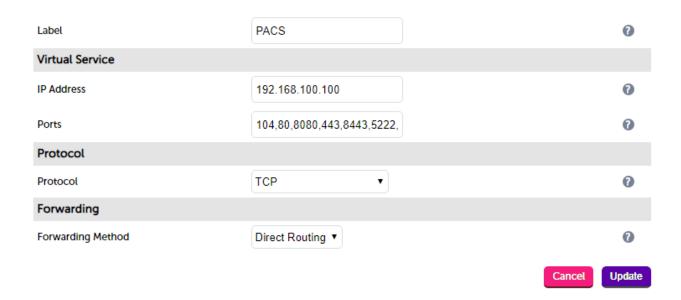
- a) Configure the External Health Check Script (used for automatic failover)
 - 1. Using the WebUI, navigate to Cluster Configuration > Health Check Scripts and click Add New Health Check.



- 2. Specify an appropriate *Name* for the health check, e.g. **IBM-WHI-MergePACS**.
- 3. Set *Type* to **Virtual Service**.
- 4. Set Template to IBM-WHI-MergePACS.
- 5. Click Update.

b) Setting up the Virtual Service (VIP)

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



3. Enter an appropriate label (name) for the VIP, e.g. PACS.

- 4. Set the Virtual Service IP address field to the required IP address, e.g. 192.168.100.100.
- 5. Set the Virtual Service Ports field to 104,80,8080,443,8443,5222,1001.
- 6. Leave Protocol set to TCP.
- 7. Leave Forwarding Method set to Direct Routing.
- 8. Click Update.
- 9. Now click **Modify** next to the newly created VIP.
- 10. Scroll to the Health Checks section.

For **automatic** failover:

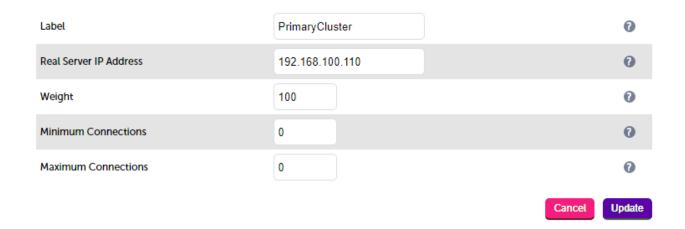
- a. Set Check Type to External Script.
- b. Set External Script to IBM-WHI-MergePACS.

For **manual** failover:

- a. Set the Check Type to No checks, Always On.
- 11. Scroll to the Fallback Server section.
 - a. Set the IP Address to the IP address of the Secondary MergePACS Cluster.
 - b. Set the *Port* to **0** (numerical zero), this ensures that the fallback server (i.e. the Secondary Cluster) can receive connections on all required ports.
- 12. Enable (check) the MASQ Fallback checkbox.
- 13. Click Update.

c) Setting up the Real Servers (RIPs)

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



- 3. Enter an appropriate label (name) for the RIP, e.g. PrimaryCluster.
- 4. Set the Real Server IP Address field to the IP address of the Primary MergePACS Cluster.
- 5. Click Update.

6.2. MergePACS Server Configuration

As mentioned in Deployment Mode, when using Layer 4 DR mode, the ARP problem must be solved. This involves configuring each MergePACS Server to be able to receive traffic destined for the VIP, and ensuring that each Server does not respond to ARP requests for the VIP address – only the load balancer should do this.

8 Note

The following steps must be performed on all MergePACS Servers.

Windows Server 2012 & Later

Windows Server 2012 and later support 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.

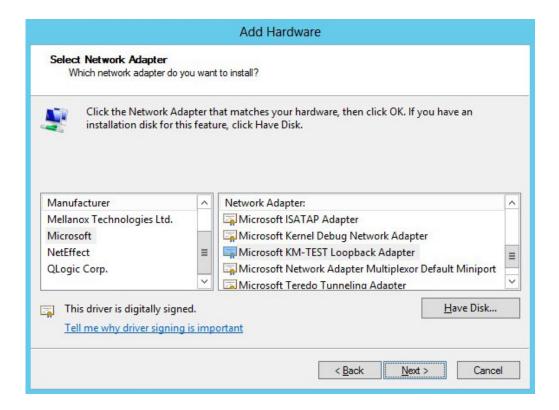
In addition, steps must be taken to set the strong/weak host behavior on each Real Server. This is used to either prevent or allow interfaces to receive packets destined for a different interface on the same server.

(!) Important

The following 3 steps must be completed on **all** Real Servers associated with the VIP.

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.



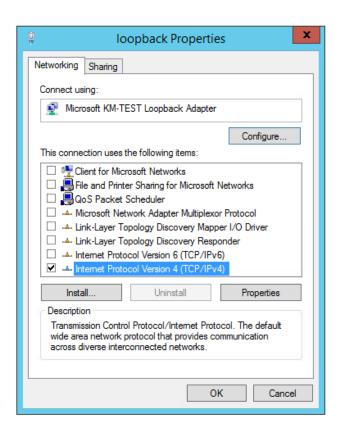
- 5. Select Microsoft & Microsoft KM-Test Loopback Adapter, click Next.
- 6. Click **Next** to start the installation, when complete click **Finish**.

Step 2 of 3: Configure the Loopback Adapter

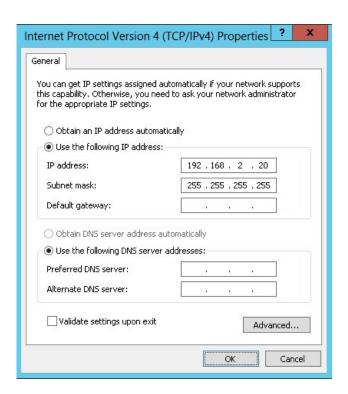
- 1. Open Control Panel and click Network and Sharing Center.
- 2. Click Change adapter settings.
- 3. Right-click the new Loopback Adapter and select Properties.
- Note You can configure IPv4 or IPv6 addresses or both depending on your requirements.

IPv4 Addresses

1. Uncheck all items except Internet Protocol Version 4 (TCP/IPv4) as shown below:



2. 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:

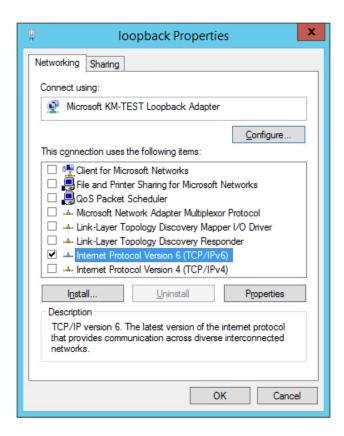


- Note 192.168.2.20 is an example, make sure you specify the correct VIP address.
- Note If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be

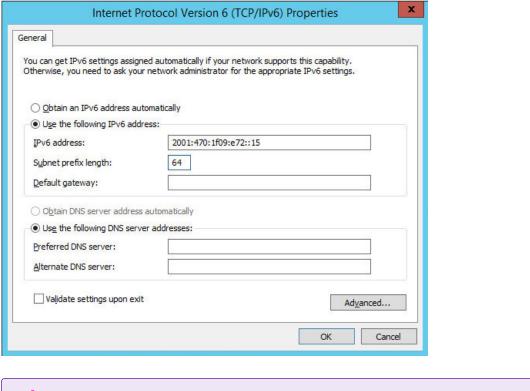
3. Click **OK** then click **Close** to save and apply the new settings.

IPv6 Addresses

1. Uncheck all items except Internet Protocol Version 6 (TCP/IPv6) as shown below:



2. Ensure that Internet Protocol Version (TCP/IPv6) is selected, click Properties and configure the IP address to be the same as the Virtual Service (VIP) and set the Subnet Prefix Length to be the same as your network setting, e.g. 2001:470:1f09:e72::15/64 as shown below:



Note 2001:470:1f09:e72::15/64 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.

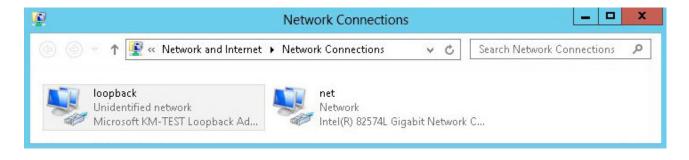
3. Click **OK** then click **Close** to save and apply the new settings.

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:



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:

For IPv4 addresses:

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

For IPv6 addresses:

```
netsh interface ipv6 set interface "net" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostsend=enabled
netsh interface ipv6 set interface "loopback" dadtransmits=0
```

Option 2 - Using PowerShell Cmdlets

For IPv4 addresses:

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

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

For IPv6 Addresses:

```
{\tt Set-NetIpInterface - InterfaceAlias \ loopback - WeakHostReceive \ enabled - WeakHostSend \ enabled - DadTransmits \ 0 \ - AddressFamily \ IPv6}
```

Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled -AddressFamily IPv6

7. Testing & Verification

A Note

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

Under normal circumstances the Primary Cluster handles all connections. Failover to the Secondary Cluster is handled automatically or manually depending on how the VIP is configured (see Virtual Service (VIP) Requirements).



7.1. Automatic Failover

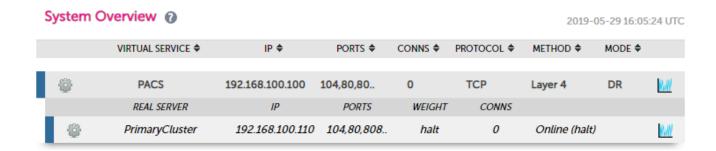
Automatic failover occurs after 5 minutes. To trigger a failover, the Primary Cluster must be continuously unavailable for this time.

7.2. Manual Failover

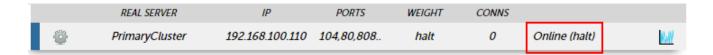
To trigger a failover to the Secondary Cluster, the 'Halt' option in the System Overview is used:



Once Halted, the VIP & RIP will be shown colored blue, connections will then be forwarded to the fallback server, I.e the Secondary Cluster:



To return to the Primary Cluster, the 'Online' option is used:



7.3. Client Connection Tests

Ensure that clients can connect via the load balancer to the MergePACS Cluster. You'll probably need to create new DNS records or modify your existing DNS records, replacing the IP addresses of individual servers or the cluster with the IP address of the Virtual Service on the load balancer.

8. Technical Support

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

9. Additional Documentation

For additional information, please refer to the Administration Manual.

10. Appendix

10.1. Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution.

We recommend that the Primary appliance is configured first and then the Secondary should be added. Once the Primary and Secondary are paired, all load balanced services configured on the Primary are automatically replicated to the Secondary over the network using SSH/SCP.

8 Note

For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

The clustered HA pair uses Heartbeat to determine the state of the other appliance. Should the active device (normally the Primary) suffer a failure, the passive device (normally the Secondary) will take over.

Non-Replicated Settings

A number of settings are not replicated as part of the Primary/Secondary pairing process and therefore must be manually configured on the Secondary appliance. These are listed by WebUI menu option in the table below:

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	All network settings including IP address(es), bonding configuration and VLANs
Local Configuration	Routing	Routing configuration including default gateways and static routes
Local Configuration	System Date & time	All time and date related settings
Local Configuration	Physical – Advanced Configuration	Various settings including Internet Proxy, Management Gateway, Firewall connection tracking table size, NIC offloading, SMTP relay, logging and Syslog Server
Local Configuration	Security	Appliance security settings
Local Configuration	SNMP Configuration	Appliance SNMP settings
Local Configuration	Graphing	Appliance graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Software Updates	Appliance software update management
Maintenance	Firewall Script	Appliance firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

(!) Important

Make sure that if these settings/updates have been configured on the Primary appliance, they're also configured on the Secondary appliance.

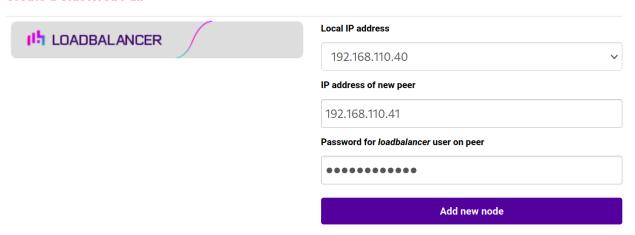
Adding a Secondary Appliance - Create an HA Clustered Pair

8 Note

If you have already run the firewall lockdown wizard on either appliance, you'll need to ensure that it is temporarily disabled on both appliances whilst performing the pairing process.

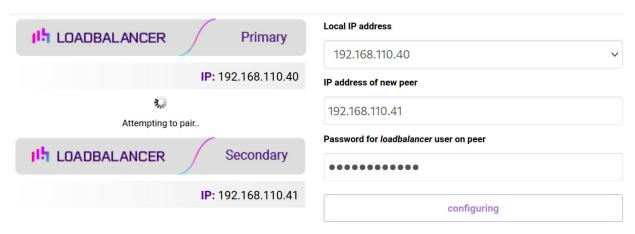
- 1. Deploy a second appliance that will be the Secondary and configure initial network settings.
- 2. Using the WebUI on the Primary appliance, navigate to: *Cluster Configuration > High-Availability Configuration*.

Create a Clustered Pair



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

Create a Clustered Pair

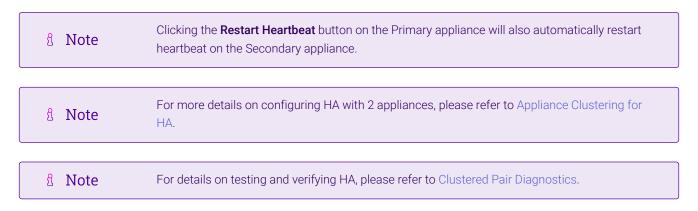


6. Once complete, the following will be displayed on the Primary appliance:

High Availability Configuration - primary



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.



11. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.1.0	2 August 2019	Styling and layout	General styling updates	RJC
1.1.1	24 August 2020	New title page	Branding update	АН
		Updated Canadian contact details	Change to Canadian contact details	
1.2.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.2.1	11 May 2022	Updated external health check related content to reflect latest software version	New software release	RJC
1.3.0	6 September 2022	Renamed document and amended references to the product	Product acquisition by Merative	AH
1.3.1	5 January 2023	Combined software version information into one section Added one level of section numbering	Housekeeping across all documentation	AH
		Added software update instructions		
		Added table of ports used by the appliance		
		Reworded 'Further Documentation' section		
1.3.2	2 February 2023	Updated screenshots	Branding update	АН
1.3.3	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.4.0	24 March 2023	New document theme Modified diagram colours	Branding update	АН



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

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.

