



Load Balancing IBM Watson Health MergePACS by Merative

Version 1.3.0

Table of Contents

1. About this Guide	3
2. Loadbalancer.org Appliances Supported	3
3. Loadbalancer.org Software Versions Supported	3
4. IBM Watson Health MergePACS by Merative Software Versions Supported	3
5. Load Balancing MergePACS	3
Port Requirements	3
Deployment Concept	3
Virtual Service (VIP) Requirements	4
Deployment Mode	4
6. Loadbalancer.org Appliance – the Basics	5
Virtual Appliance	5
Initial Network Configuration	5
Accessing the WebUI	5
Main Menu Options	6
Initial Network Configuration	7
7. Appliance & MergePACS Configuration	7
Appliance Configuration	7
Configuring VIP1 – All PACS Services	7
MergePACS Server Configuration	9
Windows Server 2012 & Later	9
8. Testing & Verification	14
Automatic Failover	14
Manual Failover	14
Client Connection Tests	15
9. Technical Support	15
10. Additional Documentation	15
11. Conclusion	15
12. Appendix	16
Configuring HA - Adding a Secondary Appliance	16
Non-Replicated Settings	16
13. Document Revision History	19

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. Loadbalancer.org Software Versions Supported

- V8.3.8 and later

Note

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

4. IBM Watson Health MergePACS by Merative Software Versions Supported

- IBM Watson Health MergePACS by Merative – all versions

5. 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.

Port Requirements

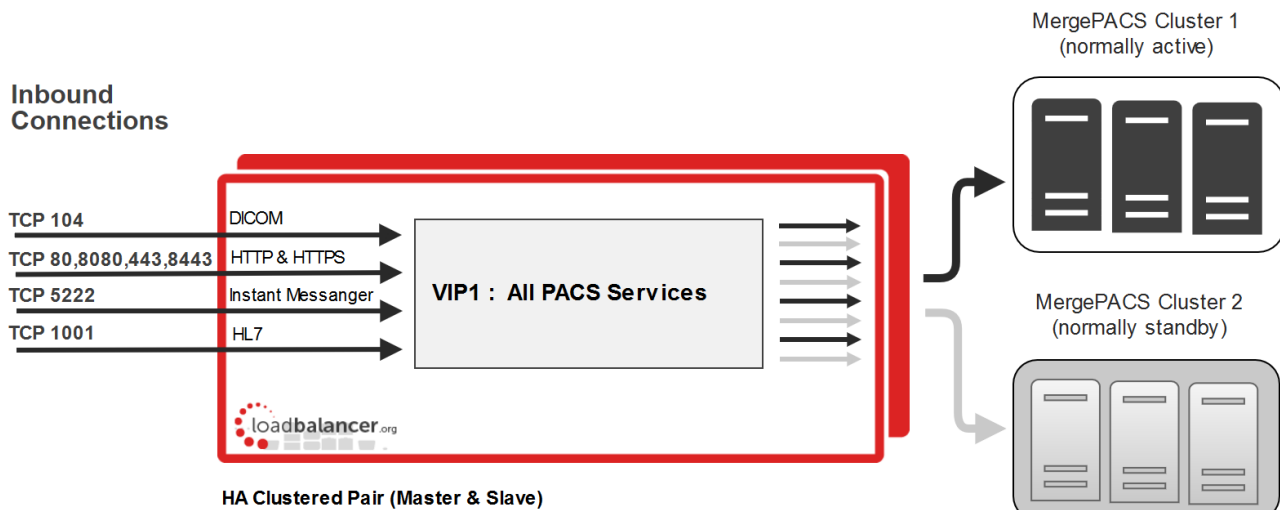
The following table shows the ports used by MergePACS. The load balancer must be configured to listen on the same ports.

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

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.



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.

Note

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

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

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](#).

6. Loadbalancer.org Appliance – the Basics

Virtual Appliance

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

Note | The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.

Note | Please refer to [Virtual Appliance Installation](#) and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.

Note | The VA has 4 network adapters. For VMware only the first adapter (eth0) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

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.

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

Accessing the WebUI

The WebUI is accessed using a web browser. By default, user authentication is based on local Apache .htaccess files. User administration tasks such as adding users and changing passwords can be performed using the WebUI menu option: *Maintenance > Passwords*.

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

Note | If required, users can also be authenticated against LDAP, LDAPS, Active Directory or Radius. For more information please refer to [External Authentication](#).

1. Using a browser, access the WebUI using the following URL:

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

2. Log in to the WebUI:

Username: loadbalancer

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

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

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

The screenshot displays the Loadbalancer.org Enterprise VA 100G WebUI. The top navigation bar includes the logo, status indicators (Primary | Secondary, Active | Passive), and a refresh button. A yellow warning banner at the top right states: "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". Below the banner, the "System Overview" section is active, showing a "Would you like to run the Setup Wizard?" dialog with "Accept" and "Dismiss" buttons. A filter bar below the dialog shows "VIRTUAL SERVICE" and "IP" filters, with a message "No Virtual Services configured." Three graphs are visible: "Network Bandwidth" showing RX and TX traffic, "System Load Average" showing 1m, 5m, and 15m averages, and "Memory Usage" showing current usage.

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

3. You'll be asked if you want to run the Setup Wizard. If you click **Accept** the Layer 7 Virtual Service configuration wizard will start. If you want to configure the appliance manually, simply click **Dismiss**.

Main Menu Options

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

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

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

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

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

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

Live Chat - Start a live chat session with one of our Support Engineers

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.

Important

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

7. Appliance & MergePACS Configuration

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**.

Health Check Details		
Name:	<input type="text" value="IBM-WHI-MergePACS"/>	?
Type:	<input type="text" value="Virtual Service"/>	?
Template:	<input type="text" value="IBM-WHI-MergePACS"/>	?

Primary Node Health Check Contents

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)

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

Label	<input type="text" value="PACS"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.100.100"/>	?
Ports	<input type="text" value="104,80,8080,443,8443,5222,"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	?

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:

Label	<input type="text" value="PrimaryCluster"/>	?
Real Server IP Address	<input type="text" value="192.168.100.110"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

3. Enter an appropriate 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**.

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.

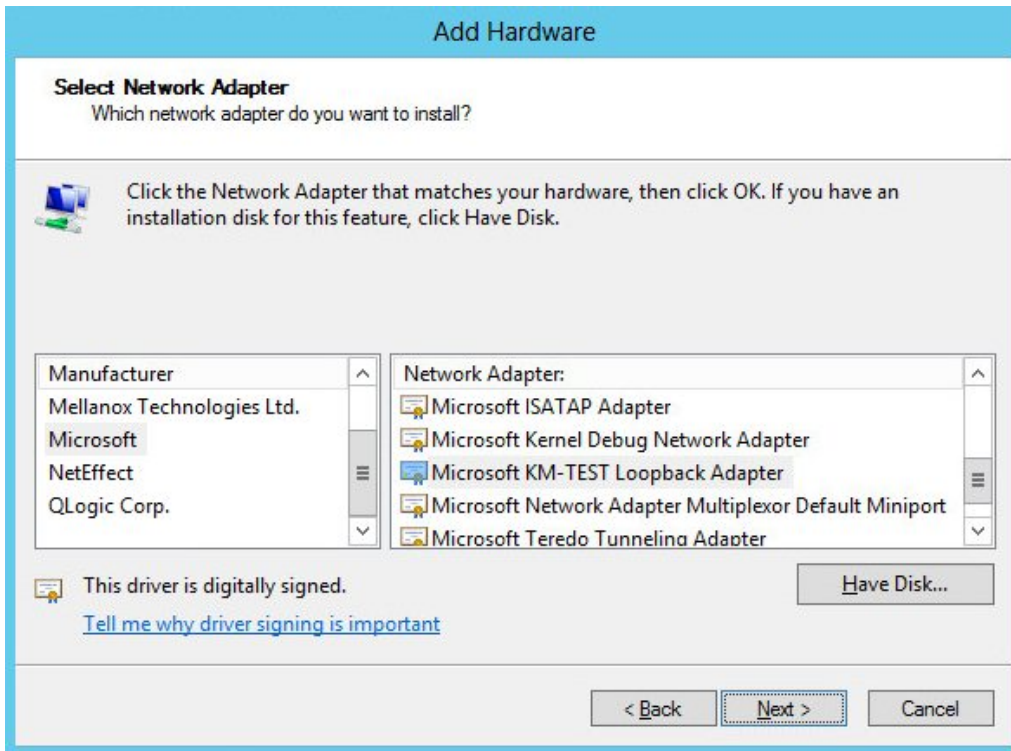
Note | The following steps must be performed on all MergePACS Servers.

Windows Server 2012 & Later

Windows Server 2012 and later support Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter. The IP address allocated to the Loopback Adapter must be the same as the Virtual Service (VIP) address. If the Real Server is included in multiple DR mode VIPs, additional IP addresses can be added to the Loopback Adapter that correspond to each VIP. In addition, steps must be taken to set the strong/weak host behavior which is used to either block or allow interfaces to receive packets destined for a different interface on the same server.

Step 1 of 3: Install the Microsoft Loopback Adapter

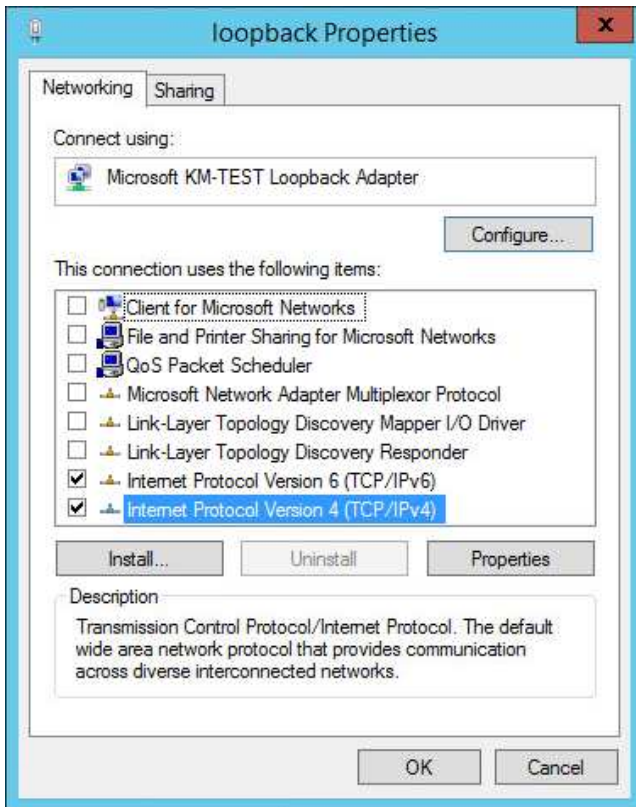
1. Click **Start**, then run `hdwwiz` to start the Hardware Installation Wizard.
2. When the Wizard has started, click **Next**.
3. Select **Install the hardware that I manually select from a list (Advanced)**, click **Next**.
4. Select **Network adapters**, click **Next**.
5. Select **Microsoft & Microsoft 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**.
4. Uncheck all items except **Internet Protocol Version 4 (TCP/IPv4)** and **Internet Protocol Version 6 (TCP/IPv6)** as shown below:



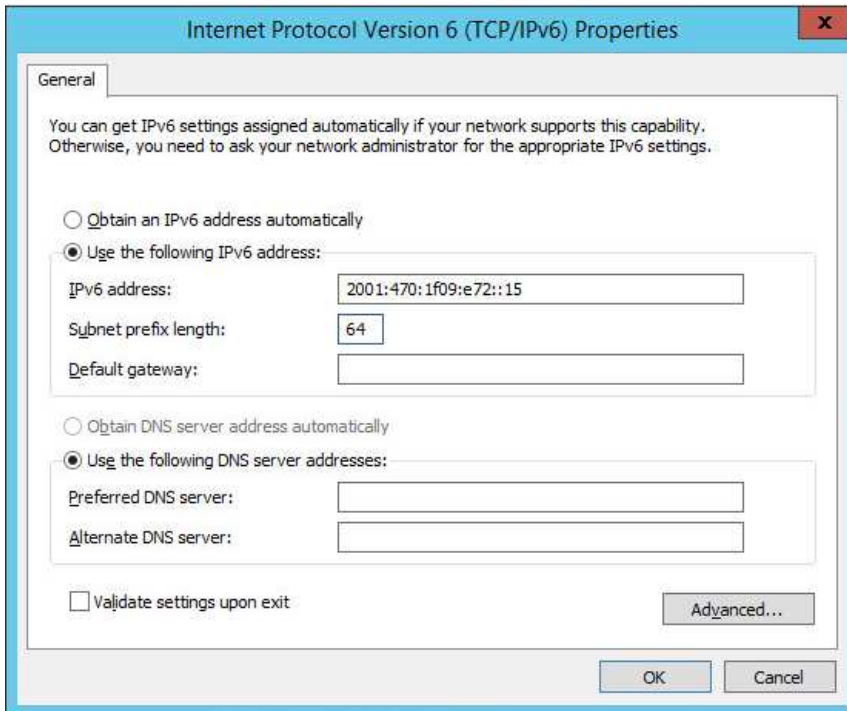
Note

Leaving both checked ensures that both IPv4 and IPv6 are supported. Select one if preferred.

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



6. If configuring IPv6 addresses select **Internet Protocol Version (TCP/IPv6)**, 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:



7. Click **OK** on TCP/IP Properties, then click **Close** on Ethernet Properties to save and apply the new settings.

Note

For Windows 2012/2016/2019, it's not necessary to modify the interface metric on the advanced tab and should be left set to Automatic.

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

To configure the correct strong/weak host behavior for Windows 2012/2016/2019, the following commands must be run on each Real Server:

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 these commands to work, the LAN connection NIC must be named "net" and the loopback NIC must be named "loopback" as shown below. If you prefer to leave your current NIC names, then the commands above must be modified accordingly. For example, if your network adapters are named "LAN" and "LOOPBACK", the commands required would be:

```
netsh interface ipv4 set interface "LAN" 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
```


For these commands to work, the LAN connection NIC must be named "net" and the loopback NIC must be named "loopback" as shown below. If you prefer to leave your current NIC names, then the commands above must be modified accordingly. For example, if your network adapters are named "LAN" and "LOOPBACK", the commands required would be:

```
netsh interface ipv6 set interface "LAN" 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
```



Note | The names for the NICs are case sensitive, so make sure that the name used for the interface and the name used in the commands match exactly.

- Start PowerShell or use a command window to run the appropriate netsh commands as shown in the example below:



Note | This shows an IPv6 example, use the IPv4 commands if you're using IPv4 addresses.

Repeat steps 1 - 3 on all remaining Windows 2012/2016/2019 Real Server(s).

If preferred you can also use the following PowerShell Cmdlets:

The following example configures both IPv4 and IPv6 at the same time:

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

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

To configure just IPv4:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled
```

```
-DadTransmits 0 -AddressFamily IPv4
```

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

To configure just IPv6:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled  
-DadTransmits 0 -AddressFamily IPv6
```

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

8. Testing & Verification

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](#)).

Automatic Failover

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

Manual Failover





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

	REAL SERVER	IP	PORTS	WEIGHT	CONNS		
	PrimaryCluster	192.168.100.110	104,80,808..	100	0	Drain	Halt 



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

System Overview

2019-05-29 16:05:24 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
	PACS	192.168.100.100	104,80,80..	0	TCP	Layer 4	DR 
	REAL SERVER	IP	PORTS	WEIGHT	CONNS		
	PrimaryCluster	192.168.100.110	104,80,808..	halt	0	Online (halt)	

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

	REAL SERVER	IP	PORTS	WEIGHT	CONNS	
	PrimaryCluster	192.168.100.110	104,80,808..	halt	0	Online (halt) 

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.

9. 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.

10. Additional Documentation

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

11. Conclusion

Loadbalancer.org appliances provide a very cost effective solution for highly available load balanced IBM Watson Health MergePACS by Merative environments.

12. Appendix

Configuring HA - Adding a Secondary Appliance

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

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

Note

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

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

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.

To add a Secondary node - i.e. create a highly available 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.org

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 above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

CREATE A CLUSTERED PAIR

192.168.110.40 loadbalancer.org

Attempting to pair..

192.168.110.41 loadbalancer.org

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

	192.168.110.40	loadbalancer.org	Break Clustered Pair
	192.168.110.41	loadbalancer.org	

7. To finalize the configuration, restart heartbeat and any other services as prompted in the blue message box at the top of the screen.

Note | Clicking the **Restart Heartbeat** button on the Primary appliance will also automatically restart heartbeat on the Secondary appliance.

Note | For more details on configuring HA with 2 appliances, please refer to [Appliance Clustering for HA](#).

Note | For details on testing and verifying HA, please refer to [Clustered Pair Diagnostics](#).

13. 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 Updated Canadian contact details	Branding update Change to Canadian contact details	AH
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

About Loadbalancer.org

Loadbalancer.org's mission is to ensure that its clients' businesses are never interrupted. The load balancer experts ask the right questions to get to the heart of what matters, bringing a depth of understanding to each deployment. Experience enables Loadbalancer.org engineers to design less complex, unbreakable solutions - and to provide exceptional personalized support.



United Kingdom

Loadbalancer.org Ltd.
Compass House, North Harbour
Business Park, Portsmouth, PO6 4PS
UK: +44 (0) 330 380 1064
sales@loadbalancer.org
support@loadbalancer.org

Canada

Loadbalancer.org Appliances Ltd.
300-422 Richards Street, Vancouver,
BC, V6B 2Z4, Canada
TEL: +1 866 998 0508
sales@loadbalancer.org
support@loadbalancer.org

United States

Loadbalancer.org, Inc.
4550 Linden Hill Road, Suite 201
Wilmington, DE 19808, USA
TEL: +1 833.274.2566
sales@loadbalancer.org
support@loadbalancer.org

Germany

Loadbalancer.org GmbH
Tengstraße 2780798,
München, Germany
TEL: +49 (0)89 2000 2179
sales@loadbalancer.org
support@loadbalancer.org