LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing eCopy ShareScan

Version 1.2.0

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced eCopy ShareScan environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any eCopy ShareScan configuration changes that are required to enable load balancing.

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 eCopy ShareScan. For full specifications of available models please refer to https://www.loadbalancer.org/products/enterprise.

Some features may not be available or fully supported in all cloud platforms due to platform specific limitations. For more details, please refer to the "Main Differences to our Standard (Non-Cloud) Product" section in the appropriate cloud platform Quick Start Guide or check with Loadbalancer.org support.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.9.1 and later

> **Note** The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If you're using an older version, or the very latest, the screenshots presented here may not match your WebUI exactly.

3.2. eCopy ShareScan

- v6.2 and later

4. eCopy ShareScan

eCopy ShareScan 6.2 is an MFP document capture solution that enables MFP users to engage their business systems and processes by completely automating document capture processes. As a result, eCopy ShareScan simplifies MFP capture workflows and enables users with advanced imaging capabilities. The eCopy ShareScan software extends the capabilities of digital copiers and scanners. When installing and setting up a ShareScan system, you must be familiar with the scanning devices that you will use with ShareScan, the ShareScan software components, and the basic installation and configuration workflow.

5. Load Balancing eCopy ShareScan

> **Note** It's highly recommended that you have a working eCopy ShareScan environment first before implementing the load balancer.

5.1. The Basics

The primary function of the load balancer is to distribute inbound requests across multiple eCopy ShareScan servers. This allows administrators to configure multiple servers and easily share the load between them. Adding additional capacity as demand grows then becomes straight forward and can be achieved by simply adding additional eCopy ShareScan servers to the load balanced cluster.

5.2. Ports & Protocols

The following table shows the ports that are normally used with eCopy ShareScan:

Port	Protocol	Use
*	TCP	Testing load balancer configuration
8080	TCP/HTTP	HTTP eCopy Tomcat application
443	TCP/HTTPS	HTTPS eCopy Tomcat application
9600	TCP	Web based MFDs
9261	TCP	Embedded MFDs i.e. Ricoh and Canon

> **Note** For the complete port list necessary to configure for a particular device vendor consult ShareScan documentation (High Availability and Load Balancing Deployment Guide, v6.2.)

5.3. eCopy ShareScan Server Health-checks

Regular eCopy ShareScan server monitoring ensures that failed servers are marked as down and client requests are only directed to functional servers. Health checks can range from a simple ICMP PING to a full negotiate check where content on a certain page is read and verified. Please refer to Real Server (eCopy ShareScan) Health Checks for more details.

5.4. SSL Termination & Certificates

It is recommended that all SSL traffic is terminated on the eCopy ShareScan servers (***SSL pass-through***).

5.5. Persistence (aka Server Affinity)

Source IP persistence is required when load balancing the eCopy ShareScan application and is the only available persistence method when load balancing at layer 4.

5.6. Load Balancer Deployment

The following diagram illustrates how the load balancer is deployed with multiple eCopy ShareScan servers.

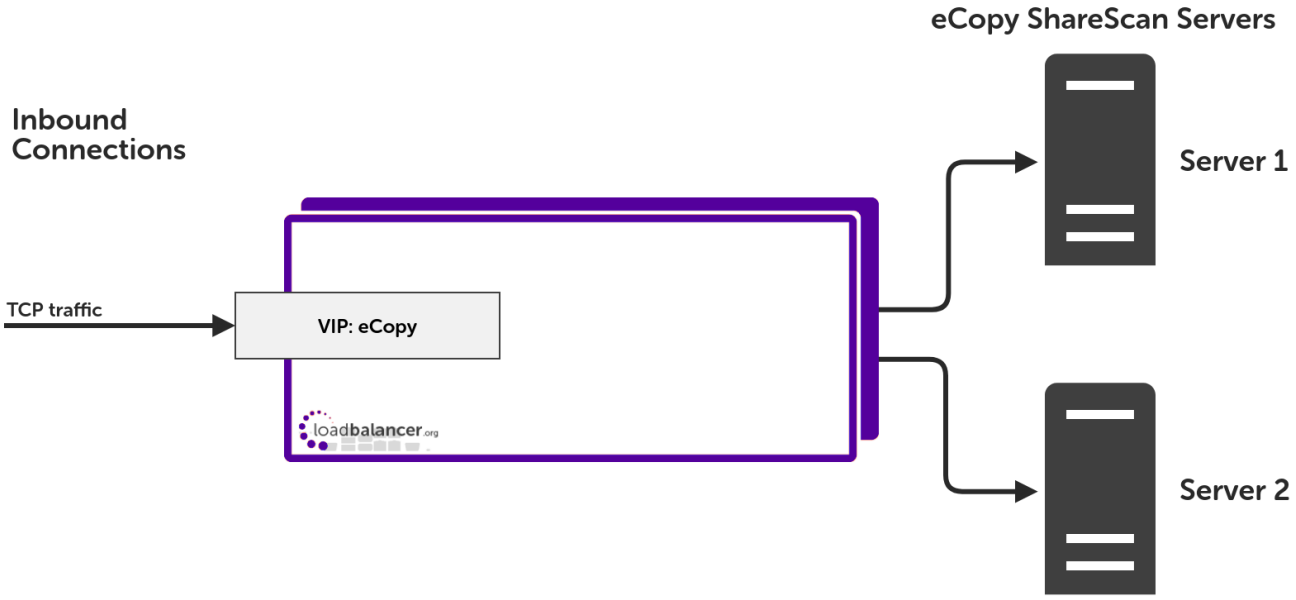

VIP = **V**irtual **IP** Address

> **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 for more details on configuring a clustered pair.

5.7. Load Balancer Deployment Modes

The load balancer can be deployed in 4 fundamental ways: ***Layer 4 DR mode***, ***Layer 4 NAT mode***, ***Layer 4 SNAT mode*** and ***Layer 7 SNAT mode***. For eCopy ShareScan, Layer 4 DR mode and Layer 4 NAT mode are recommended. These modes are described below and are used for the configurations presented in this guide. For configuring using DR mode, please refer to Appliance & eCopy ShareScan Server Configuration – Using Layer 4 DR Mode, for configuring using NAT mode, refer to Appliance & eCopy ShareScan Server Configuration – Using Layer 4 NAT Mode.

5.7.1. 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***.

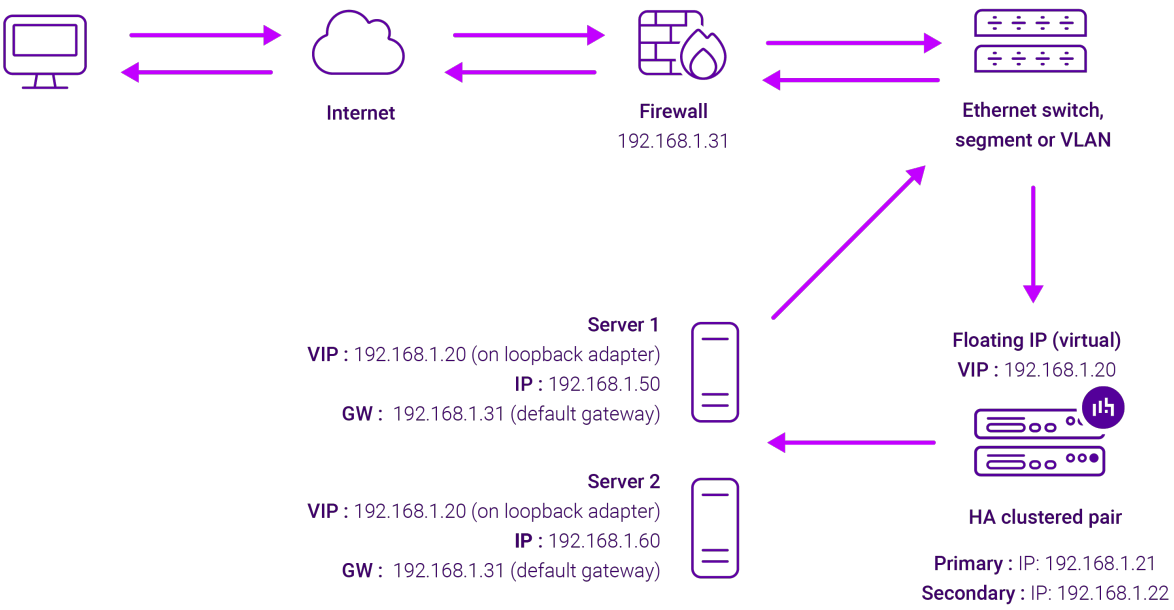

- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that each Real Server (and the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

5.7.2. Layer 4 NAT Mode

Layer 4 NAT mode is a high performance solution, although not as fast as layer 4 DR mode. This is because real server responses must flow back to the client via the load balancer rather than directly as with DR mode. The image below shows an example network diagram for this mode.

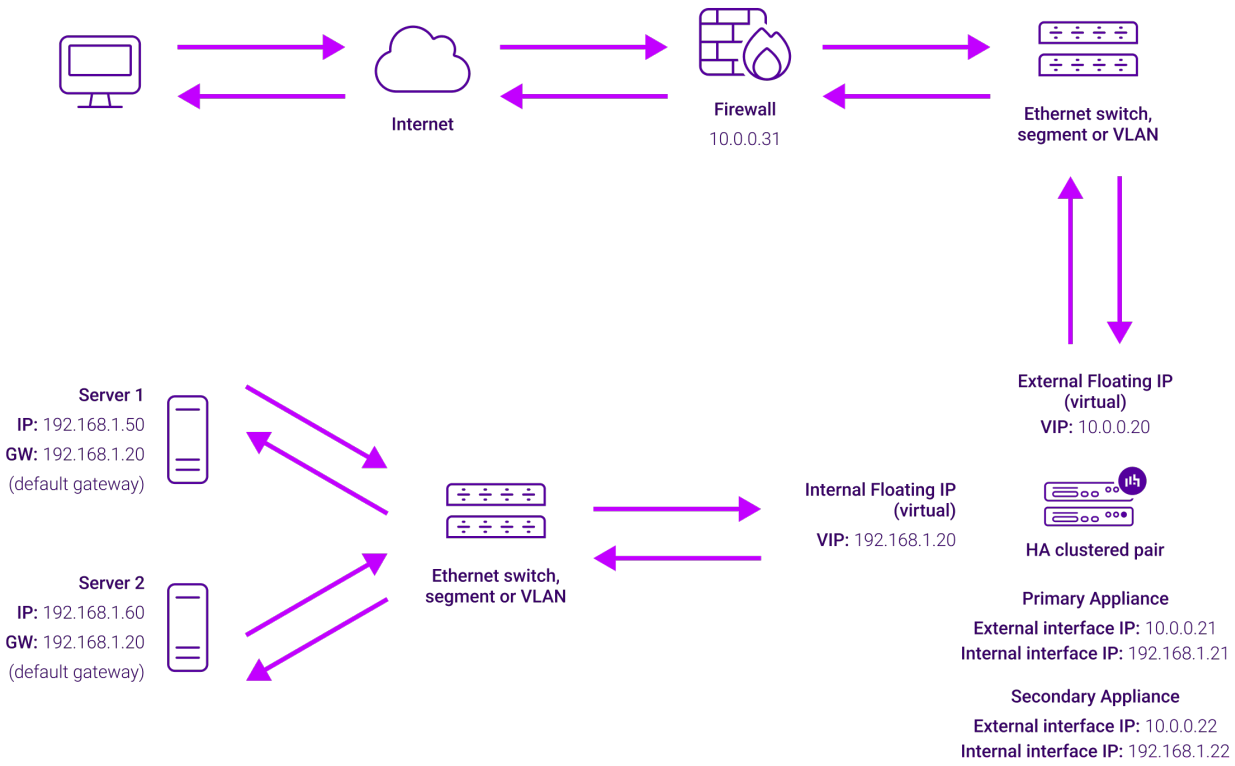

- The load balancer translates all requests from the Virtual Service to the Real Servers.
- NAT mode can be deployed in the following ways:
  - **Two-arm (using 2 Interfaces)** (as shown above) - Here, 2 subnets are used. The VIP is located in one subnet and the load balanced Real Servers are located in the other. The load balancer requires 2 interfaces, one in each subnet.

> **Note** This can be achieved by using two network adapters, or by creating VLANs on a single adapter.

    - Normally **eth0** is used for the internal network and **eth1** is used for the external network, although this is not mandatory since any interface can be used for any purpose.
    - If the Real Servers require Internet access, *Auto-NAT* should be enabled using the WebUI menu option: *Cluster Configuration > Layer 4 - Advanced Configuration*, the external interface should be selected.
    - The default gateway on the Real Servers must be set to be an IP address on the load balancer.

> **Note** For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can "float" (move) between Primary and Secondary appliances.

    - Clients can be located in the same subnet as the VIP or any remote subnet provided they can route to the VIP.
  - **One-arm (using 1 Interface)** - Here, the VIP is brought up in the same subnet as the Real Servers.

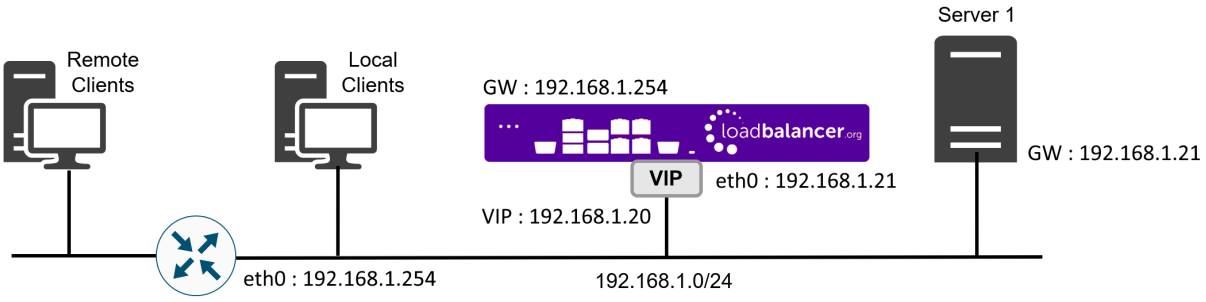

- To support remote clients, the default gateway on the Real Servers must be an IP address on the load balancer and routing on the load balancer must be configured so that return traffic is routed back via the router.

> **Note** For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can "float" (move) between Primary and Secondary appliances.

- To support local clients, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer. For more information please refer to One-Arm (Single Subnet) NAT Mode.

- If you want Real Servers to be accessible on their own IP address for non-load balanced services, e.g. RDP, you will need to setup individual SNAT and DNAT firewall script rules for each Real Server or add additional VIPs for this.
- Port translation is possible with Layer 4 NAT mode, e.g. VIP:80 → RIP:8080 is supported.
- NAT mode is transparent, i.e. the Real Server will see the source IP address of the client.

NAT Mode Packet re-Writing

In NAT mode, the inbound destination IP address is changed by the load balancer from the Virtual Service IP address (VIP) to the Real Server. For outbound replies the load balancer changes the source IP address of the Real Server to be the Virtual Services IP address.

The following table shows an example NAT mode setup:

Protocol	VIP	Port	RIP	Port
TCP	10.0.0.20	80	192.168.1.50	80

In this simple example all traffic destined for IP address 10.0.0.20 on port 80 is load-balanced to the real IP address 192.168.1.50 on port 80.

Packet rewriting works as follows:

1) The incoming packet for the web server has source and destination addresses as:

Source	x.x.x.x:34567	**Destination**	10.0.0.20:80

2) The packet is rewritten and forwarded to the backend server as:

Source	x.x.x.x:34567	**Destination**	192.168.1.50:80

3) Replies return to the load balancer as:

Source	192.168.1.50:80	**Destination**	x.x.x.x:34567

4) The packet is written back to the VIP address and returned to the client as:

Source	10.0.0.20:80	**Destination**	x.x.x.x:34567

5.7.3. Loadbalancer.org Recommended Mode

Where possible we recommend that Layer 4 Direct Routing (DR) mode is used. This mode offers the best possible performance since replies go directly from the eCopy ShareScan servers to the client, not via the load balancer. It's also relatively simple to implement. Ultimately, the final choice does depend on your specific requirements and infrastructure.

6. Loadbalancer.org Appliance – the Basics

6.1. 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.

> **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.

6.2. Initial Network Configuration

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

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

6.3. 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.

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:

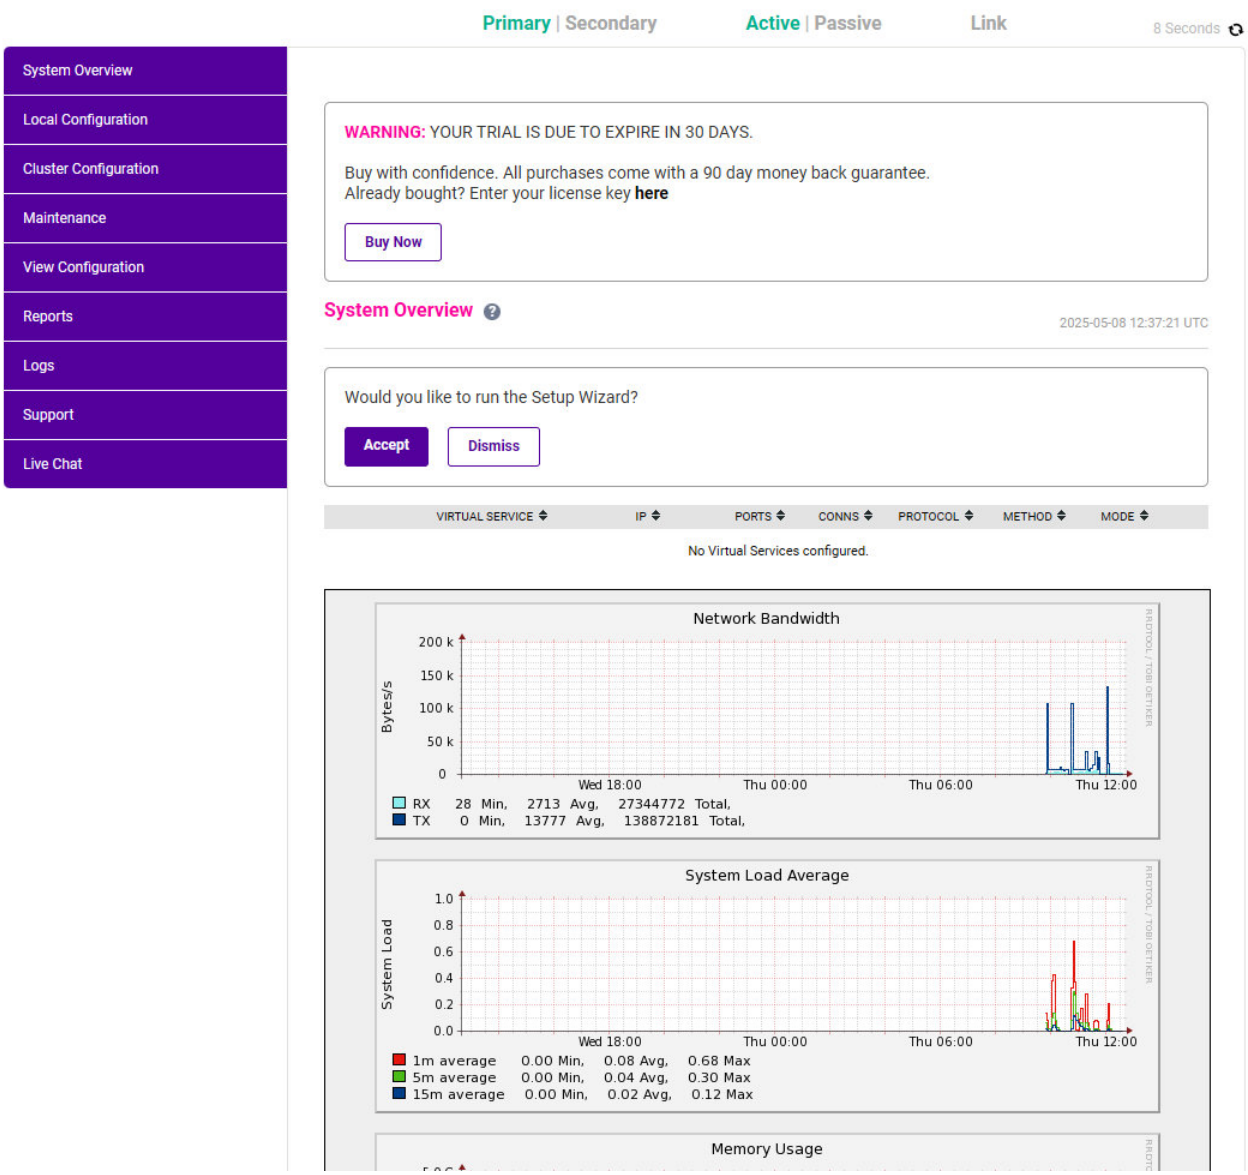

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

6.3.1. 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

6.4. 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.

6.4.1. 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:

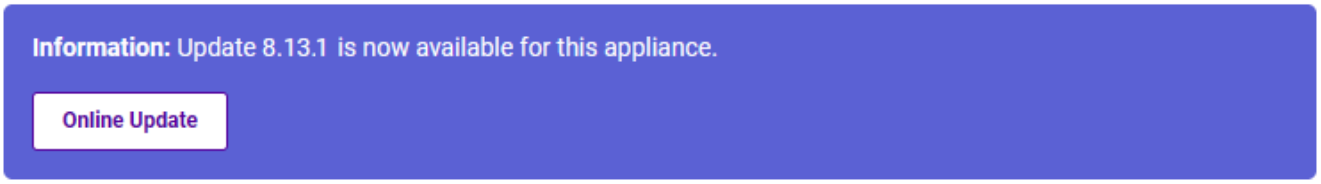

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

> **Important** Do not navigate away whilst the update is ongoing, this may cause the update to fail.

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

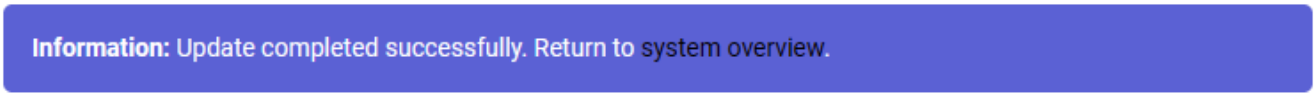

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

6.4.2. 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:

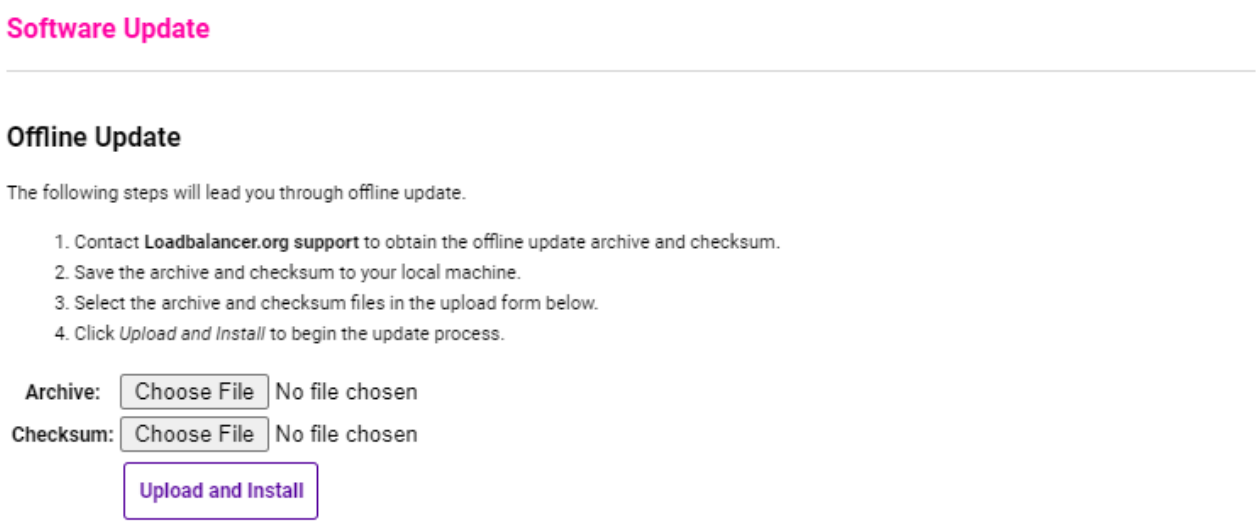

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

6.5. 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.

6.6. HA Clustered Pair Configuration

Loadbalancer.org recommend that load balancer appliances are deployed in pairs for high availability. In this guide a single unit is deployed first, adding a secondary unit is covered in Configuring HA - Adding a Secondary Appliance.

7. Appliance & eCopy ShareScan Server Configuration – Using Layer 4 DR Mode

> **Note** It's highly recommended that you have a working eCopy ShareScan environment first before implementing the load balancer and you must ensure that the DNS name points to the load balancer VIP.

7.1. Overview

This is our recommended deployment mode for eCopy ShareScan. It's ideal when you want the fastest possible deployment and cannot make any network changes on the eCopy ShareScan servers.

7.2. Load Balancer Configuration

7.2.1. Configure the Network Interface

1. One interface is required. For more information on configuring network settings please refer to Initial Network Configuration.

7.2.2. Configure the Virtual Service (VIP)

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

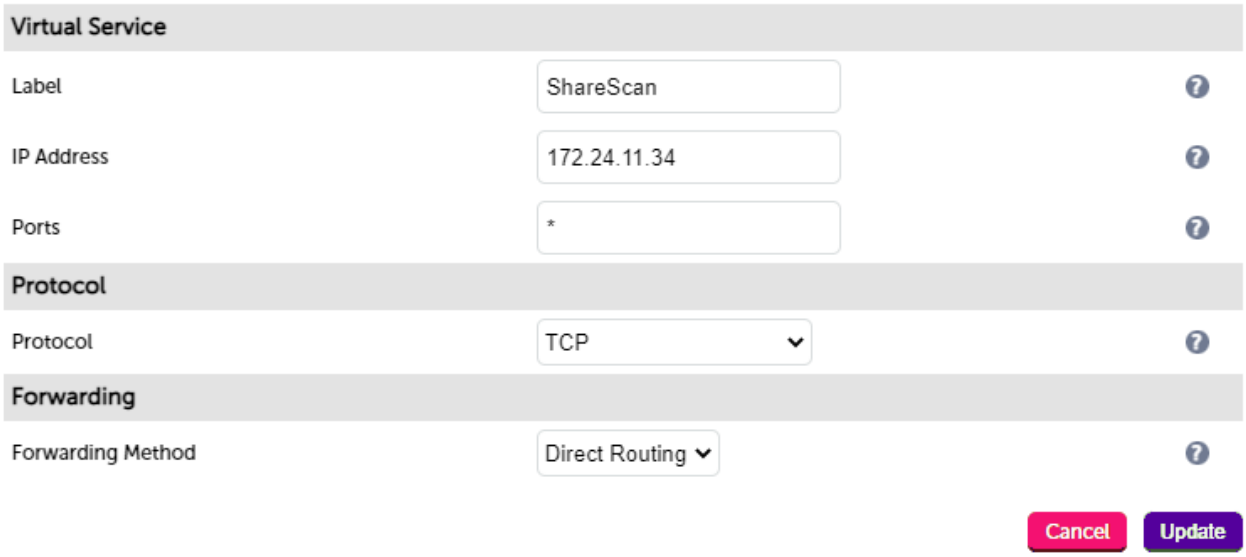

3. Enter an appropriate name (Label) for the VIP, e.g. **ShareScan**.
4. Set the Virtual Service IP address field to the required IP address, e.g. **172.24.11.34**.
5. Set the Virtual Service Ports field to *.
6. Leave *Protocol* set to **TCP**.
7. Ensure that *Forwarding Method* is set to **Direct Routing**.
8. Click **Update**.
9. Now click **Modify** next to the newly created Virtual Service.
10. Set *Balance Mode* (the load balancing algorithm) according to your requirements. "Weighted least connection" is the default and recommended method.
11. Persistence is enabled by default for new layer 4 VIPs and is based on source IP address. The persistence timeout can be set using the *Persistence Timeout* field, the default is 5 minutes which is normally fine for HTTP/HTTPS traffic.
12. Set the *Health Checks Check type* menu to **Negotiate**.
13. Set *Check Port* to **443**.
14. Set *Protocol* to **HTTPS**.
15. Ensure that *Request to send* and *Response expected* are both blank.
16. Click **Update**.

> **Note** For eCopy health check you can either monitor ports 8080 or 443 for the Tomcat service (Konica Minolta) or, ports 9600 for web-based devices or 9261 for embedded devices, like Ricoh and Canon as per the table in Ports & Protocols.

7.2.3. Configure the Real Servers (RIPs)

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

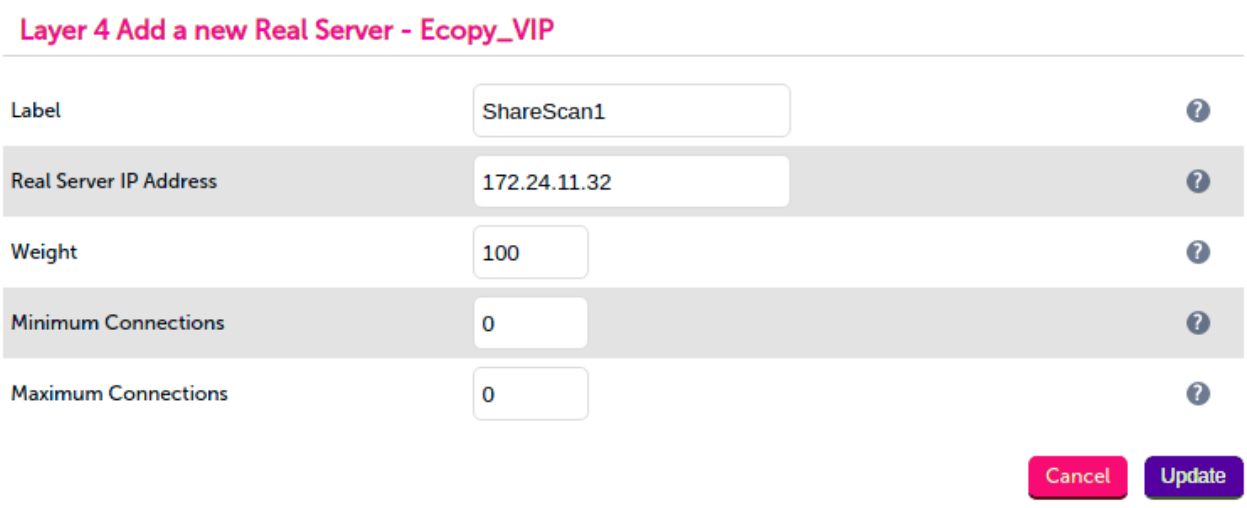

3. Enter an appropriate name (Label) for the first eCopy ShareScan server, e.g. **ShareScan1**.
4. Change the ***Real Server IP Address*** field to the required IP address, e.g. **172.24.11.32**.
5. Leave other settings at their default values.
6. Click **Update**.
7. Repeat the above steps for your other eCopy ShareScan server(s).

7.3. eCopy ShareScan Server Configuration

7.3.1. Solve the 'ARP Problem'

As mentioned previously, DR mode works by changing the destination MAC address of the incoming packet to match the selected ShareScan server on the fly which is very fast. When the packet reaches the ShareScan server it expects the ShareScan server to own the Virtual Services IP address (VIP). This means that you need to ensure that the ShareScan server (and the load balanced application) respond to both the ShareScan servers own IP address and the VIP. The ShareScan server should not respond to ARP requests for the VIP. Only the load balancer should do this.

To achieve this, a loopback adapter is added to the ShareScan servers. The IP address is set to be the same as the Virtual Service and the loopback adapter is configured so that it does not respond to ARP requests. Please refer to Solving the ARP Problem for full details of solving the ARP problem for Windows 2012 & later.

7.4. DR Mode – Key Points

- You ***must*** solve the 'ARP Problem' on all eCopy ShareScan servers in the cluster (please refer to Solving the ARP Problem for more information)
- Virtual Services & Real Servers (i.e. the eCopy ShareScan servers) must be within the same switch fabric. They can be on different subnets but this cannot be across a router – this is due to the way DR mode works, i.e. by changing MAC addresses to match the destination server
- Port translation is not possible, e.g. VIP:80 → eCopy ShareScan:82 is not allowed. The port used for the VIP & RIP must be the same

7.5. Configure ShareScan server registry settings

The following registry changes should be made on the ShareScan servers:

1. Using the Start menu, enter **regedit** to access the registry.
2. Open/expand the tree on the left pane and select **HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Nuance\ShareScan**.
3. In the right pane, choose the string **ManagerIP**. Double click on it and enter the IP address of the load balancer VIP.
4. Next, right click on the right side pane and select "new string". Name the new string value as **ClusterNodeIP**. Double click on the new string and enter the IP address of the main network adapter, i.e. the real server IP address.
5. In the left pane, now navigate to ***ShareScanManager***. In the right pane, right click and choose ***New String*** and enter the name **ClusterName**.

6. Double click on **ClusterName** and change the value to the FQDN of the load balancer VIP and click **OK**.
7. Reboot the server to apply the registry changes.

8. Appliance & eCopy ShareScan Server Configuration – Using Layer 4 NAT Mode

> **Note** It's highly recommended that you have a working eCopy ShareScan environment first before implementing the load balancer.

8.1. Overview

If the load balancer and the eCopy ShareScan servers are not part of the same layer 2 network, then DR mode cannot be used. If you require a high performance solution that is transparent by default (i.e. the client IP address is maintained through the load balancer) then layer 4 NAT mode can be used. Layer 4 NAT mode is also a high performance solution, although not as fast as layer 4 DR mode. This is because eCopy ShareScan server responses must flow back to the client via the load balancer rather than directly as with DR mode.

8.2. Load Balancer Configuration

8.2.1. Configure the Network Interfaces

Layer 4 NAT mode is typically used in a 2-arm configuration where the VIP is located in one subnet and the load balanced Real Servers are located in another. This can be achieved by using two network adapters, or by creating VLAN's on a single adapter. Single arm configuration is also supported under certain conditions - for more information please refer to Layer 4 NAT Mode.

To configure an additional network interface for a 2-arm configuration:

1. Using the WebUI, navigate to *Local Configuration > Network Interface Configuration*.
2. Scroll to the *IP Address Assignment* section.

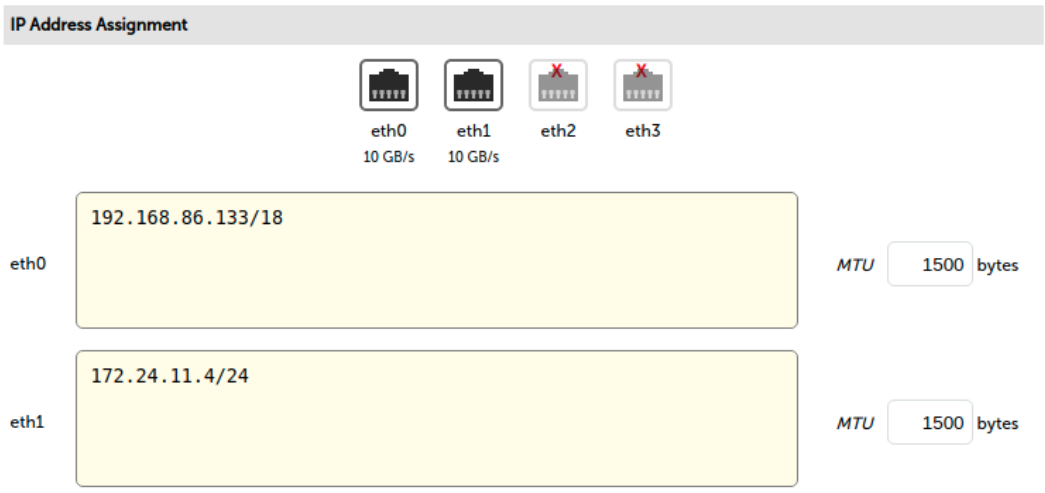

3. Specify an appropriate IP address for **eth1** in CIDR format as shown above.

4. Click **Configure Interfaces**.

> **Note** There are no restrictions on which interface is used for each requirement.

8.2.2. Configure the Virtual Service (VIP)

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

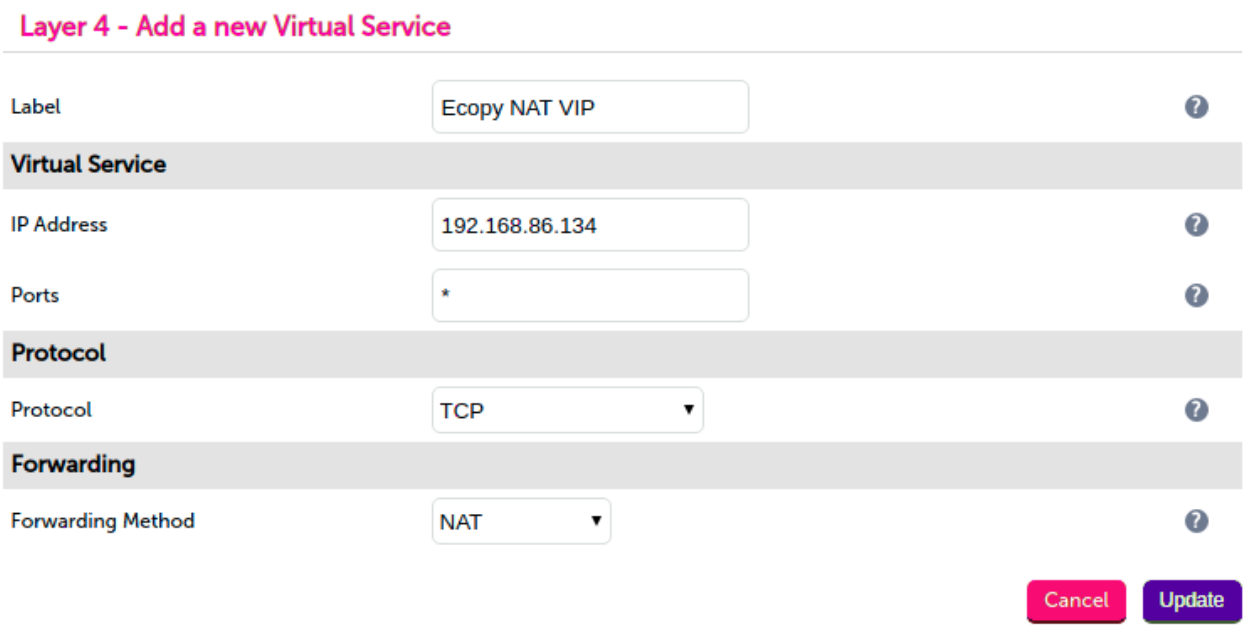

3. Enter an appropriate name (Label) for the VIP, e.g. **eCopy NAT VIP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.86.134**.
5. Set the *Virtual Service Ports* field to *.
6. Leave *Protocol* set to **TCP**.
7. Set the *Forwarding Method* to **NAT**.
8. Click **Update**.
9. Now click **Modify** next to the newly created Virtual Service.
10. Set *Balance Mode* (the load balancing algorithm) according to your requirements. "Weighted least connection" is the default and recommended method.
11. Persistence is enabled by default for new layer 4 VIPs and is based on source IP address. The persistence timeout can be set using the *Persistence Timeout* field, the default is 5 minutes which is normally fine for HTTP/HTTPS traffic.
12. Set the *Health Checks Check type* menu to **Negotiate**.
13. Set *Check Port* to **443**.
14. Set *Protocol* to **HTTPS**.

15. Ensure that the fields ***Request to send*** and ***Response expected*** are blank.
16. Click **Update**.

> **Note** For eCopy health check you can either monitor ports 8080 or 443 for the Tomcat service (Konica Minolta) or, ports 9600 for web-based devices or 9261 for embedded devices, like Ricoh and Canon as per the table in Ports & Protocols.

8.2.3. Configure the Real Servers (RIPs)

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

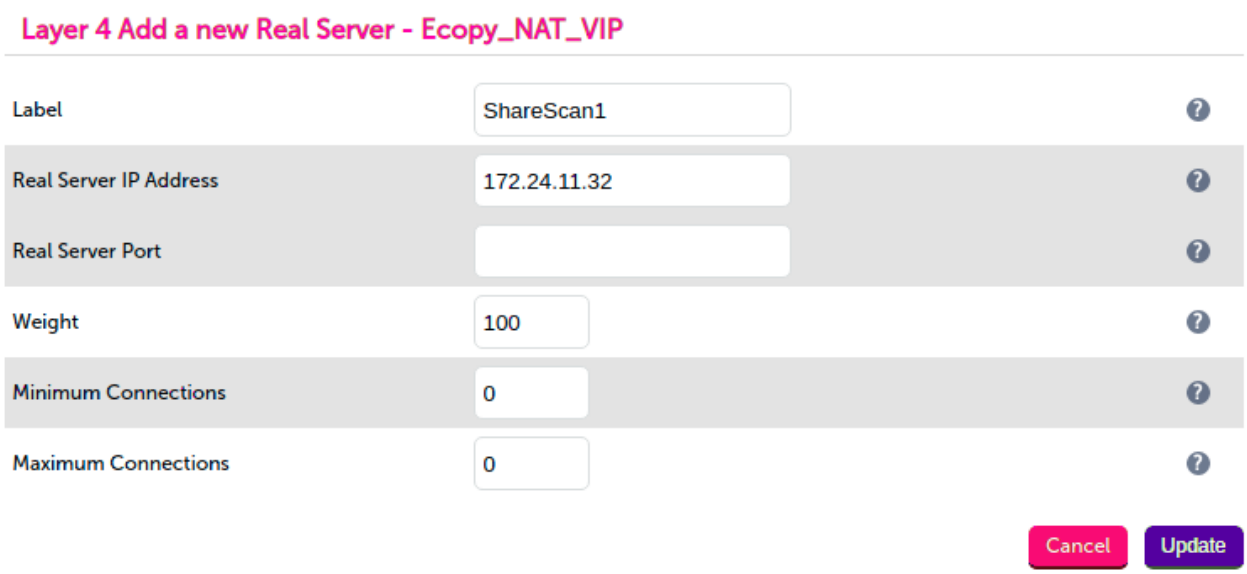

3. Enter an appropriate name (Label) for the first eCopy ShareScan server, e.g. **ShareScan1**.
4. Change the ***Real Server IP Address*** field to the required IP address , e.g. **172.24.11.32**.
5. Leave the Real Server Port ***blank***.
6. Leave other settings at their default values.
7. Click **Update**.
8. Repeat the above steps for your other eCopy ShareScan server(s).

8.2.4. Create a Floating IP to use for the eCopy ShareScan server server's Default Gateway

The default gateway on each eCopy ShareScan server server must be configured to be an IP address on the load balancer. It's possible to use the IP address assigned to the internal facing interface (**eth0** in this example) for the default gateway, although it's recommended that an additional floating IP is created for this purpose. This is required if two load balancers (our recommended configuration) are used. In this scenario if the Primary unit fails, the floating IP will be brought up on the Secondary.

To create a floating IP address on the load balancer:

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IPs*.
2. Enter the required IP address to be used for the default gateway, e.g. **172.24.11.35**.
3. Click **Update**.

   Once added, there will be two floating IP's, one for the Virtual Service (**192.168.86.134**) and one for the default gateway (e.g. **172.24.11.35**) as shown below:

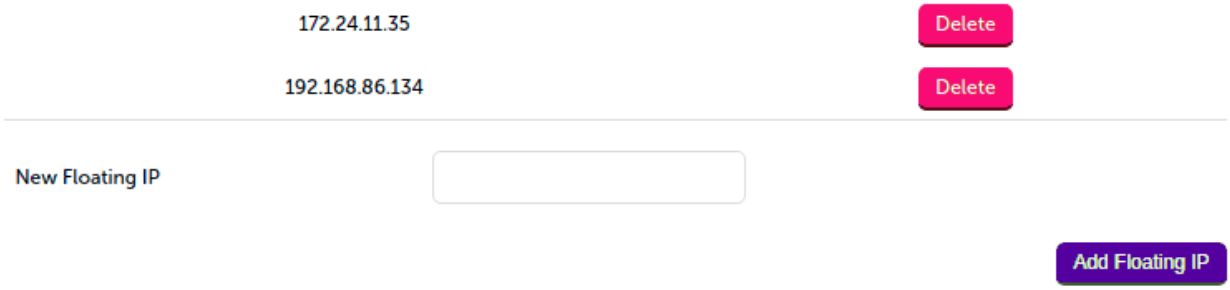

8.3. eCopy ShareScan Server Configuration

8.3.1. Default Gateway

To ensure return traffic passes back to the client via the load balancer, set the default gateway of each eCopy ShareScan server to be the floating IP address added in the previous step, in this example **172.24.11.35**.

8.3.2. NAT Mode – Key Points

- Virtual Services & Real Servers (i.e. the eCopy ShareScan servers) must be on different subnets
- The default gateway on the eCopy ShareScan servers should be an IP address on the load balancer (for an HA pair this must be a floating IP address)
- Port translation is possible, e.g. *VIP:80 → RIP:8080* is allowed

8.4. Configure ShareScan server registry settings

The following registry changes should be made on the ShareScan servers:

1. Using the Start menu, enter **regedit** to access the registry.
2. Open/expand the tree on the left pane and select **HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Nuance\ShareScan**.
3. In the right pane, choose the string **ManagerIP**. Double click on it and enter the IP address of the load balancer VIP.
4. Next, right click on the right side pane and select "new string". Name the new string value as **ClusterNodeIP**. Double click on the new string and enter the IP address of the main network adapter, i.e. the real server IP address.
5. In the left pane, now navigate to *ShareScanManager*. In the right pane, right click and choose *New String* and enter the name **ClusterName**.

6. Double click on **ClusterName** and change the value to the FQDN of the load balancer VIP and click **OK**.
7. Reboot the server to apply the registry changes.

8.5. Real Server (eCopy ShareScan) Health Checks

The load balancer performs regular health checks to ensure that each server in the cluster is healthy and able to accept client connections. The health check options at layer 4 have been outlined below.

8.5.1. Layer 4

By default, a TCP connect health check is used for newly created layer 4 Virtual Services. The following tables lists all options available:

Check Type	Description
Negotiate	Sends a request and looks for a specific response. This option enables the load balancer to perform a more robust check. For example, an HTTP check can be configured that requests a certain page and then looks for a specific word on that page.
Connect to port	Just do a simple connect to the specified port/service & verify that it's able to accept a connection.
Ping server	Sends an ICMP echo request packet to the Real Server.
External check	Use a custom script for the health check.
No checks, always Off	All Real Servers are off.
No checks, always On	All Real Servers are on (no checking).
5 Connects, 1 Negotiate	Do 5 connect checks and then 1 negotiate check.
10 Connects, 1 Negotiate	Do 10 connect checks and then 1 negotiate check.

> **Note** For full details on the options available, please refer to Real Server Health Monitoring & Control.

8.6. Server Feedback Agent

The load balancer can modify the weight (amount of traffic) of each server by gathering data from either a custom agent or an HTTP server. For layer 4 VIPs the feedback method can be set to either agent or HTTP, for Layer 7 VIPs, only the agent method is supported.

A telnet to port 3333 on a Real Server with the agent installed will return the current idle stats as an integer value in the range 0 – 100. The figure returned can be related to CPU utilization, RAM usage or a combination of both. This can be configured using the XML configuration file located in the agents installation folder (by default *C:\ProgramData\LoadBalancer.org\LoadBalancer*).

The load balancer typically expects a 0-99 integer response from the agent which by default relates to the current CPU idle state, e.g. a response of 92 would imply that the Real Servers CPU is 92% idle. The load balancer will then use the formula (92/100*requested_weight) to find the new optimized weight.

> **Note** The 'Requested Weight' is the weight set in the WebUI for each Real Server. For more information please also refer to this blog.

8.6.1. Agent Download

The latest Windows feedback agent can be downloaded from here. To install the agent, run loadbalanceragent.msi on each eCopy ShareScan Server:

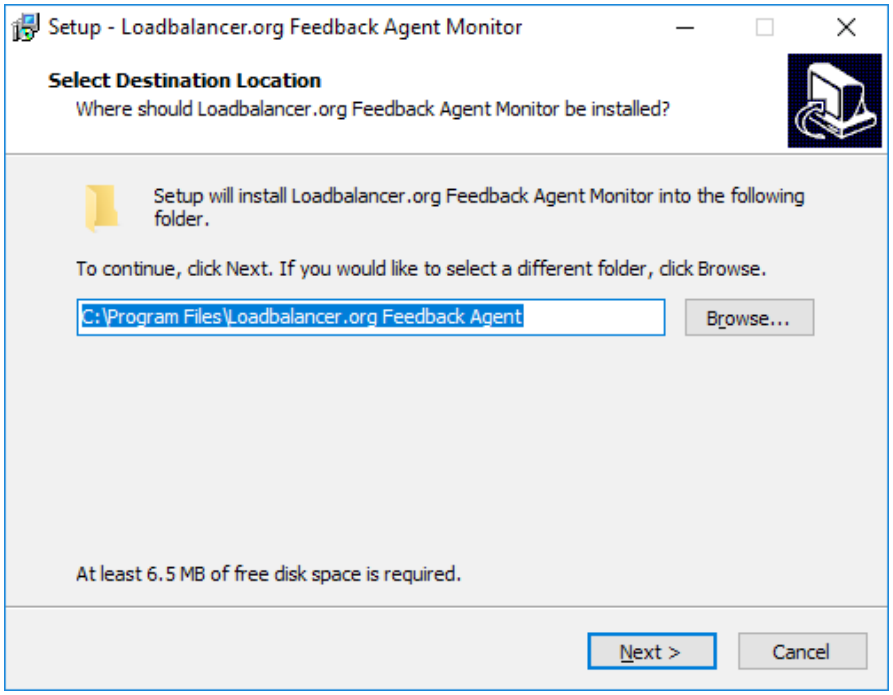

Leave the default location or change according to your requirements, click **Next**.

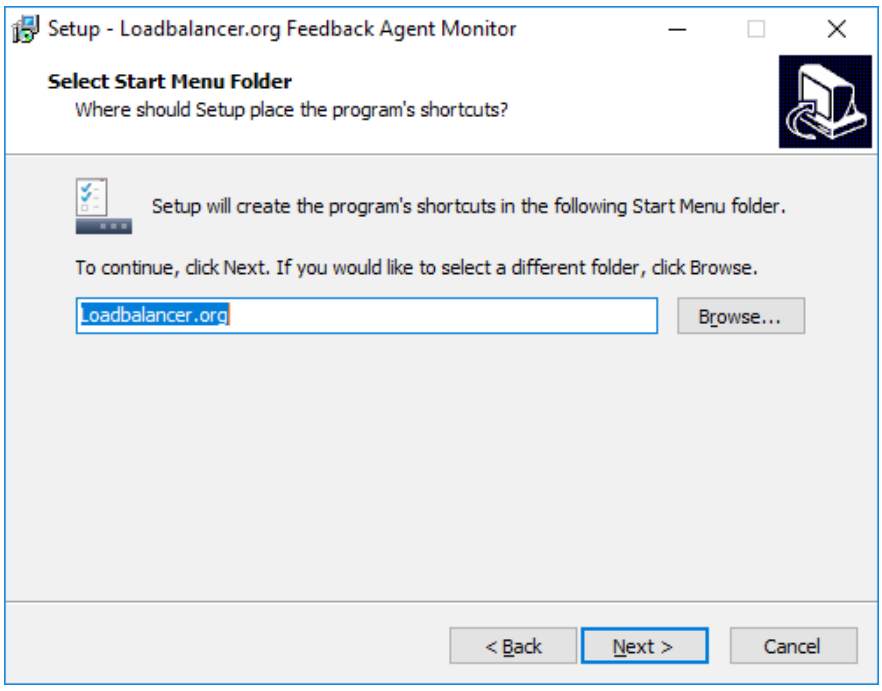

Leave the default location or change according to your requirements, click **Next**.

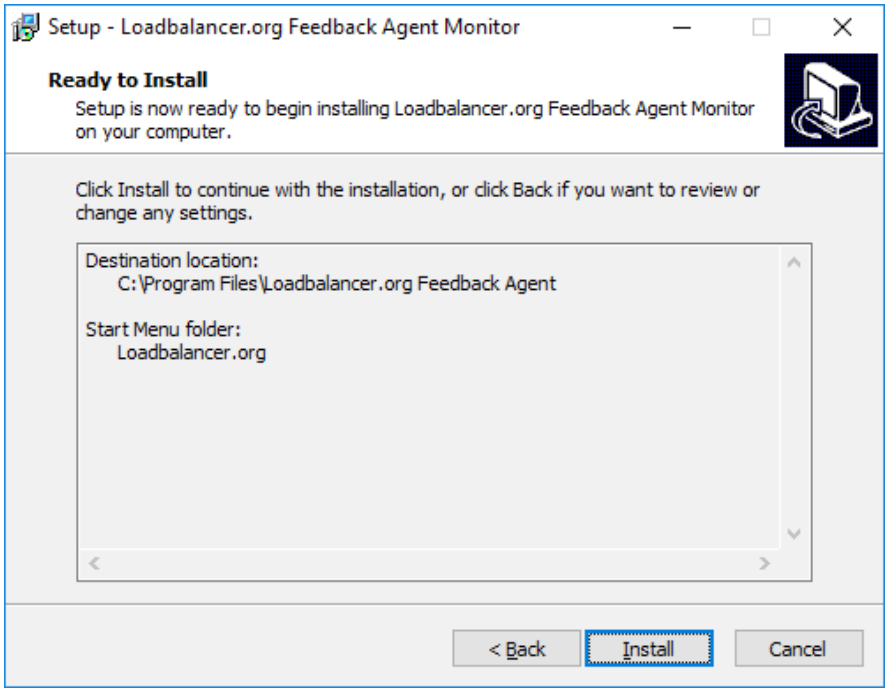

Click **Install** to start the installation process.

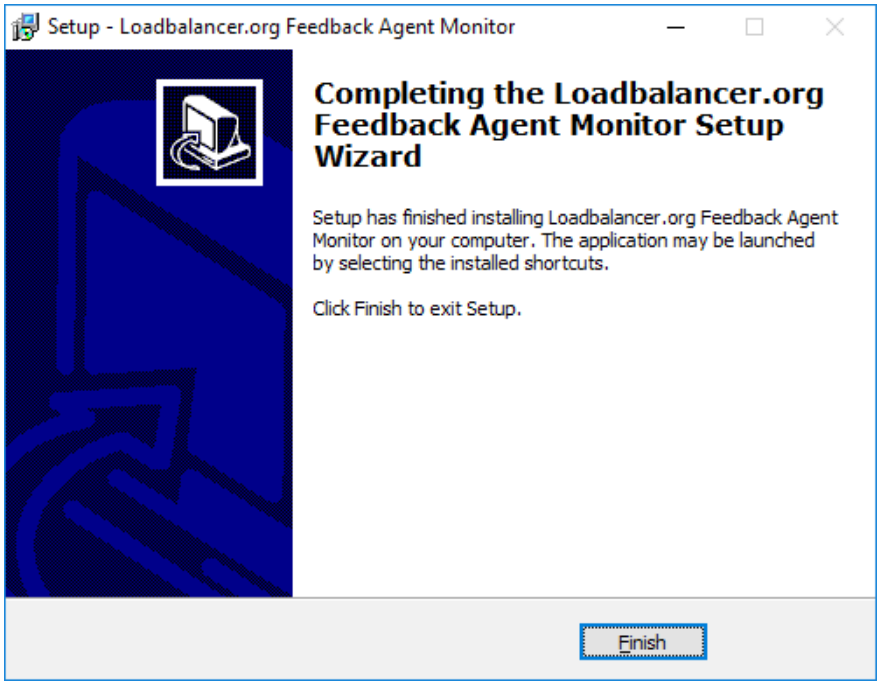

Click **Finish**.

> **Note** The agent should be installed on all eCopy ShareScan Servers in the cluster.

8.6.2. Starting the Agent

Once the installation has completed, you'll need to start the service on the Real Servers. The service is controlled by the Feedback Agent monitor & control program that is also installed along with the Agent. This can be accessed on the Windows server from: ***Start> Loadbalancer.org > Loadbalancer.org Feedback Agent***. It's also possible to start the service using the services snap-in – the service is called 'LBCPUMon'.

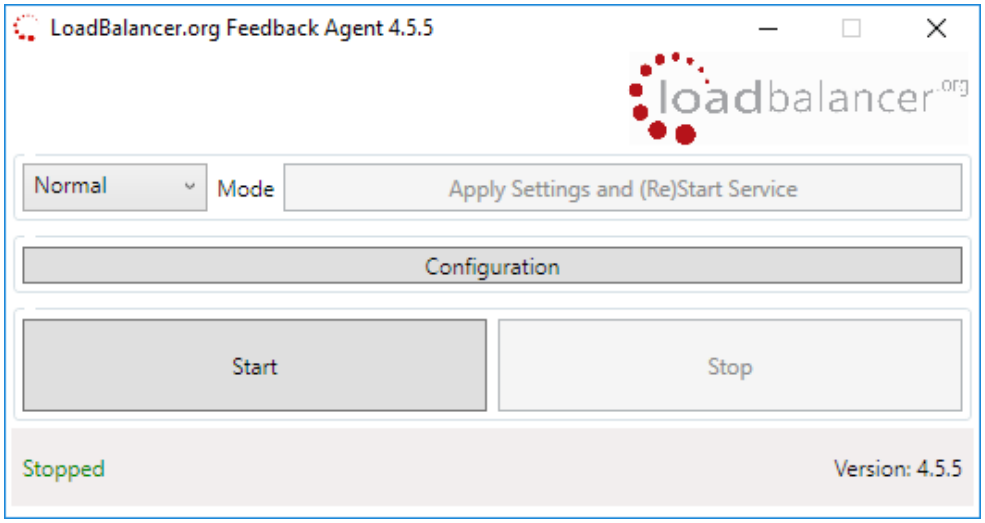

- To start the service, click the **Start** button.
- To stop the service, click the **Stop** button.

8.6.3. Configuration

To Configure Virtual Services to use the feedback agent, follow the steps below:

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 Virtual Services* and click **Modify** next to the Virtual Service.

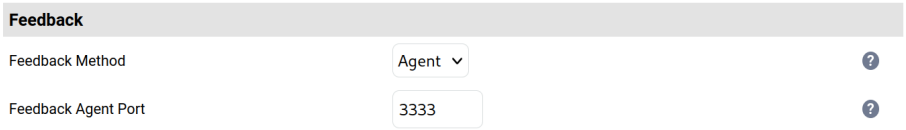

2. Change the Feedback Method to **Agent**.
3. Click **Update**.
4. Reload/Restart services as prompted.

8.7. Load Balancer Transparency

8.7.1. Layer 4

Both Layer 4 DR mode and layer 4 NAT mode are transparent by default. This means that ShareScan will log the actual IP address of the client rather than the IP address of the load balancer.

9. Testing & Verification

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

9.1. Using the System Overview

The System Overview can be accessed via the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the eCopy

ShareScan servers) and shows the state/health of each server as well as the state of the cluster as a whole. The example below shows that all ShareScan servers are healthy (green) and available to accept connections:

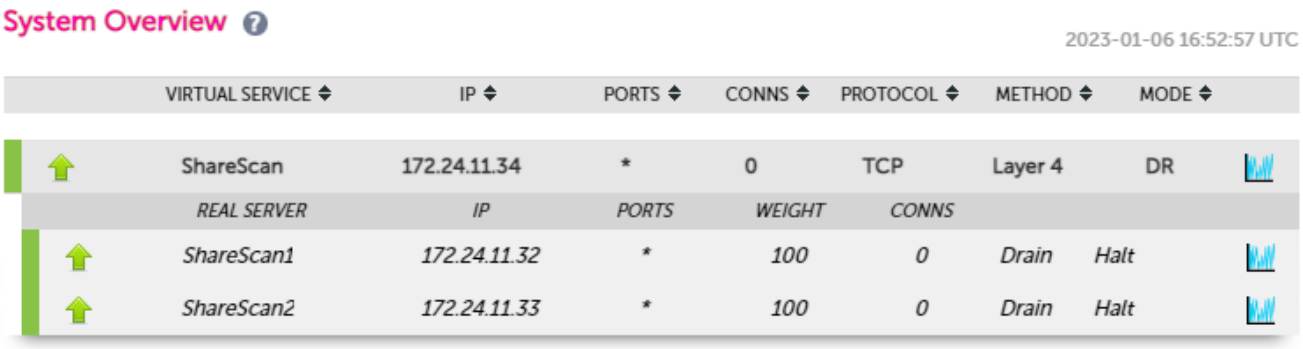

If one of the servers within the cluster fails its health check, that server will be colored red and the cluster will be colored yellow as shown below:

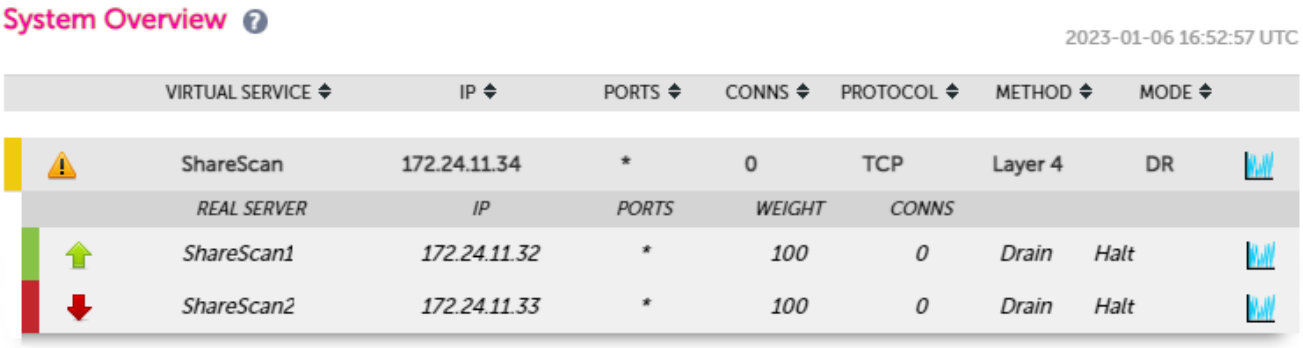

Make sure that all servers are up (green) and verify that clients can connect to the VIP and access all load balanced services.

> **Note** Make sure that DNS points at the VIP rather than individual servers.

9.2. Using the eCopy ShareScan Troubleshooting Tool

eCopy ShareScan has an application troubleshooting tool that can be utilised to test connectivity via the load balancer VIP to the eCopy Sharescan real servers. As such, to initiate the troubleshooter follow the defined steps below.

1. Start ShareScan Troubleshooter on all Sharescan server nodes and one Sharescan client that has access to the Sharescan load balanced VIP.
2. On the ShareScan server nodes click the *Advanced drop down menu > Network tests > Server side* network test. It will automatically start listening on port 9600.
3. On the non-load balanced client PC click the *Advanced drop down menu > Network tests > Client side* network test.
4. Enter the ShareScan Virtual IP into the Server address / hostname, and click Connect.
5. The request should now connect to one of the Sharescan servers via the load balancer VIP resulting in a

connection message in one of the open dialogues on one of the servers.

6. Confirm that the IP shown by the connection message is the IP of the is of the client PC and NOT the IP of the load balancer.

10. Technical Support

If you have any questions regarding the appliance or would like assistance designing your deployment, please don't hesitate to contact our support team: support@loadbalancer.org.

11. Further Documentation

For additional information, please refer to the Administration Manual.

12. Appendix

12.1. Solving the ARP Problem

When using Layer 4 DR mode, the ARP problem must be solved. This involves configuring each eCopy ShareScan server to be able to receive traffic destined for the VIP, and ensuring that each eCopy ShareScan server does not respond to ARP requests for the VIP address – only the load balancer should do this. The steps below are for Windows 2012 & later.

12.1.1. 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, the strong/weak host behavior must be configured on each Real Server. The weak host model allows packets with any IP to be sent or received via an interface. The strong host model only allows packets with an IP belonging to the interface to be sent or received.

> **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**.

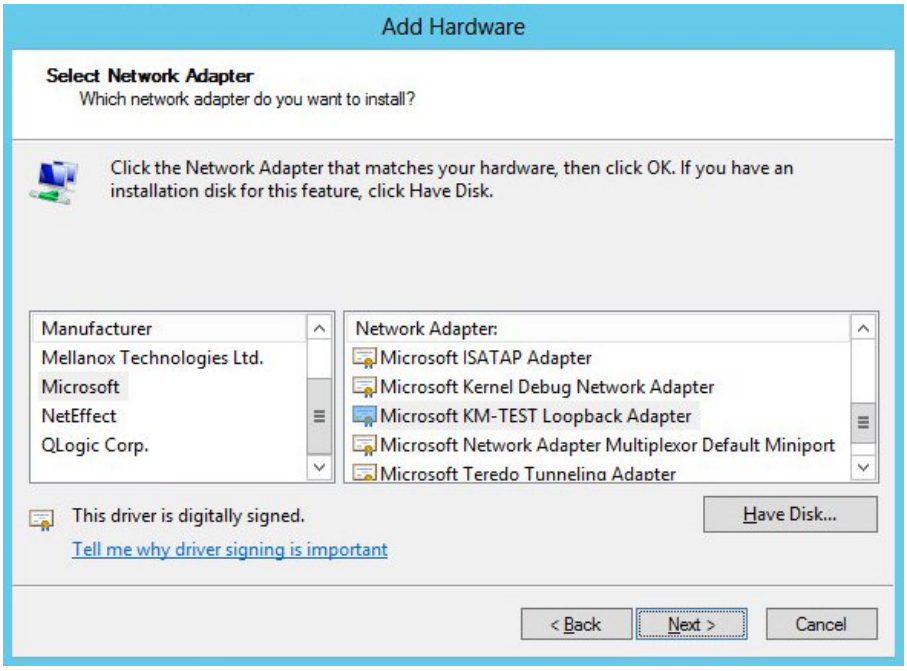

5. Select **Microsoft** & **Microsoft KM-Test Loopback Adapter**, click **Next**.
6. Click **Next** to start the installation, when complete click **Finish**.

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:

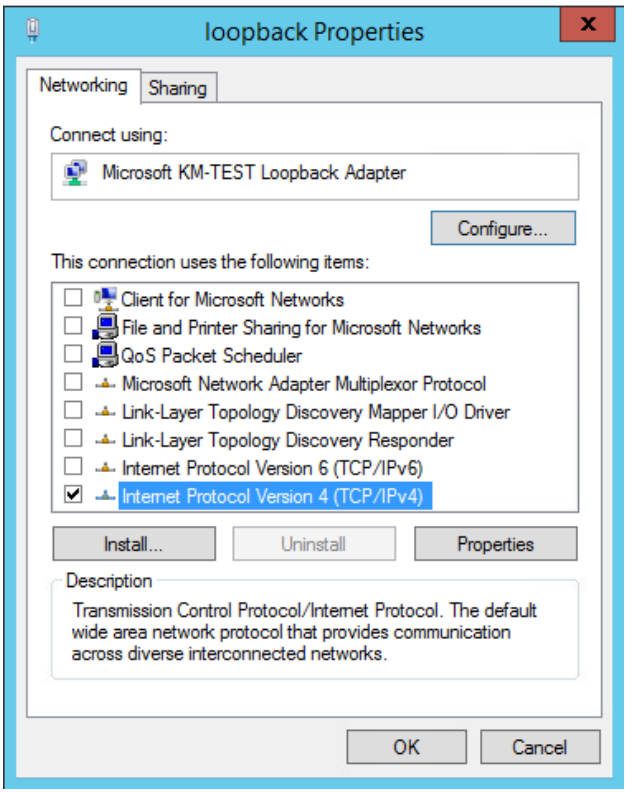

2. Ensure that **Internet Protocol Version (TCP/IPv4)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service address (VIP) with a subnet mask of **255.255.255.255**, e.g. **192.168.2.20/255.255.255.255** as shown below:

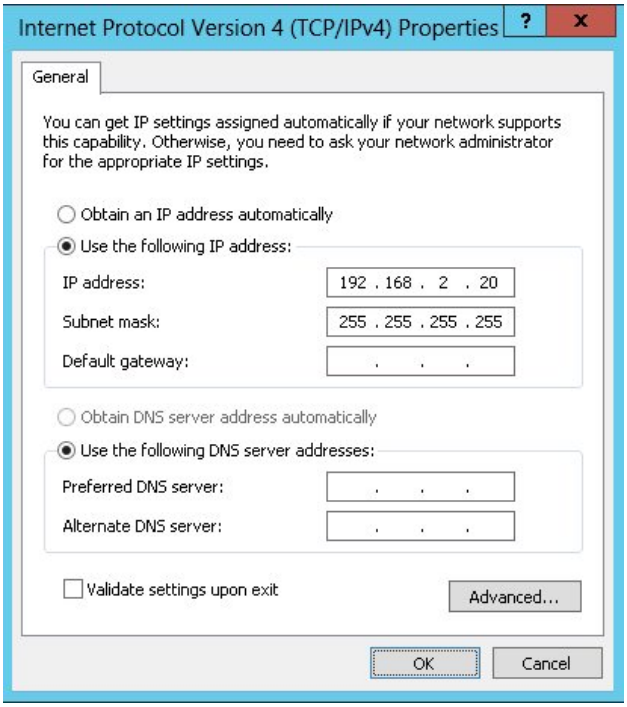

> **Note** **192.168.2.20** is an example, make sure you specify the correct VIP address.

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

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:

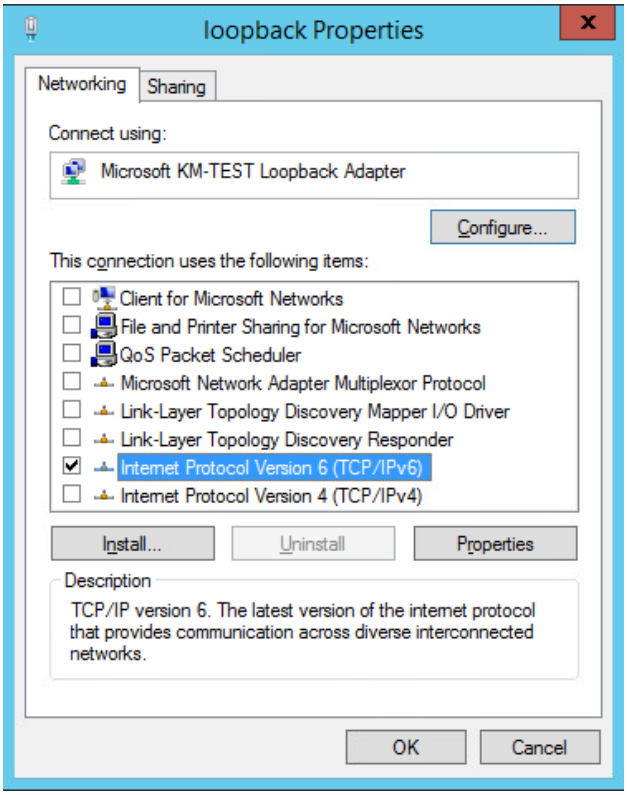

2. Ensure that **Internet Protocol Version (TCP/IPv6)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) and set the ***Subnet Prefix Length*** to be the same as your network setting, e.g. **2001:470:1f09:e72::15/64** as shown below:

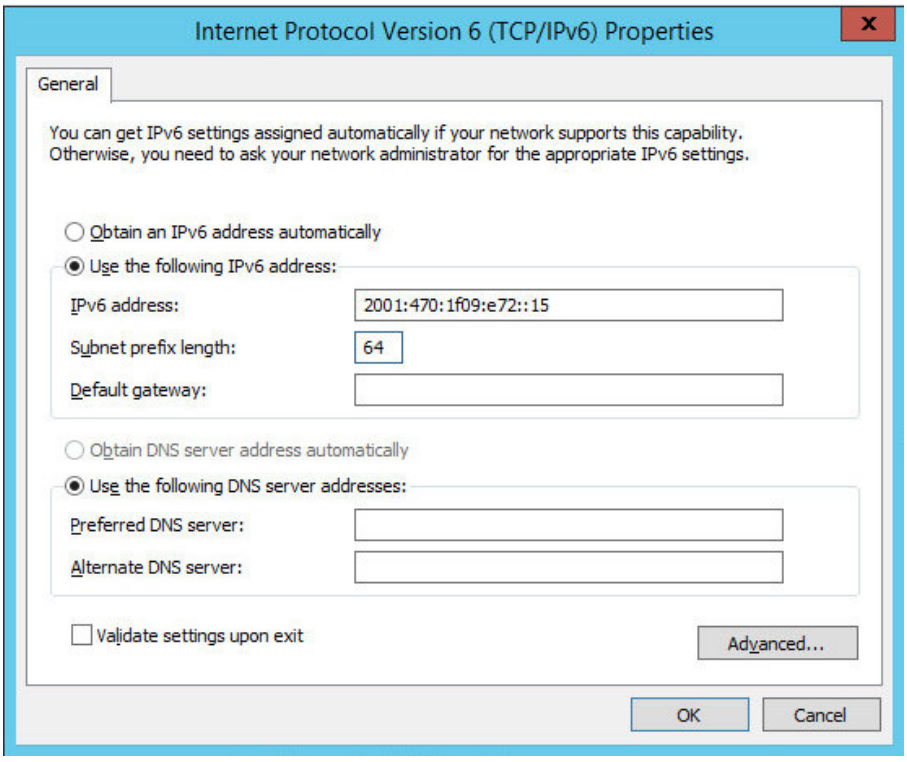

> **Note** **2001:470:1f09:e72::15/64** is an example, make sure you specify the correct VIP address.

> **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:

> **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:

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:

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

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

12.2. 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.

> **Note** For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the documentation library

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.

12.2.1. 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

WebUI Main Menu Option	Sub Menu Option	Description
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	Fallback Page	Fallback page configuration
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.

12.2.2. 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*.

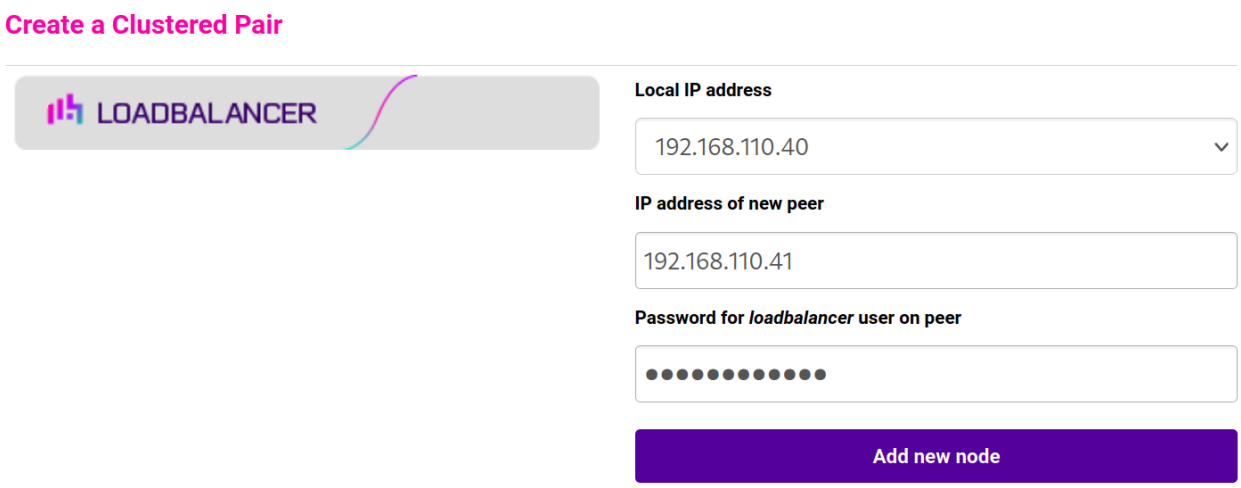

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

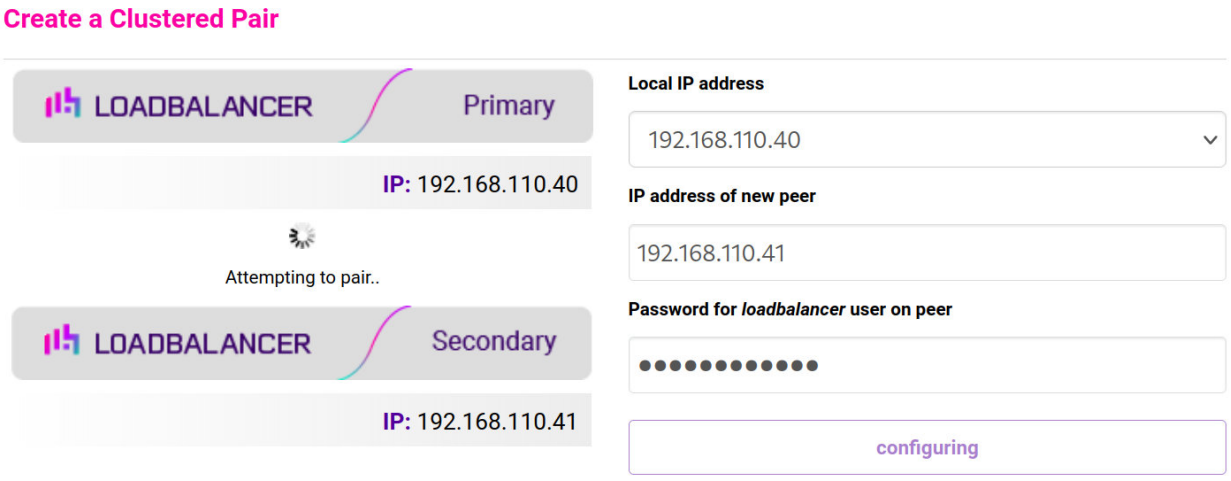

6. Once complete, the following will be displayed on the Primary appliance:

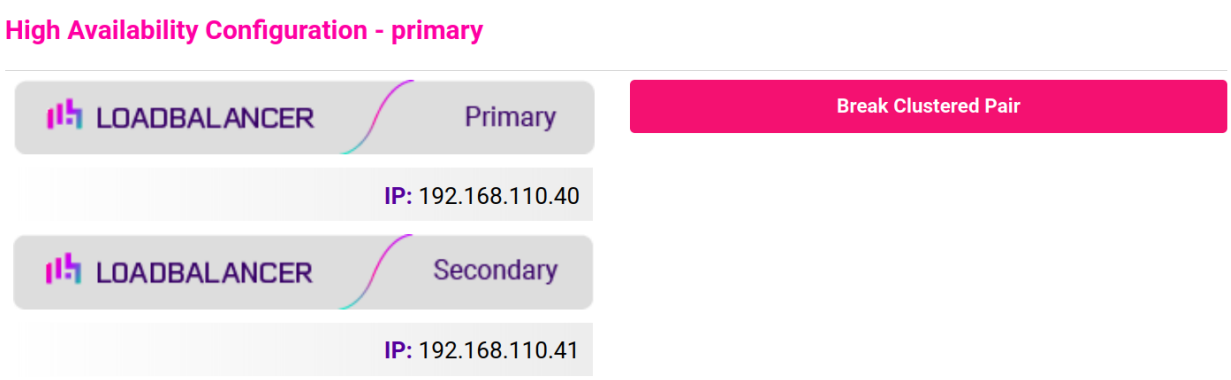

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	27 November 2019	Initial draft	Initial draft	IBG, AH
1.0.1	20 December 2019	Guide update	Health checks updated	IBG
1.0.2	9 March 2020	Guide update	Health checks updated	IBG
1.0.3	3 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.1.0	1 November 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.1.1	5 January 2023	Updated Testing & Verification section	General Improvements	RJC
1.1.2	5 January 2023	Combined software version information into one section Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section	Housekeeping across all documentation	AH
1.1.3	2 February 2023	Updated screenshots	Branding update	AH
1.1.4	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.2.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH

LOADBALANCER

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.