LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing PaperCut

Version 2.0.0

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced PaperCut environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Microsoft print Server & Papercut server 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 for load balancing PaperCut print servers. 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. Papercut

- Papercut NG/MF : v20 and later
- Mobility Print : v1.0.3400 and later

4. PaperCut NG/MF and Mobility Print

- **PaperCut NG** is a software application that helps organizations manage printing. It helps to minimise waste, save paper and toner/ink, improve document security, save the system administrator's time and encourage end-users to improve printing behavior.
- **Papercut MF** has the same features as NG, but also has the additional ability to integrate directly with Multifunction Devices (MFDs) and other hardware to deliver extra features. The embedded software available with PaperCut MF runs on the MFD and allows you to track and control printing, copying, faxing, and scanning.
- **PaperCut Mobility Print** simplifies the printing process for bring your own devices (BYOD) and other end-user managed devices, such as smartphones, tablets, laptops or Chromebooks. End users can quickly discover and set up printers on their own, regardless of their operating system or the brand of printer.

4.1. PaperCut Server Types

4.1.1. Primary Server / Application Server

One of the servers in the deployment must be nominated as the Primary Server. This system runs the Application Server software that is responsible for providing the user interface, storing the data, and managing the application logic. The system nominated for this task is usually a print server (also refer to the note in Application Servers below).

4.1.2. Secondary Server / Print Server

Other print servers in the deployment are known as Secondary Servers. These servers run the Print Provider component and typically the Mobility Print server component if deployed. Secondary servers communicate back to the Primary Server to enable print job control and management.

4.1.3. Site Server

Site Servers duplicate the key features of the Primary Server to a local site during an outage. MFDs are configured to connect to a Site Server as if it were the primary server to remove their reliance on WAN links. PaperCut secondary servers (Print Providers) are also aware of their local Site Server, providing a failover server if the primary server cannot be contacted.

5. Load Balancing PaperCut

Note	It's highly recommended that you have a working PaperCut version 20 or later environment first before implementing the load balancer.

5.1. Load Balancing & HA Requirements

The following Papercut server roles can be load balanced to provide improved performance, HA and resilience:

5.1.1. Application Servers

While not strictly load balancing, this feature improves system resilience by allowing multiple PaperCut Application Servers to be positioned behind a Network Load Balancer in an active/passive configuration. This ensures high availability of the PaperCut system as discussed here.

Important	As mentioned here under ***Requirements***, when Application Servers are deployed in active/passive mode behind a load balancer, all print queues and PaperCut Mobility Print installations must be hosted on PaperCut Secondary or Site Servers. These components cannot be hosted on the Application Server when using this solution.

5.1.2. Print Servers

For environments with large numbers of users or high print volumes, PaperCut supports distributing print jobs across multiple print servers using a network load balancer as mentioned here.

5.1.3. Mobility Print Servers

Mobility Print can be installed on multiple servers and can be configured to work with a Network Load Balancer to

distribute print requests. More detailed are available here.

> **Note** This configuration is only possible with on-network methods of printer discovery (mDNS, DNS, Known Host) and is not currently supported for Mobility Print Cloud Print.

5.2. Virtual Service (VIP) Requirements

The VIPs required are slightly different depending on which load balancing method is used. This is due to two main reasons:

1. Layer 7 SNAT mode is not transparent. This means that the source IP address of the client is lost and replaced by the load balancer's own IP address. However, as mentioned here under ***Set up trusted proxy servers***, some Papercut features require the client's IP address to be available to function correctly. As an alternative, `X-Forward-For` headers can be used. HTTPS traffic on port 9192 must first be decrypted on the load balancer to allow the headers to be inserted.
2. Layer 7 SNAT mode does not support UDP. This means that DNS printer discovery can only be implemented using TCP and mDNS which relies solely on UDP and therefore cannot be load balanced when using layer 7 SNAT mode.

5.2.1. When using Layer 4 DR Mode

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	App_Servers_HTTP-HTTPS	L4 DR	TCP 9191,9192	None	HTTP (GET)
VIP 2	App_Servers_OtherPorts	L4 DR	TCP 9193	None	Connect to Port
VIP 3	Print_Servers	L4 DR	TCP 445	Source IP	Connect to Port
VIP 4	Mobility_Print_Servers	L4 DR	TCP 9163,TCP 9164	Source IP	HTTP (GET)
VIP 5	Mobility_Discovery_DNS	L4 DR	TCP/UDP 53	Source IP	Connect to Port
VIP 6	Mobility_Discovery_mDNS	L4 DR	UDP 5353	Source IP	Connect to Port

5.2.2. When using Layer 7 SNAT Mode

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	App_Servers_HTTP-HTTPS	L7 SNAT (HTTP/HTTPS)	TCP 9191 (HTTP), TCP 9192 (HTTPS - via HAProxy SSL termination)	None	HTTP (GET)
VIP 2	App_Servers_OtherPorts	L7 SNAT (TCP)	TCP 9193	None	Connect to Port
VIP 3	Print_Servers	L7 SNAT (TCP)	TCP 445	Source IP	Connect to Port

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 4	Mobility_Print_Servers	L7 SNAT (TCP)	TCP 9163,TCP 9164	Source IP	HTTP (GET)
VIP 5	Mobility_Discovery_DNS	L7 SNAT (TCP)	TCP 53	Source IP	Connect to Port

Important The ports specified above are the standard Papercut ports. It's possible to change some of these ports. If this has been done, the VIP ports must be modified accordingly.

Note VIP 2 may require additional ports to be specified depending on which MFD is used. More details are available here.

Note VIP 5 is only required if DNS printer discovery is used for Mobility Print. If Layer 7 SNAT mode is used, UDP is not supported so DNS printer discovery only works using TCP.

Note VIP 6 is only required if mDNS printer discovery is used for Mobility Print. If Layer 7 SNAT mode is used, UDP is not supported so mDNS printer discovery cannot be used.

5.3. SSL Termination

When layer 7 SNAT mode is used, SSL Termination must be configured on the load balancer for the HTTPS Application Server traffic. As mentioned above, this enables `X-Forward-For` headers to be inserted.

In this guide, HAProxy is used to terminate SSL. The SSL termination is automatically configured when layer 7 VIP 1 is created. This adds an HTTPS listener on TCP port 9192 to the Virtual Service.

If required, VIP 1 can be configured to redirect all HTTP connections to HTTPS.

Note If HTTP connections are redirected to HTTPS on the load balancer, you'll need to configure the PaperCut Print Provider to support HTTPS as mentioned here under ***Configure PaperCut Application Server Failover > Additional details on load balancer configurations > SSL Configurations***. To configure this, see this article.

Certificates in PEM or PFX format can be uploaded to the load balancer.

6. Deployment Concept

6.1. When using Layer 4 DR Mode

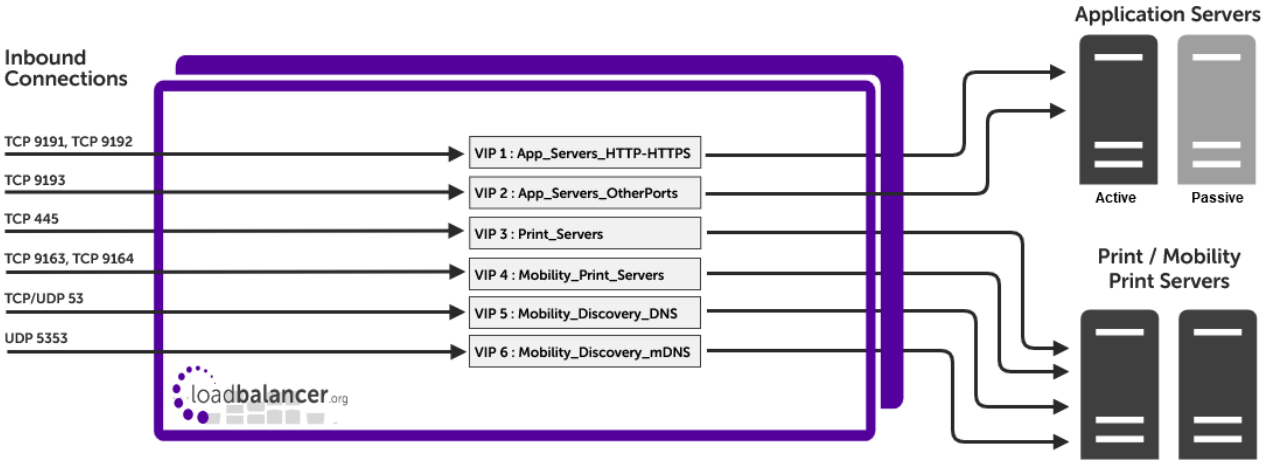

6.2. When using Layer 7 SNAT Mode

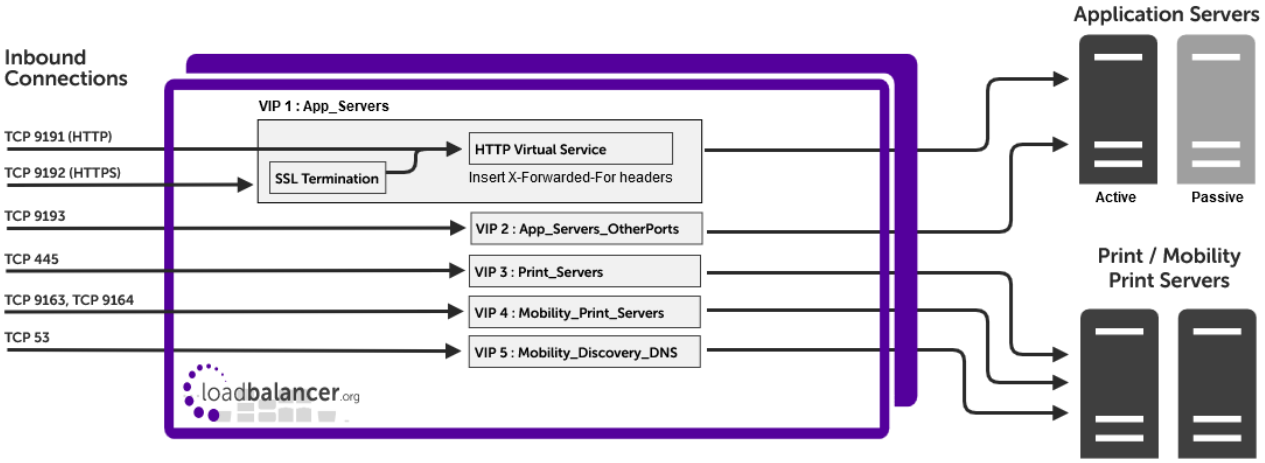

> **Note** The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to Configuring HA - Adding a Secondary Appliance for more details on configuring a clustered pair.

7. Load Balancer Deployment Methods

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 PaperCut, layer 4 DR mode is recommended. This mode offers the best possible performance since replies go directly from the Real Servers to the client, not via the load balancer. It is also transparent meaning that the client source IP address is maintained through to the load balanced servers.

If Layer 4 DR mode cannot be used due to Real Server or network topology reasons, then layer 7 SNAT mode is recommended.

> **Note** Layer 7 SNAT mode is compatible with the base features of PaperCut NG and MF, however, certain embedded applications may not be prepared to handle application-level load balancing as mentioned here under ***Configure PaperCut Application Server Failover > Additional details on***

load balancer configurations > Using a Layer 7 (Application Layer) load balancer.

> **Note** Since PaperCut NG/MF uses the originating IP address to help identify the calling device in some situations, the load balancer should be added as a trusted proxy as mentioned here.

> **Note** As mentioned above in Virtual Service (VIP) Requirements, if Layer 7 is used, mDNS is not supported.

These modes are described below and are used for the configurations presented in this guide.

For configuring using DR mode refer to Configure Load Balancing for PaperCut – Using Layer 4 DR Mode.

For configuring using layer 7 SNAT mode refer to Configure Load Balancing for PaperCut – Using Layer 7 SNAT Mode.

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*.

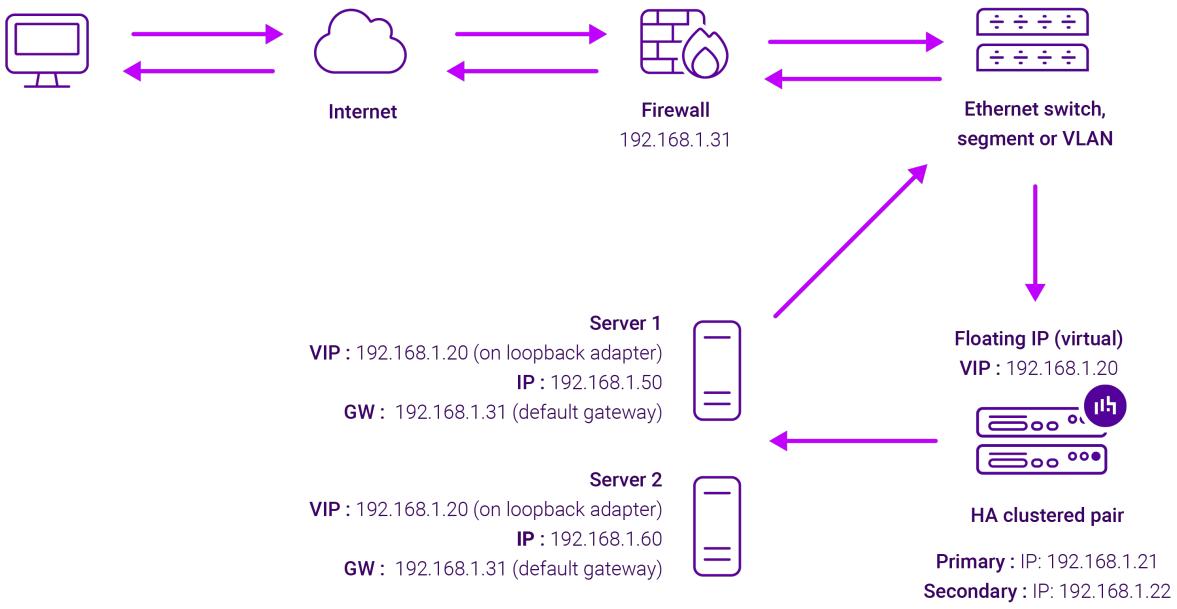

- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that each Real Server (and the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

7.2. Layer 7 SNAT Mode

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

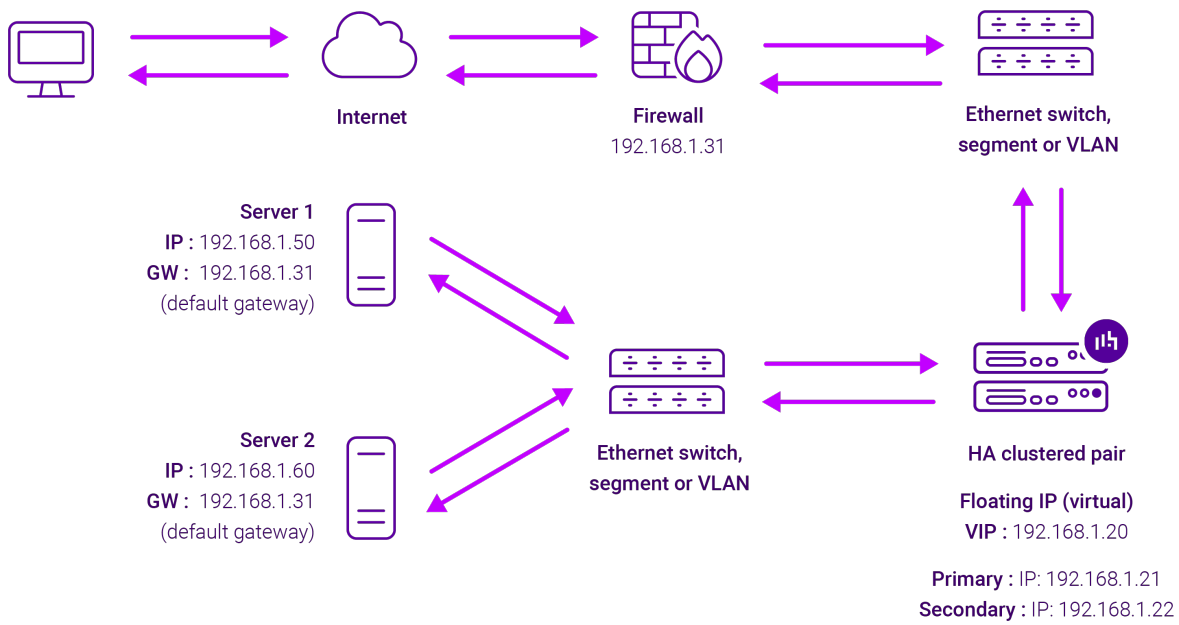

- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's own IP address by default, or any other local appliance IP address if preferred (e.g. the VIP address). This can be configured per layer 7 VIP. If required, the load balancer can be configured to provide the actual client IP address to the Real Servers in 2 ways. Either by inserting a

header that contains the client's source IP address, or by modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. For more information on these methods please refer to Transparency at Layer 7.

- Layer 7 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, **eth1** is typically used for client side connections and **eth0** is used for Real Server connections, although this is not mandatory since any interface can be used for any purpose.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

8. Configuring PaperCut for Load Balancing - Overview

8.1. Microsoft Print Servers

When load balancing Microsoft print servers, a number of additional configuration steps must be followed to allow them to be load balanced and accessed via a shared name. A number of registry changes are required to enable the load balanced print servers to be accessed via a shared name. Also, a DNS Host (A) record that points to the VIP address must be created. The hostname used must match the name used for the registry changes mentioned above.

For details, please refer to Microsoft Print Server Configuration.

8.2. Deployment Method

Depending on the load balancer deployment method used, certain additional configuration changes are required.

8.2.1. Layer 4 DR Mode

If layer 4 DR mode is used, the "ARP problem" must be solved on each load balanced Real Server. This enables DR mode to work correctly.

For more information on the "ARP problem", please refer to DR Mode Considerations in the Administration Manual.

For details on how to solve the "ARP problem", please refer to Solve the "ARP Problem".

8.2.2. Layer 7 SNAT Mode

If layer 7 SNAT mode is used, the load balancer should be added as a trusted Proxy as mentioned above to enable `X-Forwarded-For` headers to be used to extract the client source IP address.

For details, please refer to Add the Load Balancer as a Trusted Proxy.

8.3. Configure Papercut Components to Point at the VIPs on the Load Balancer

All PaperCut components must be configured to point at the load balancer rather than individual servers.

9. Loadbalancer.org Appliance – the Basics

9.1. 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.

9.2. 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.

9.3. 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:

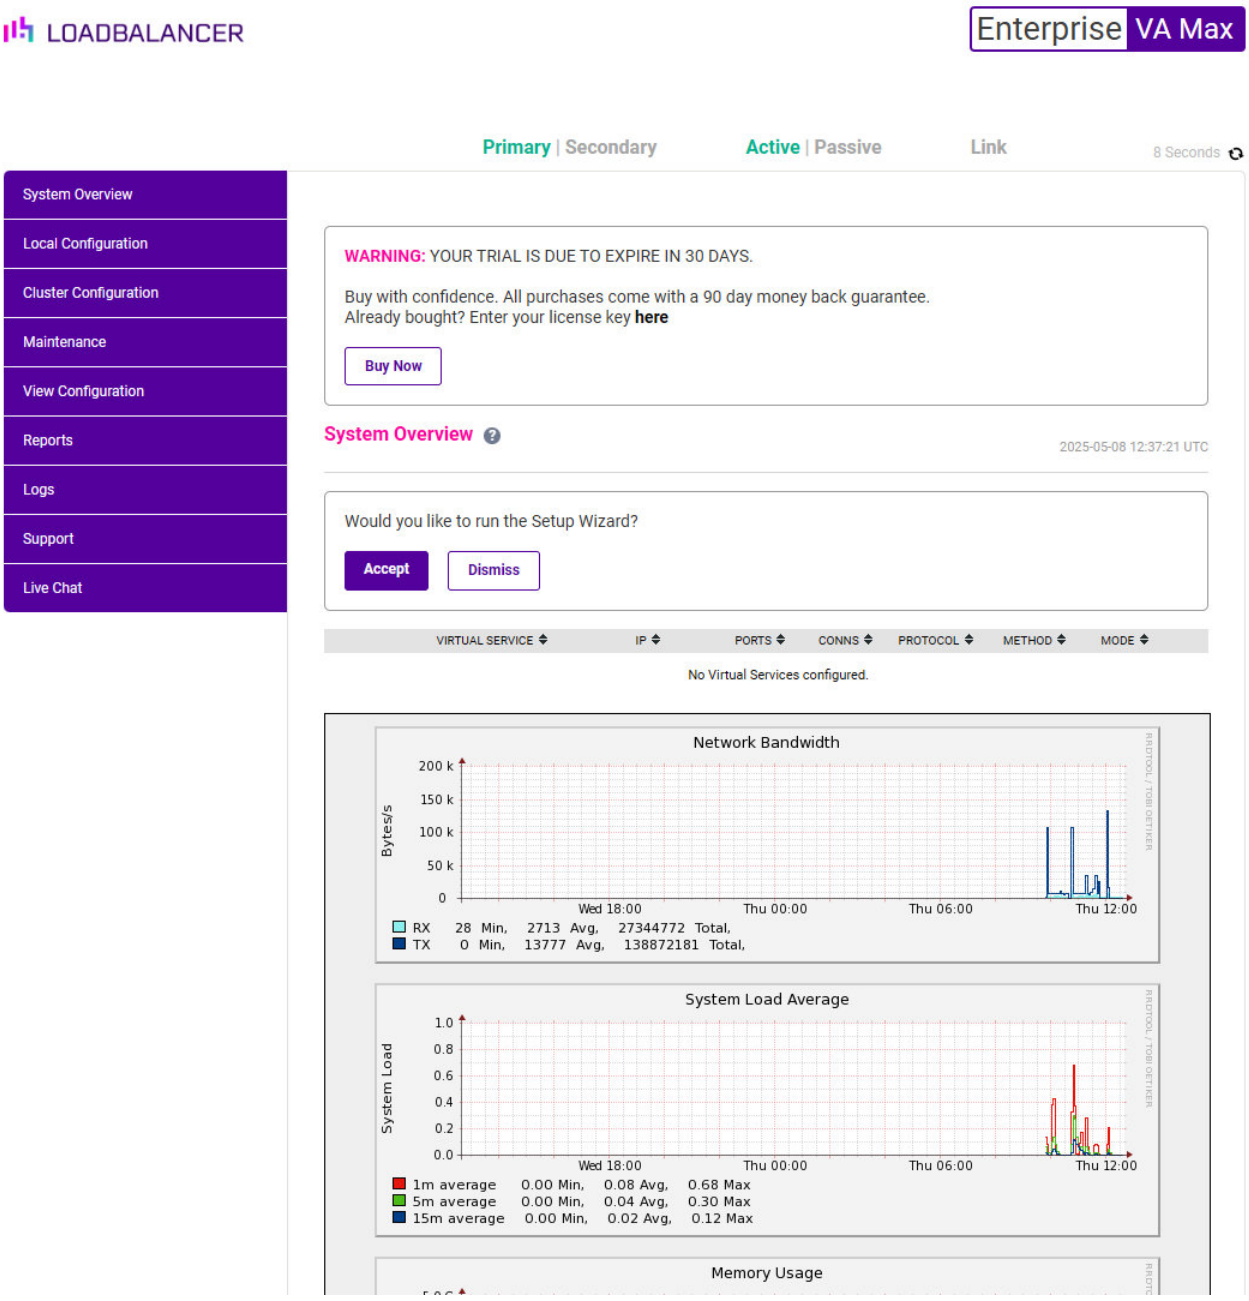

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

9.3.1. 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

9.4. 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.

9.4.1. 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:

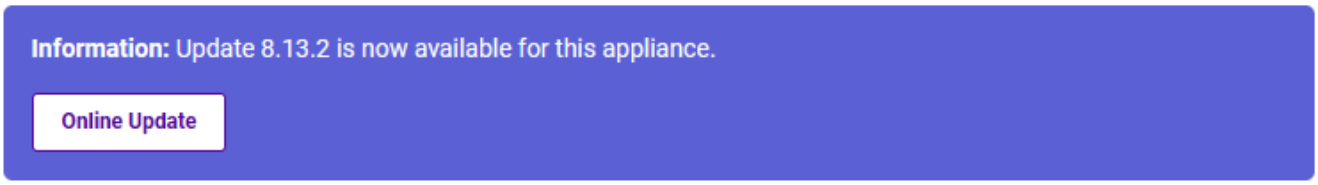

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

> **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:

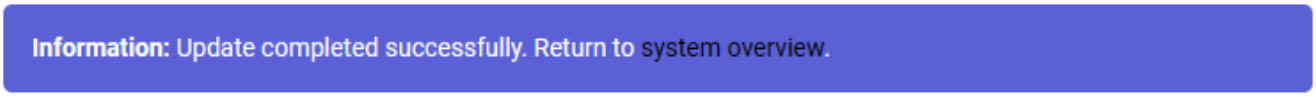

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

9.4.2. 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:

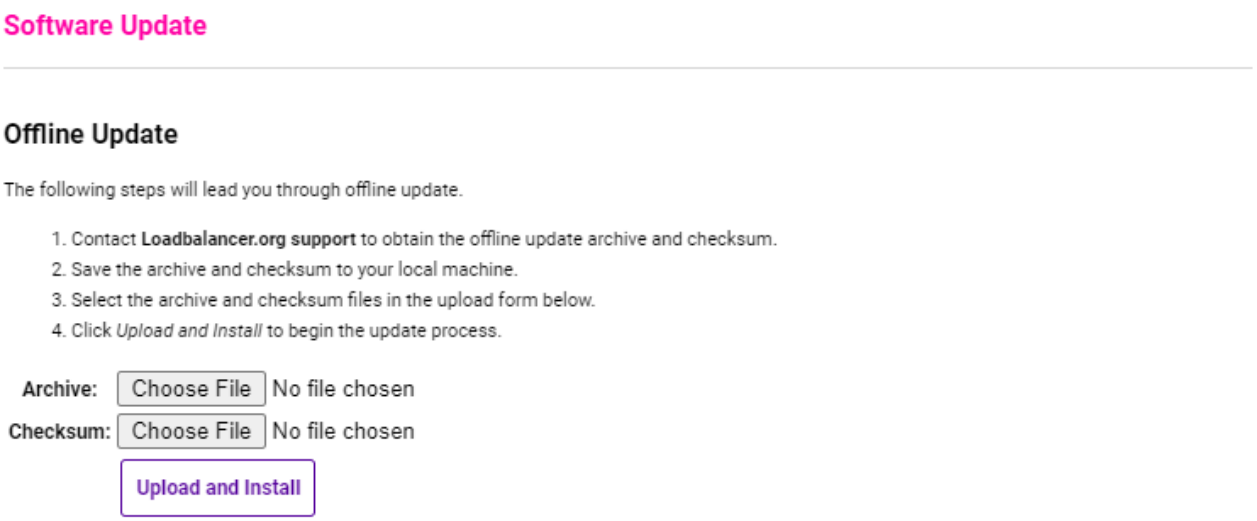

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

9.5. 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

Protocol	Port	Purpose
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.

9.6. 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.

10. Configure Load Balancing for PaperCut – Using Layer 4 DR Mode

10.1. Appliance Configuration

10.1.1. Configure the Health Check Interval

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Advanced Configuration*.
2. Set the *Check interval* to **30**, i.e. 30 seconds.

> **Note** This is the interval recommended by Papercut as mentioned here under *Configure PaperCut Application Server Failover > Configure NLB Health Checks*.

3. Click **Update**.

10.1.2. VIP 1 – App_Servers_HTTP-HTTPS

Configure The Virtual Service (VIP)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.

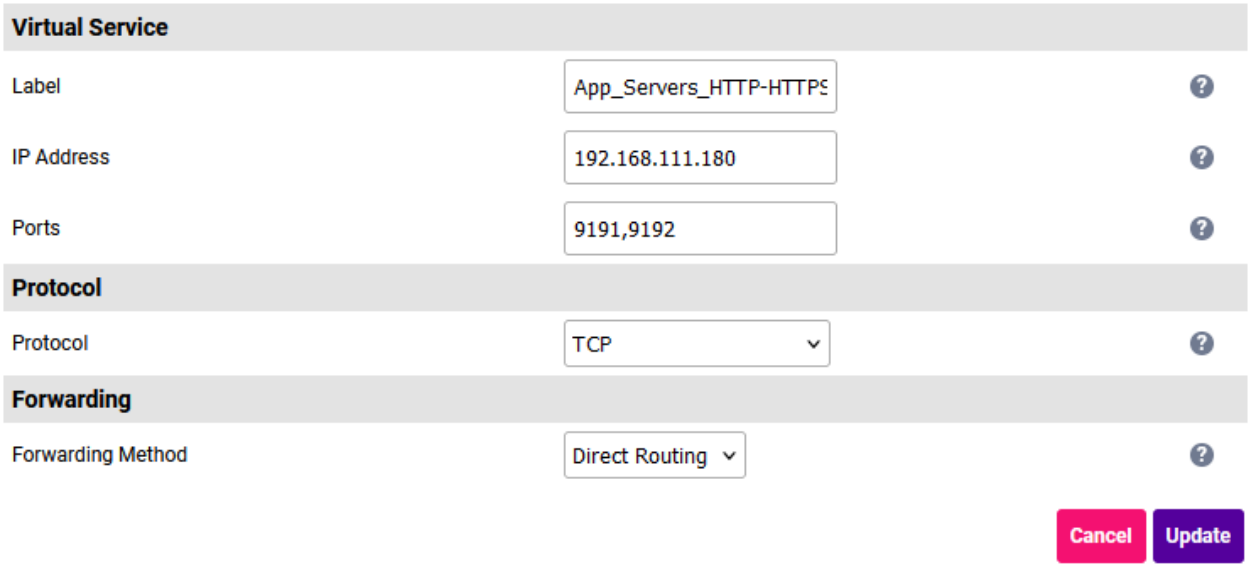

2. Specify the required *Label* (name) for the VIP, e.g. **App_Servers_HTTP-HTTPS**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **9191,9192**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the Virtual Service.
8. Now click **Modify** next to the newly created Virtual Service.
9. Scroll to the *Persistence* section.
   - Disable *Persistence* by unchecking the *Enable* check box.
10. Scroll to the *Health Checks* section.
    - Set the *Check Type* to **Negotiate** and the *Protocol* to **HTTP**.
    - Set the ***Request to send*** to the following value substituting <AUTHORIZATION KEY> with authorization key that can be found in the HTTP header:

      ```
      /api/health/application-server/status?disk-threshold-mb=1&Authorization=<AUTHORIZATION KEY>
      ```

      The HTTP header can be found on the Application Server under *Options > Advanced* in the *System Health Monitoring* section.

      HTTP header

      ```
      Authorization:JjtxY8ztIIZhAOKtGHs2swxw7Q3eyVXH
      ```

11. Click **Update**.

Define the Associated Real Servers (RIPs)

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

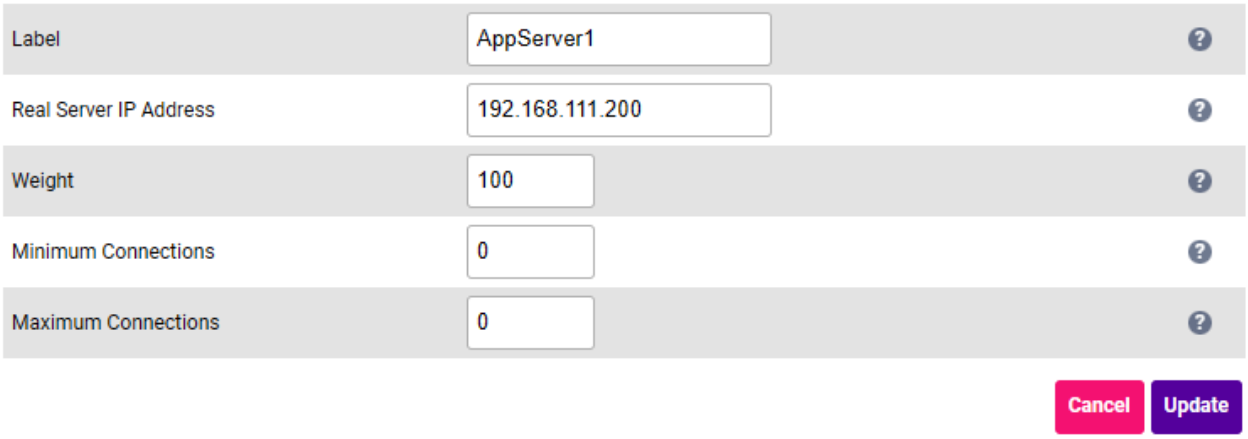

3. Specify an appropriate label for the Real Server, e.g. **AppServer1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.200**.
5. Click **Update**.
6. Repeat the above steps to add the remaining (passive) Application Server(s).

10.1.3. VIP 2 – App_Servers_OtherPorts

Configure The Virtual Service (VIP)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.

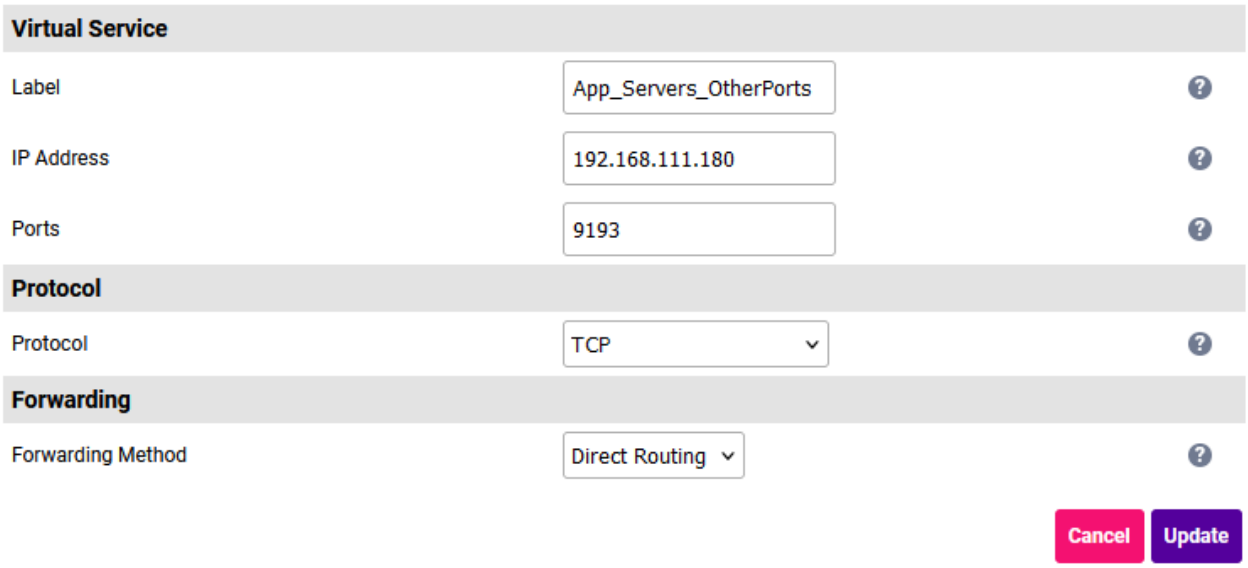

2. Specify the required *Label* (name) for the VIP, e.g. **App_Servers_OtherPorts**.

3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **9193**.

> **Note** Additional ports may need to be specified depending in which MFD is used. More details are available here.

5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the Virtual Service.
8. Now click **Modify** next to the newly created Virtual Service.
9. Scroll to the *Persistence* section.
   - Disable *Persistence* by unchecking the *Enable* check box.
10. Click **Update**.

Define the Associated Real Servers (RIPs)

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

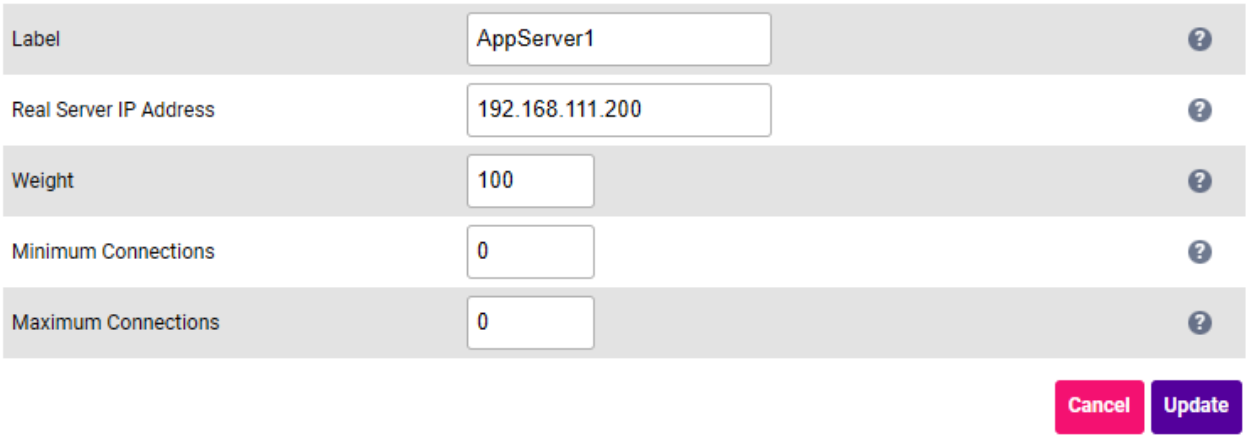

3. Specify an appropriate label for the Real Server, e.g. **AppServer1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.200**.
5. Click **Update**.
6. Repeat the above steps to add the remaining (passive) Application Server(s).

10.1.4. VIP 3 – Print_Servers

Configure The Virtual Service (VIP)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.

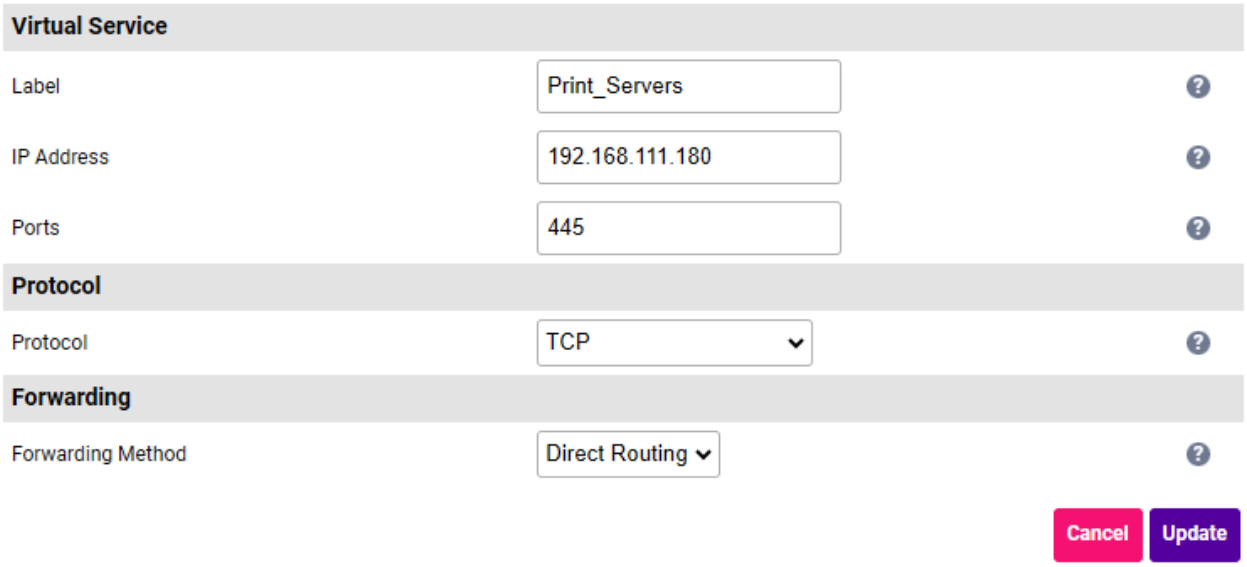

2. Specify the required *Label* (name) for the VIP, e.g. **Print_Servers**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **445**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the Virtual Service.

Define the Associated Real Servers (RIPs)

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

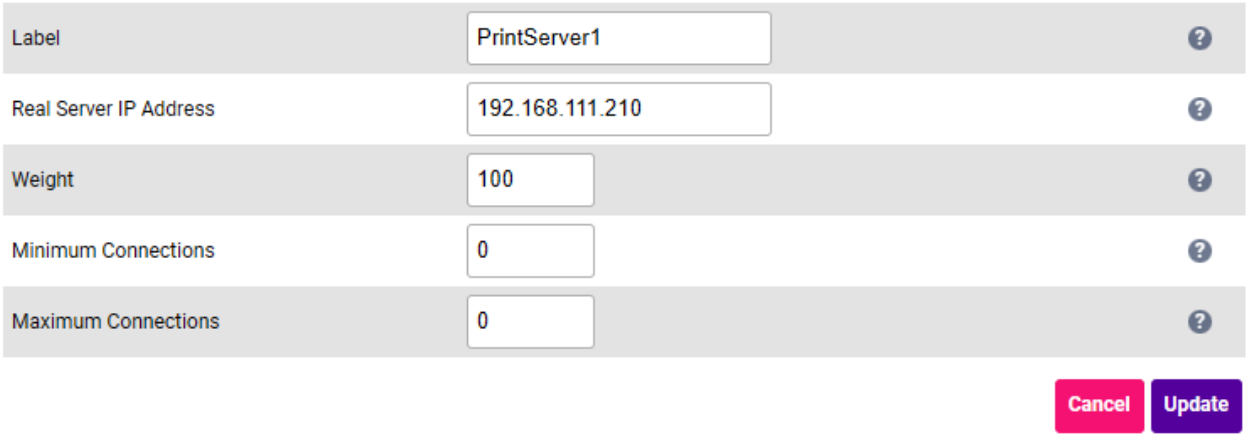

3. Specify an appropriate label for the Real Server, e.g. **PrintServer1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Click **Update**.

6. Repeat the above steps to add additional Print Server(s).

10.1.5. VIP 4 – Mobility_Print_Servers

Configure The Virtual Service (VIP)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.

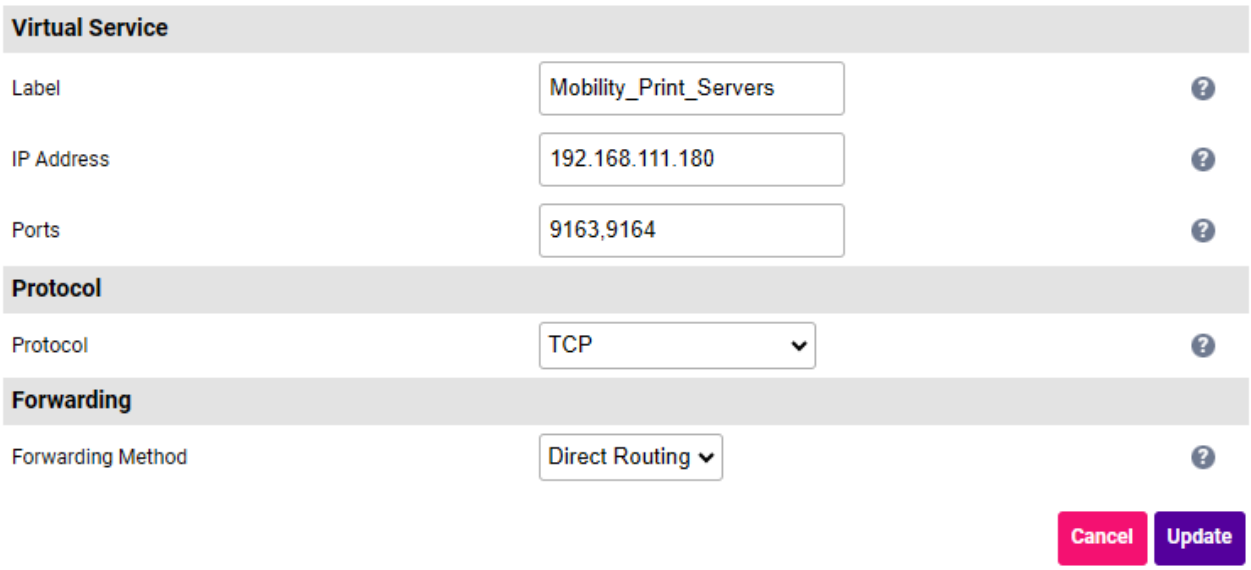

2. Specify the required *Label* (name) for the VIP, e.g. **Mobility_Print_Servers**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **9163,9164**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the Virtual Service.
8. Now click **Modify** next to the newly created Virtual Service.
9. Scroll to the *Health Checks* section.
    - Set the *Check Type* to **Negotiate** and the *Protocol* to **HTTP**.
    - Set the *Check Port* to **9163**.
    - Set the *Request to send* to **/health**.
10. Click **Update**.

Define the Associated Real Servers (RIPs)

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

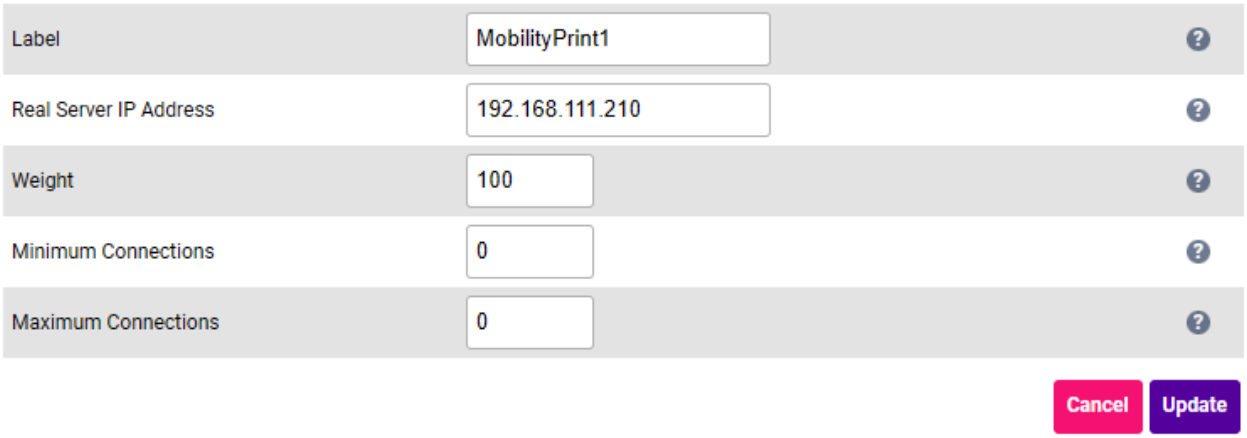

3. Specify an appropriate label for the Real Server, e.g. **MobilityPrint1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Click **Update**.
6. Repeat the above steps to add additional Mobility Print Server(s).

10.1.6. VIP 5 – Mobility_Discovery_DNS

> **Note** VIP 5 is only required if DNS printer discovery is used for Mobility Print.

> **Note** By default, the GSLB service running on the load balancer is bound to all IPs on port 53. This must be changed to a specific IP address so that port 53 can be specified for VIP 4. This can be done using the WebUI menu option: *Local Configuration > Physical - Advanced configuration >*, scrolling to the *Service Socket Addresses* section and setting the GSLB service to one of the interface IP addresses.

Configure The Virtual Service (VIP)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.

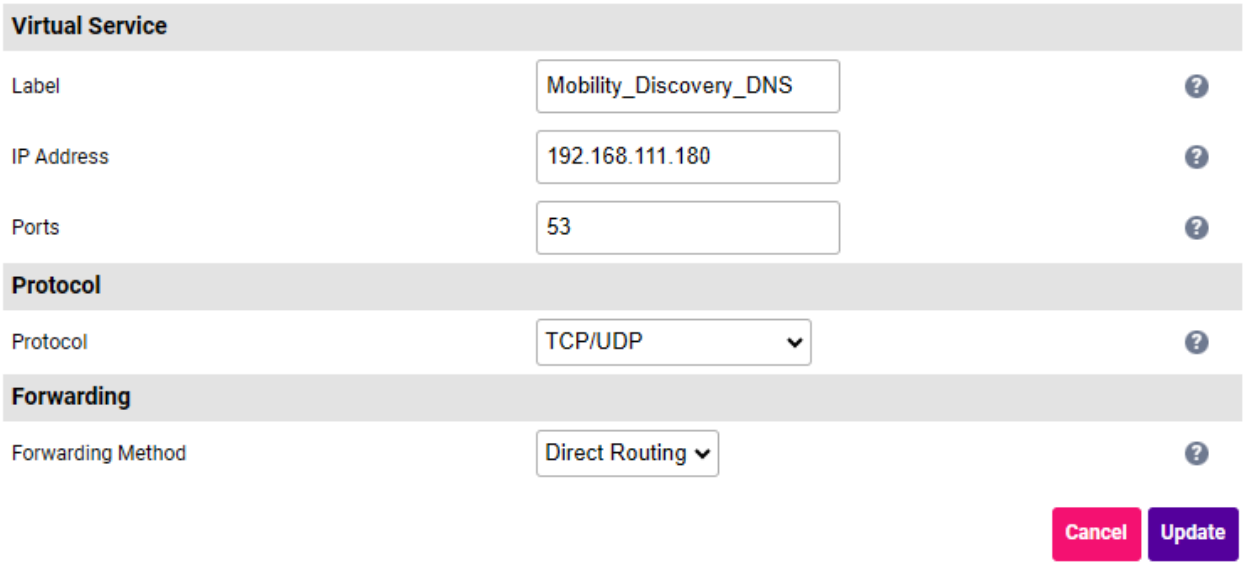

2. Specify the required *Label* (name) for the VIP, e.g. **Mobility_Discovery_DNS**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **53**.
5. Leave the *Protocol* set to **TCP/UDP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the Virtual Service.

Define the Associated Real Servers (RIPs)

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

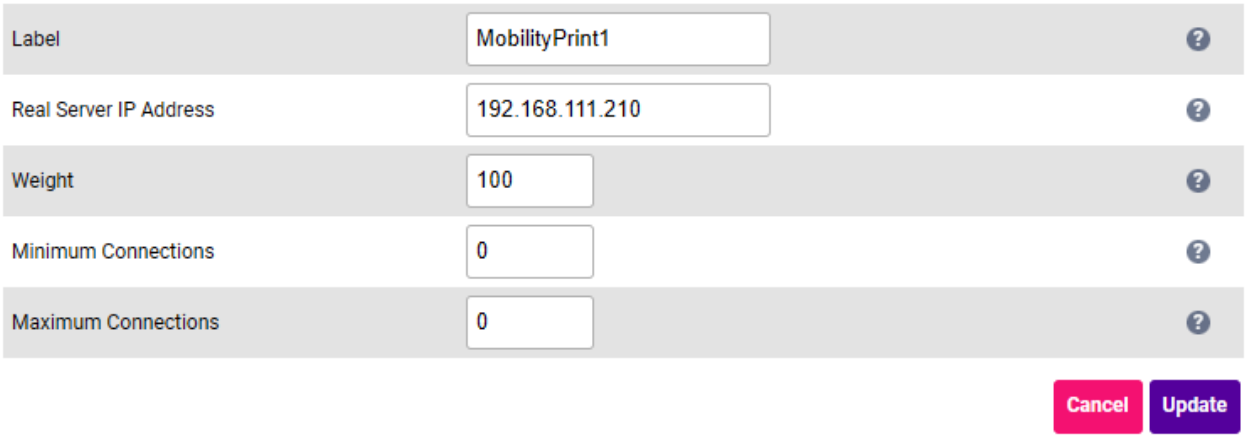

3. Specify an appropriate label for the Real Server, e.g. **MobilityPrint1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Click **Update**.
6. Repeat the above steps to add additional Mobility Print Server(s).

10.1.7. VIP 6 – Mobility_Discovery_mDNS

> **Note** VIP 6 is only required if mDNS printer discovery is used for Mobility Print.

Configure The Virtual Service (VIP)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.

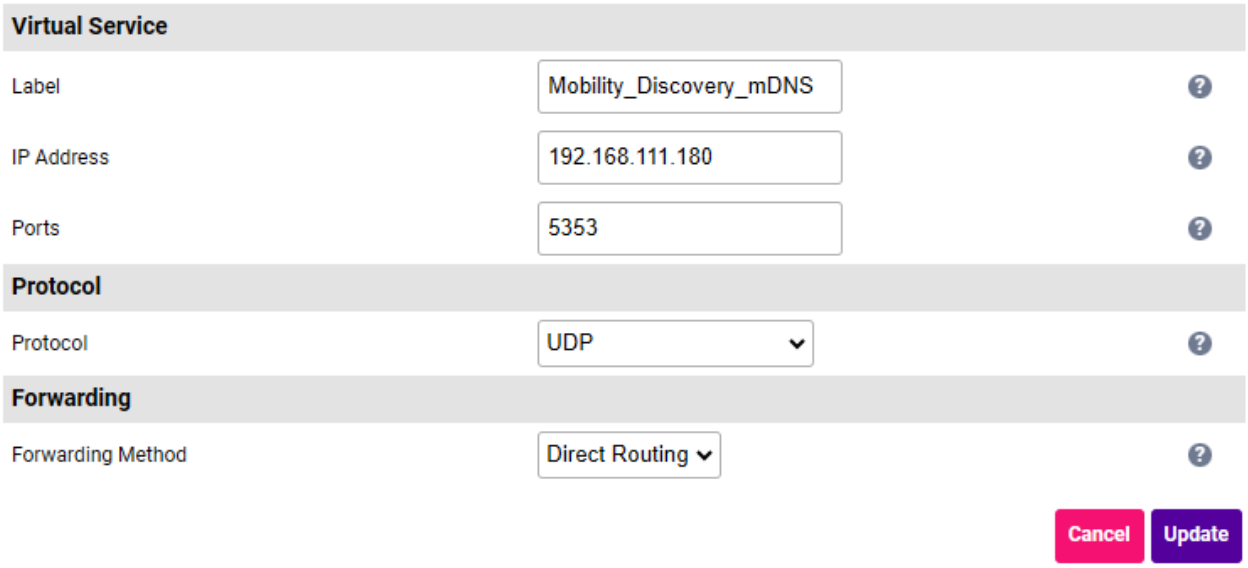

2. Specify the required *Label* (name) for the VIP, e.g. **Mobility_Discovery_mDNS**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **5353**.
5. Leave the *Protocol* set to **UDP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the Virtual Service.

Define the Associated Real Servers (RIPs)

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

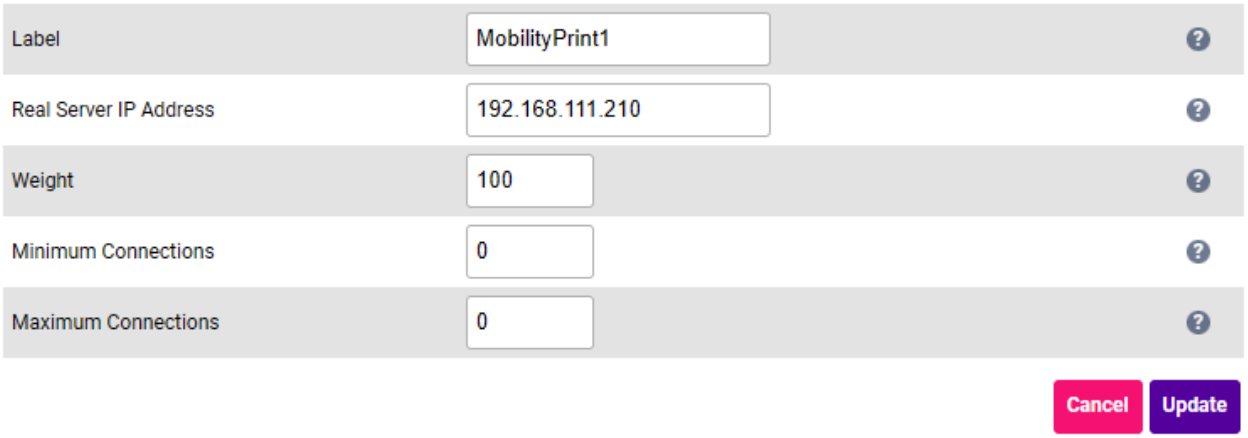

3. Specify an appropriate label for the Real Server, e.g. **MobilityPrint1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Click **Update**.

6. Repeat the above steps to add additional Mobility Print Server(s).

10.2. Papercut Configuration

10.2.1. Configure Papercut Components to Point at the VIPs on the Load Balancer

All PaperCut components must be configured to point at the Virtual Services (VIPs) on the load balancer rather than individual servers.

10.2.2. Solve the "ARP Problem"

When using layer 4 DR mode, the "ARP Problem" must be solved on each load balanced Real Server to enable DR mode to work correctly. The exact steps required depend on the particular operating system in use. The section below detail the steps for Windows 2012 & later. For other operating systems, please refer to DR Mode Considerations in the Administration Manual.

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**.

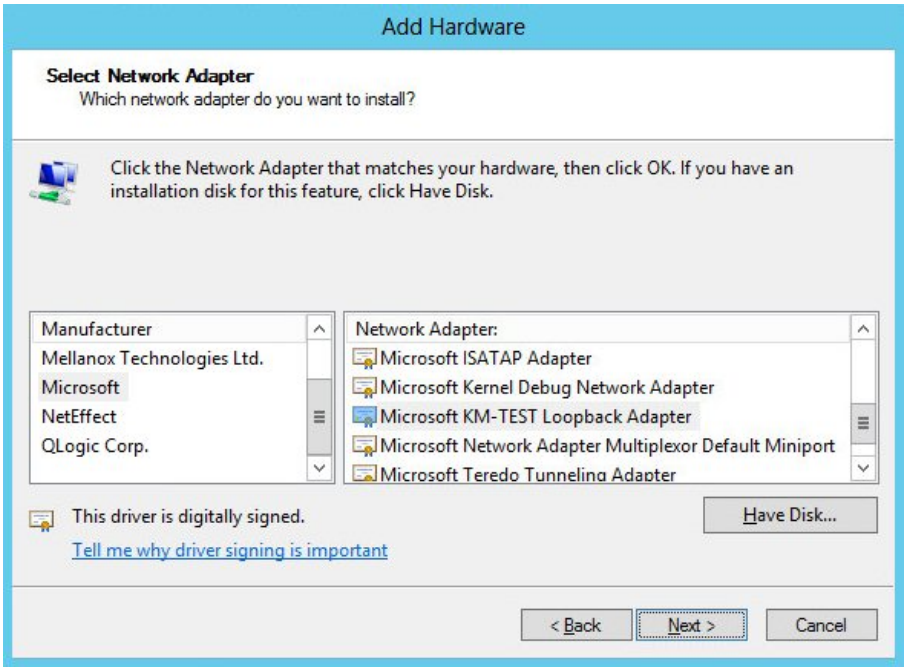

5. Select **Microsoft** & **Microsoft KM-Test Loopback Adapter**, click **Next**.
6. Click **Next** to start the installation, when complete click **Finish**.

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.

> **Important** When configuring the loopback adapter properties, make sure that **Client for Microsoft Networks** and **File & Printer Sharing for Microsoft Networks** is also checked as shown below.

IPv4 Addresses

1. Uncheck all items except **Client for Microsoft Networks**, **File & Printer Sharing for Microsoft Networks** and **Internet Protocol Version 4 (TCP/IPv4)** as shown below:

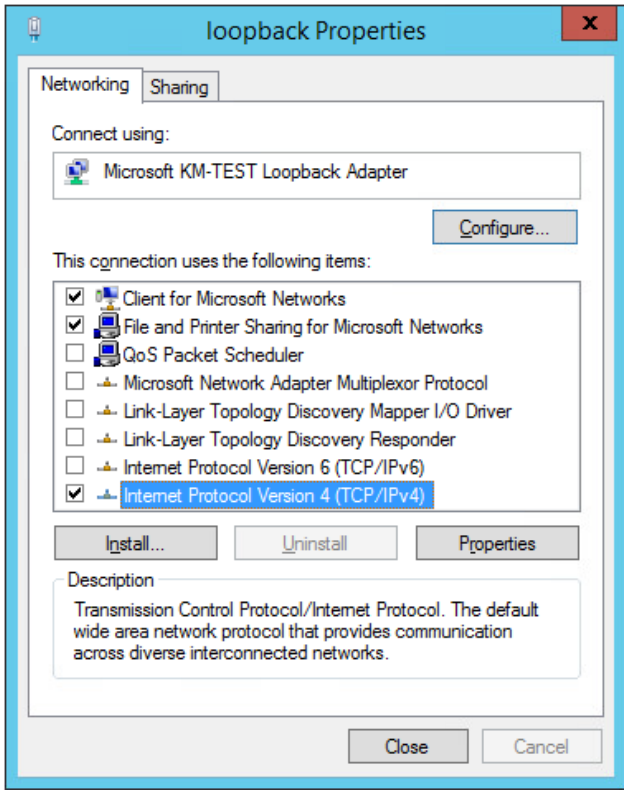

2. Ensure that **Internet Protocol Version (TCP/IPv4)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service address (VIP) with a subnet mask of **255.255.255.255**, e.g. **192.168.2.20/255.255.255.255** as shown below:

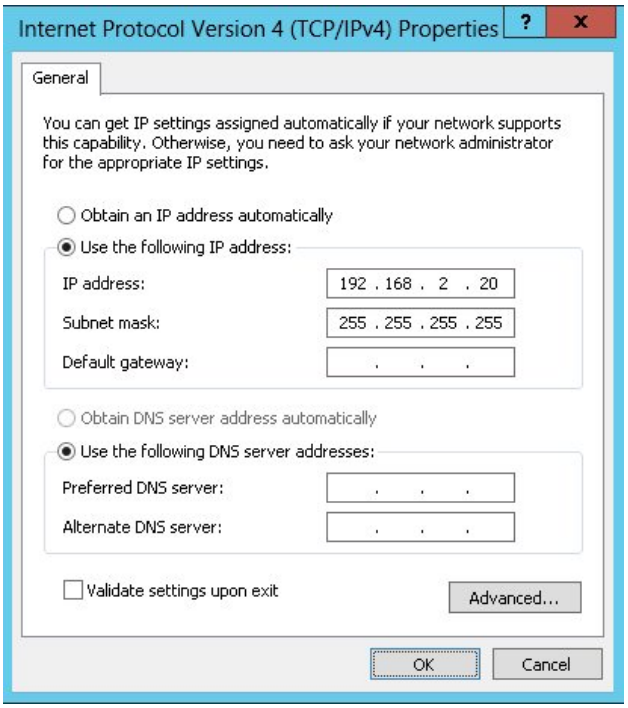

> **Note** **192.168.2.20** is an example, make sure you specify the correct VIP address.

> **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 **Client for Microsoft Networks**, **File & Printer Sharing for Microsoft Networks** and **Internet Protocol Version 6 (TCP/IPv6)** as shown below:

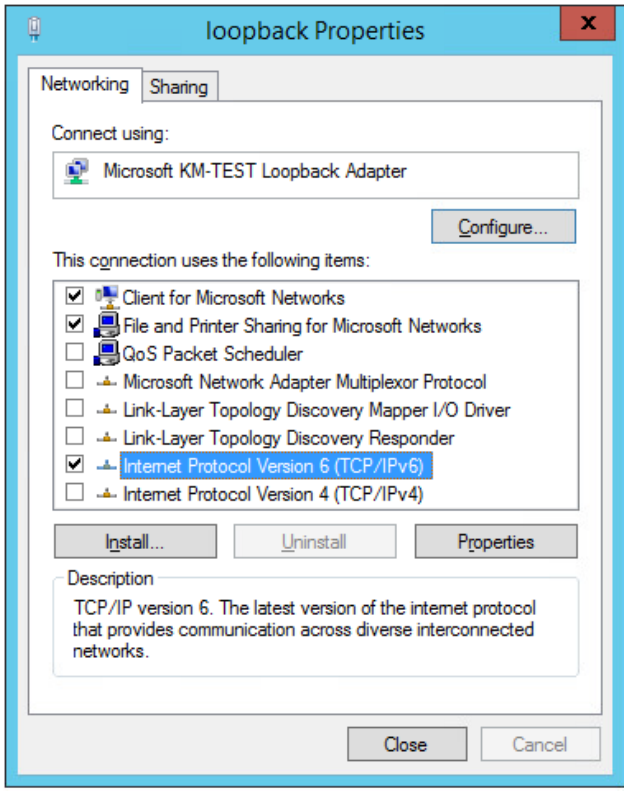

2. Ensure that **Internet Protocol Version (TCP/IPv6)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) and set the *Subnet Prefix Length* to be the same as your network setting, e.g. **2001:470:1f09:e72::15/64** as shown below:

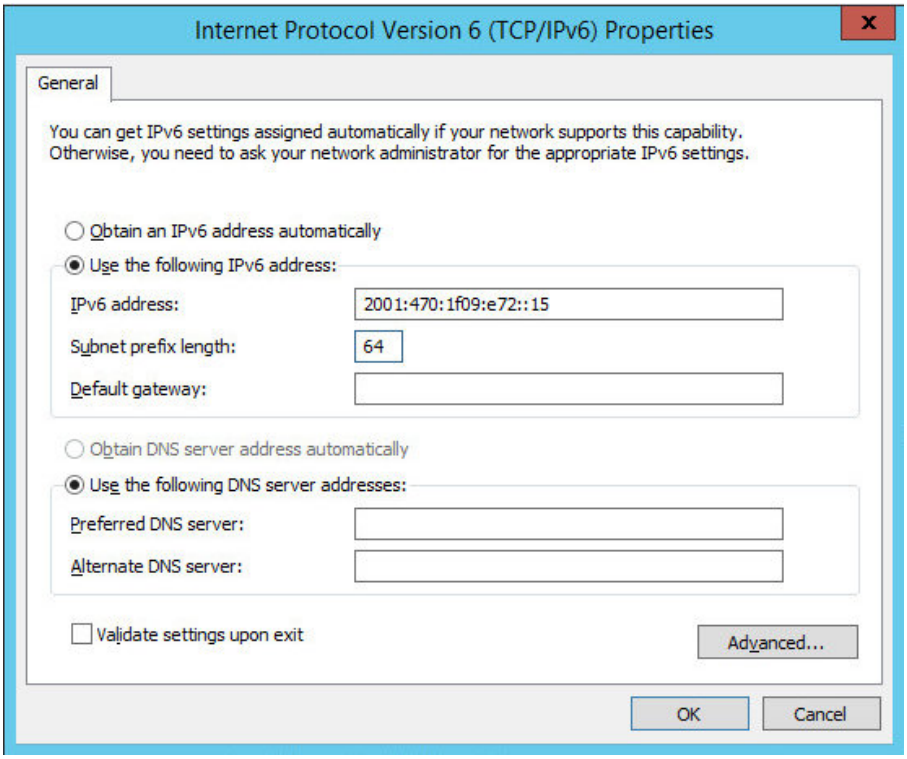

> **Note** **2001:470:1f09:e72::15/64** is an example, make sure you specify the correct VIP address.

> **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
```

10.3. Configure Microsoft Print Servers

If Microsoft print servers are used, follow the steps in Microsoft Print Server Configuration.

11. Configure Load Balancing for PaperCut – Using Layer 7 SNAT Mode

11.1. Appliance Configuration

11.1.1. Configure the Health Check Interval

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Advanced Configuration*.
2. Set the *Interval* to **30000**, i.e. 30 seconds.

> **Note** This is the interval recommended by Papercut as mentioned here under *Configure PaperCut Application Server Failover > Configure NLB Health Checks*.

3. Click **Update**.

11.1.2. Upload the SSL Certificate for use with the SSL Termination

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

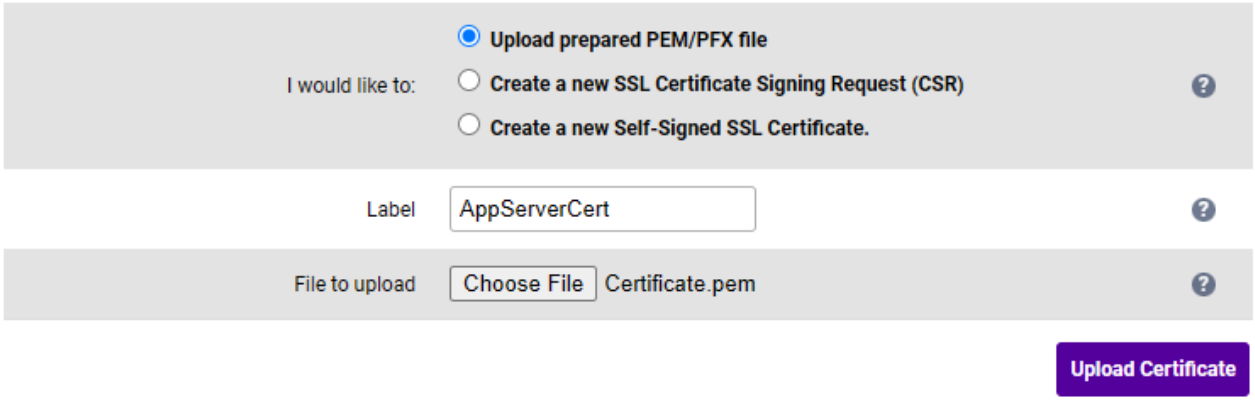

   - Specify an appropriate *Label*, e.g. **AppServerCert**.
   - Click **Choose File**.
   - Browse to and select the relevant PEM or PFX file.
   - For PFX files specify the password if required.

4. Click **Upload Certificate**.

11.1.3. VIP 1 – App_Servers

Configure The Virtual Service (VIP) & SSL Termination

This step creates a Virtual Service that listens on HTTP port 9191 and a SSL termination for HTTPS traffic on port 9192.

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

3. In the *Termination* section, enable (check) the *Create HAProxy SSL Termination* checkbox.

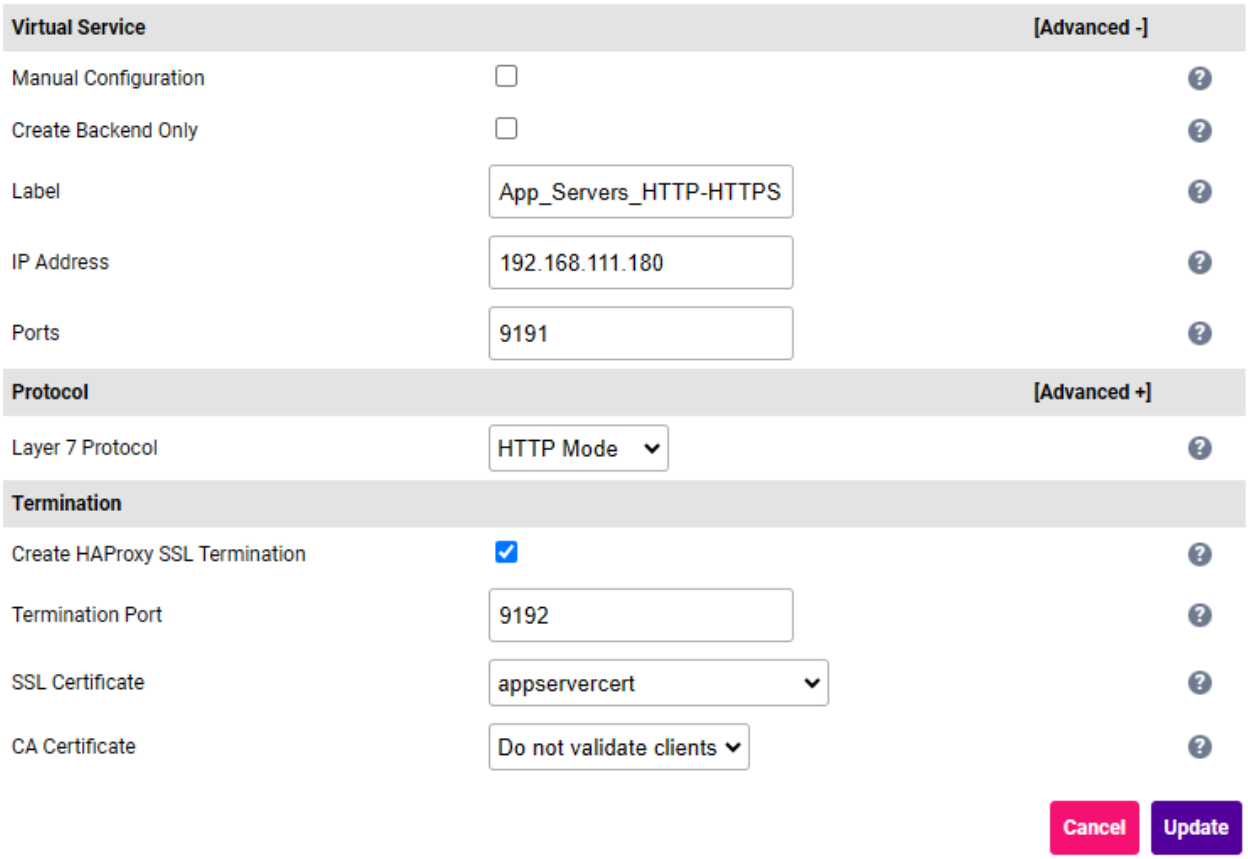

4. Specify the required *Label* (name) for the VIP, e.g. **App_Servers_HTTP-HTTPS**.
5. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
6. Set the *Ports* field to **9191**.
7. Leave the *Protocol* set to **HTTP Mode**.
8. In the *Termination* section.
9. Set the *Termination Port* to **9192**.
10. Select the *SSL Certificate* uploaded previously.
11. Click **Update** to create the Virtual Service.
12. Click **Modify** next to the newly created Virtual Service.
13. Scroll to the *Persistence* section.
    - Set the *Persistence Mode* to **None**.
14. Scroll to the *Health Checks* section.
    - Set the *Check Type* to **Negotiate HTTP (GET)**.
15. Set the *Request to send* to the following value substituting <AUTHORIZATION KEY> with authorization key that can be found in the HTTP header:

```
/api/health/application-server/status?disk-threshold-mb=1&Authorization=<AUTHORIZATION KEY>
```

The HTTP header can be found on the Application Server under *Options > Advanced* in the *System Health Monitoring* section.

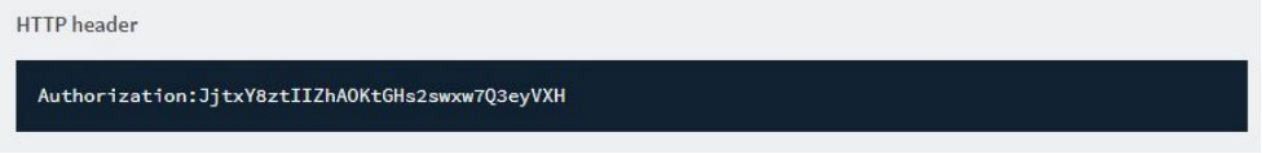

16. Click **Update**.

Define the Associated Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

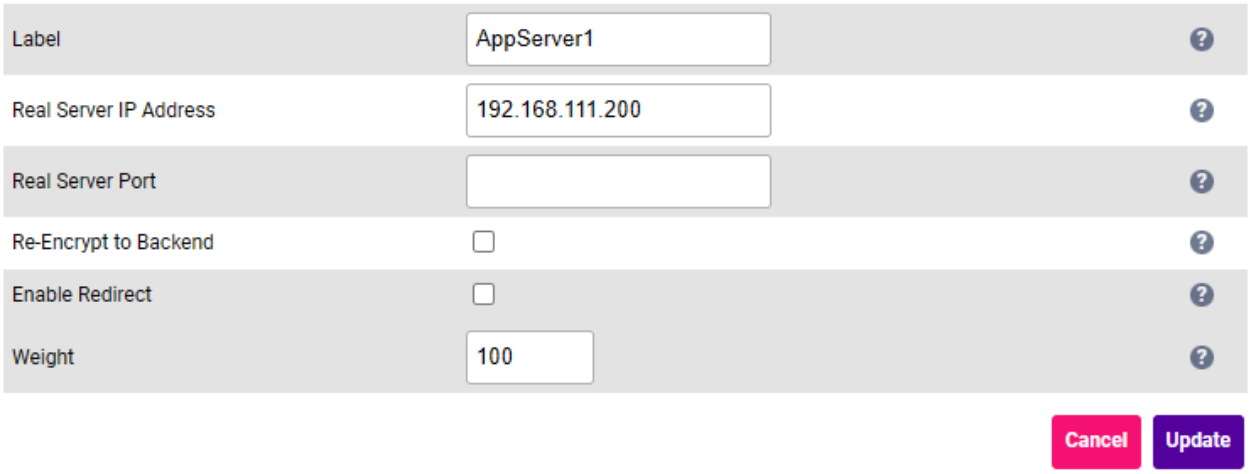

3. Specify an appropriate *Label* for the Real Server, e.g. **AppServer1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.200**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps for the remaining (passive) Application Server(s).

11.1.4. VIP 2 – App_Servers_OtherPorts

Configure The Virtual Service (VIP)

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

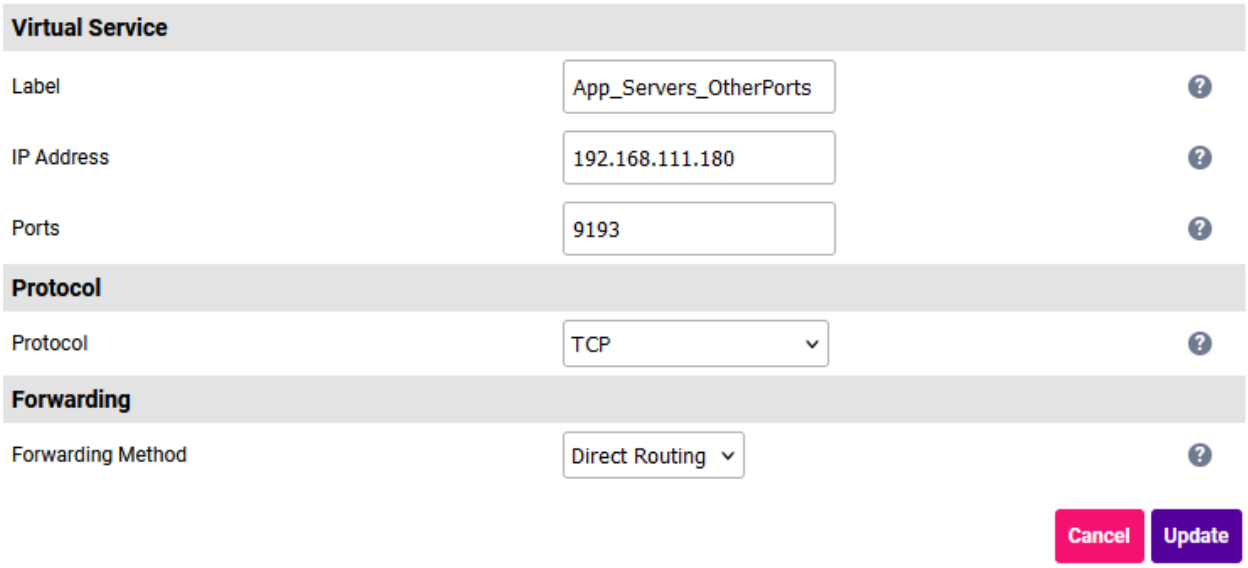

2. Specify the required *Label* (name) for the VIP, e.g. **App_Servers_OtherPorts**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **9193**.

> **Note** Additional ports may need to be specified depending in which MFD is used. More details are available here.

5. Set the *Protocol* to **TCP** Mode.
6. Click **Update** to create the Virtual Service.
7. Click **Modify** next to the newly created Virtual Service.
8. Scroll to the *Persistence* section.
   - Set the *Persistence Mode* to **None**.
9. Click **Update**.

Define the Associated Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

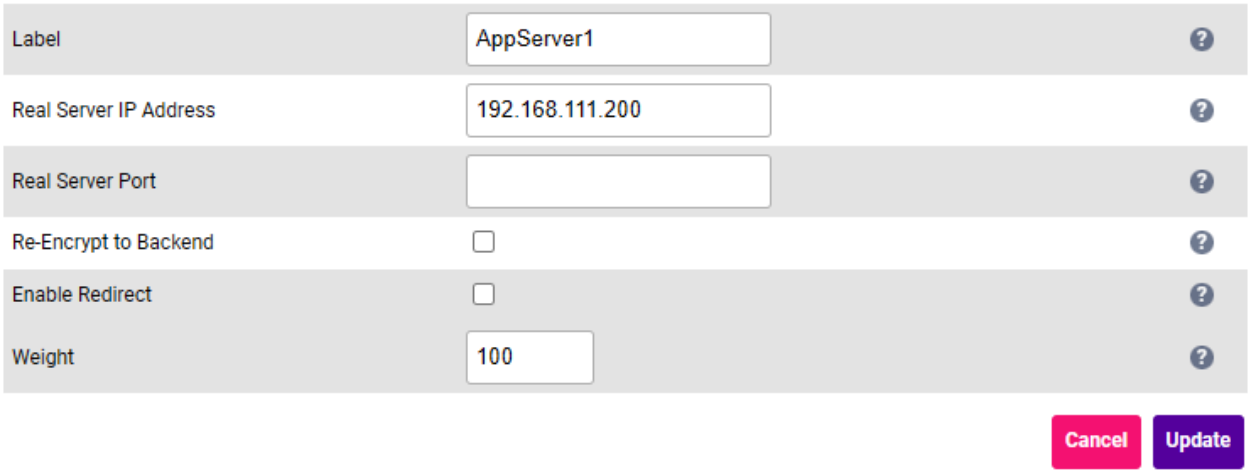

3. Specify an appropriate *Label* for the Real Server, e.g. **AppServer1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.200**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps for the remaining (passive) Application Server(s).

11.1.5. VIP 3 – Print_Servers

Configure the Virtual Service (VIP)

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

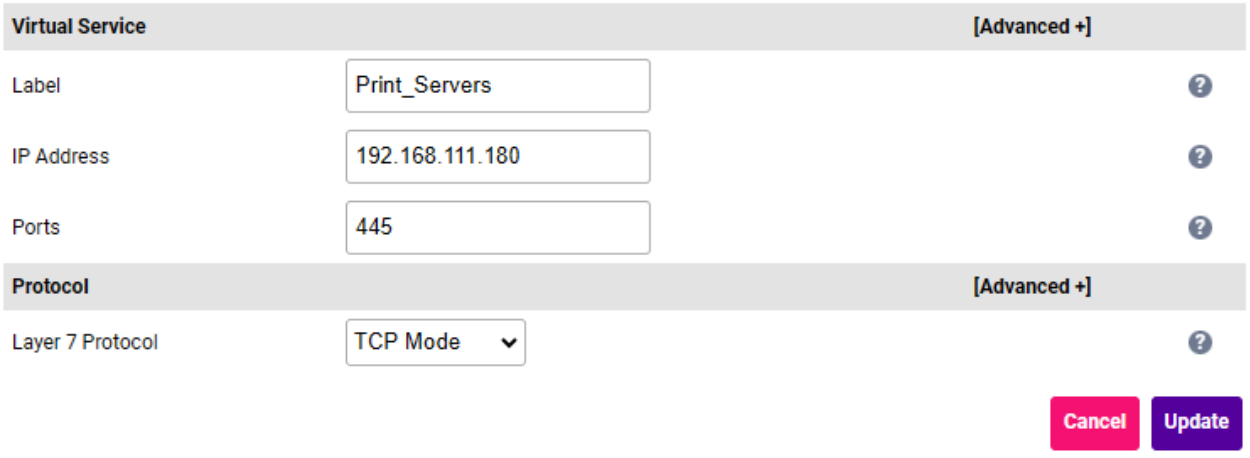

2. Define the required *Label* (name) for the VIP, e.g. **Print_Servers**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **445**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update** to create the Virtual Service.

Define the Associated Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

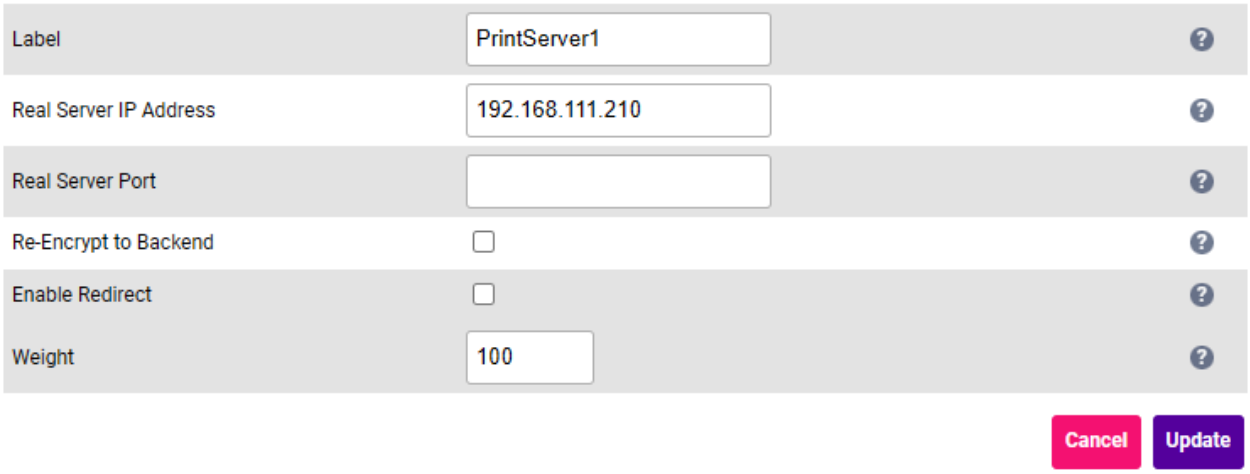

3. Specify an appropriate *Label* for the Real Server, e.g. **PrintServer1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add additional Print Server(s).

11.1.6. VIP 4 – Mobility_Print_Servers

Configuring The Virtual Service (VIP)

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

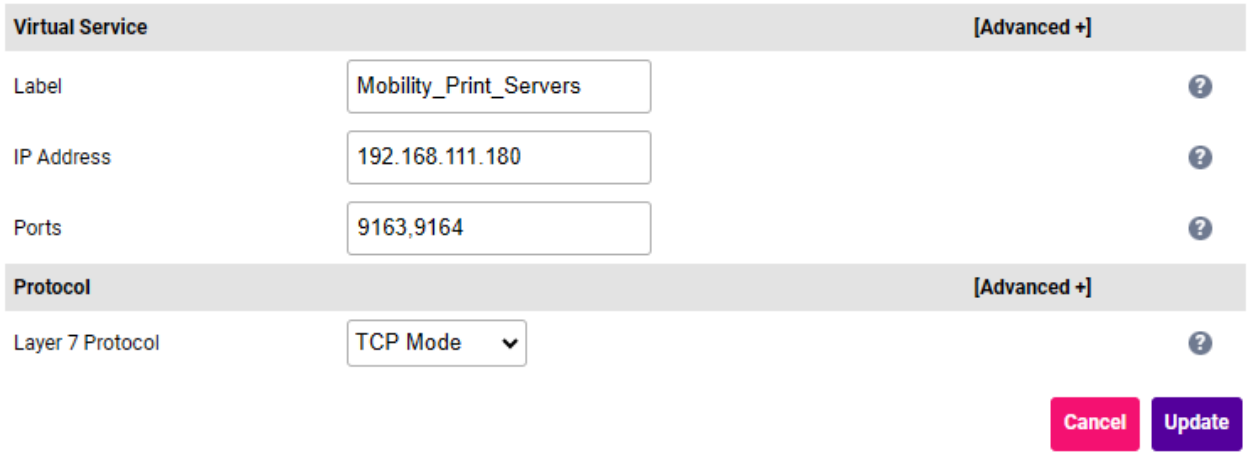

2. Define the required *Label* (name) for the VIP, e.g. **Mobility_Print_Servers**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.

4. Set the *Ports* field to **9163,9164**.
5. Set the *Protocol* to **TCP Mode**.
6. Click **Update** to create the Virtual Service.

Define the Associated Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

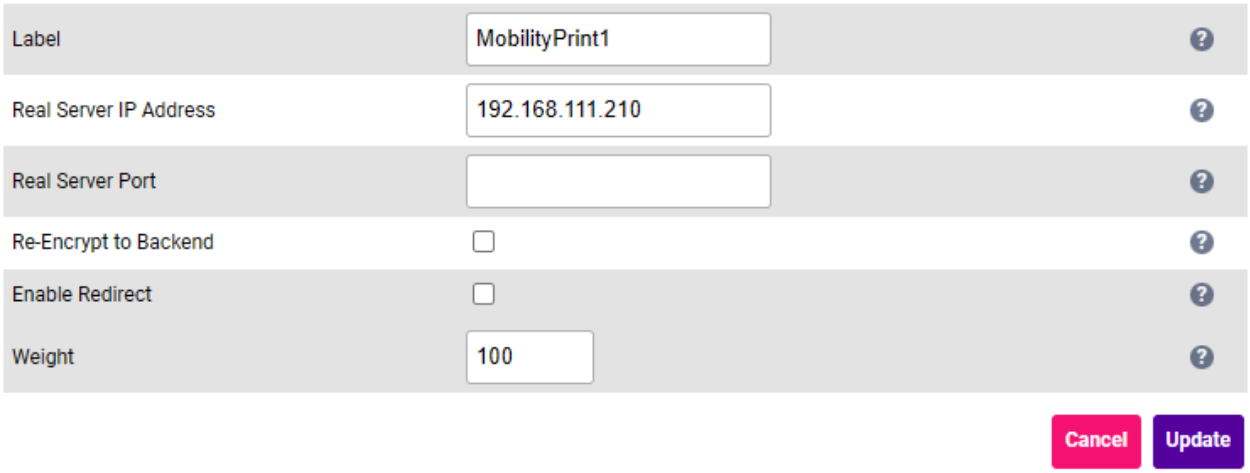

3. Specify an appropriate *Label* for the Real Server, e.g. **MobilityPrint1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add additional Mobility Print Server(s).

11.1.7. VIP 5 – Mobility_Discovery_DNS

> **Note** VIP 5 is only required if DNS printer discovery is used for Mobility Print.

> **Note** By default, the GSLB service running on the load balancer is bound to all IPs on port 53. This must be changed to a specific IP address so that port 53 can be specified for VIP 4. This can be done using the WebUI menu option: *Local Configuration > Physical - Advanced configuration >*, scrolling to the *Service Socket Addresses* section and setting the GSLB service to one of the interface IP addresses.

Configuring The Virtual Service (VIP)

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

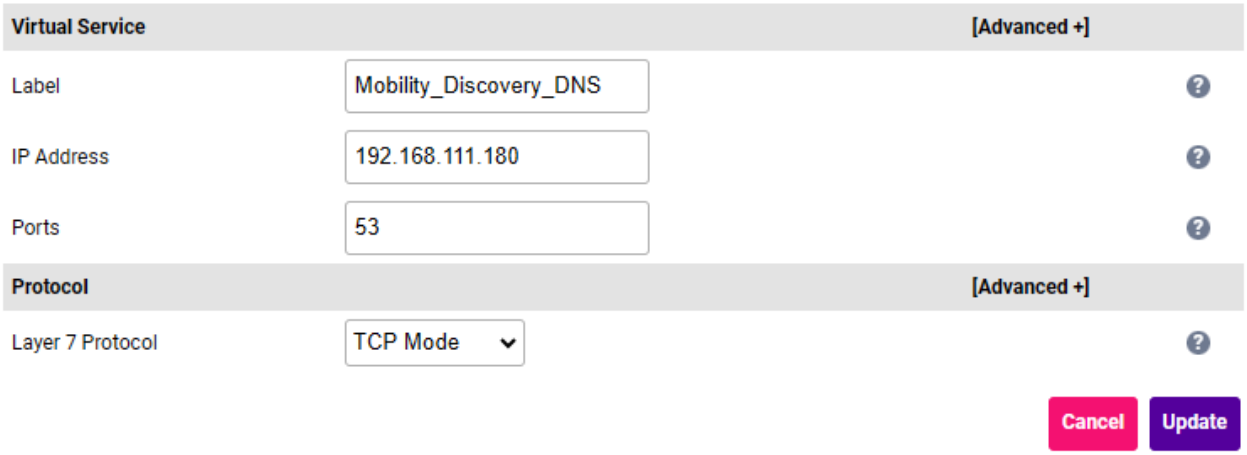

2. Define the required *Label* (name) for the VIP, e.g. **Mobility_Discovery_DNS**.
3. Set the *IP Address* field to the required IP address, e.g. **192.168.111.180**.
4. Set the *Ports* field to **53**.
5. Set the *Protocol* to **TCP Mode**.
6. Click **Update** to create the Virtual Service.

Define the Associated Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

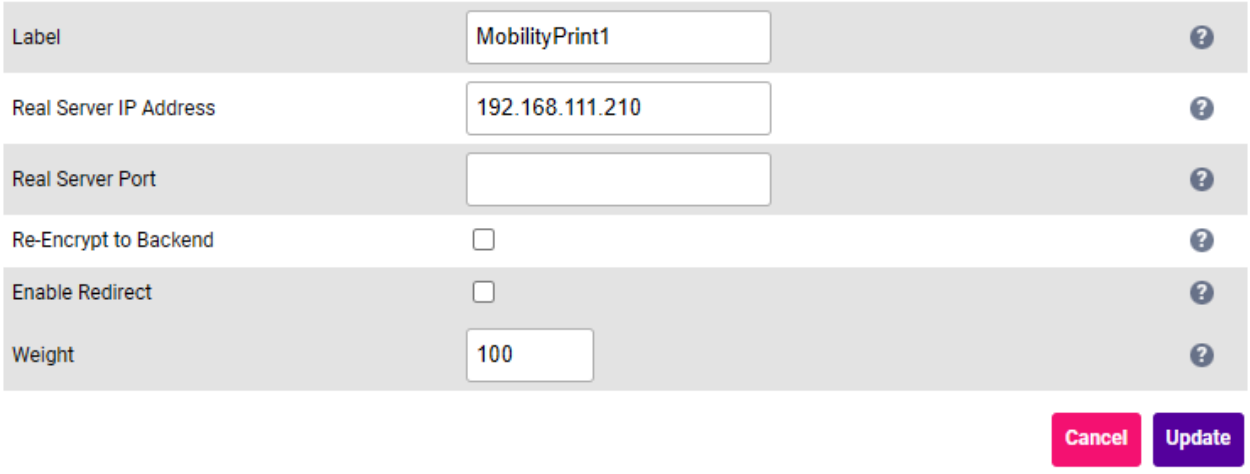

3. Specify an appropriate *Label* for the Real Server, e.g. **MobilityPrint1**.
4. Set the *Real Server IP Address* field to the required address, e.g. **192.168.111.210**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add additional Mobility Print Server(s).

11.2. Papercut Configuration

11.2.1. Configure Papercut Components to Point at the VIPs on the Load Balancer

All PaperCut components must be configured to point at the Virtual Services (VIPs) on the load balancer rather than individual servers.

11.2.2. Add the Load Balancer as a Trusted Proxy

When layer 7 SNAT mode is used, the load balancer should be added as a trusted Proxy to enable the `X-Forwarded-For` headers from the load balancer to be trusted and used by Papercut.

For details, please refer to Set up trusted proxy servers.

> **Note** By default the source IP address for layer 7 VIPs is the interface IP address - in this case, specify the interface address as the trusted proxy. The source address can also be set to any other address that the load balancer owns (typically the VIP address) using the ***Set Source Address*** field for the VIP. In this case, specify this address as the trusted proxy.

11.3. Configure Microsoft Print Servers

If Microsoft print servers are used, follow the steps in Microsoft Print Server Configuration.

12. Microsoft Print Server Configuration

12.1. Enable Print and Document Server Load Balancing

When load balancing Microsoft print and document servers, a number of additional configuration steps must be followed to allow them to be load balanced and accessed via a shared name. The exact steps required depend on the particular version of Windows Server being used as detailed below.

12.1.1. Pre-Requisites

1. Each Server must be joined to the same domain as the client PCs.
2. Each Server must have the **Print and Document Service** role installed.
3. All printers must be installed & shared on each Server using exactly the same share names, settings and permissions.

> **Note** A number of issues have been reported when using Type 4 print drivers, so whenever possible we recommend using Type 3 drivers. Type 4 drivers are usually bundled with the operating system or are downloaded from Windows update, whereas Type 3 drivers are typically downloaded from the printer manufacturer's website.

12.1.2. Enable access via Hostname

To enable the load balanced Print and Document Servers to be accessed via an appropriate hostname, complete the following steps:

> **Note** The configuration steps below assume the hostname for the VIP is **PapercutPrintService** and the domain name is **lbtestdom.com**. Change these to suit your environment.

Windows 2019 & Later

For Windows 2019 & later, local host file entries and a single Registry Key must be added to each Server:

1. Add the following host entries to the local hosts file on each Server:

```
<Real Server IP address> PapercutPrintService
<Real Server IP address> PapercutPrintService.lbtestdom.com
```

For example, if you have 2 Print and Document Servers - 192.168.111.210 and 192.168.111.211, the following entries must be added:

On the 192.168.111.210 server:

```
192.168.111.210 PapercutPrintService
192.168.111.210 PapercutPrintService.lbtestdom.com
```

On the 192.168.111.211 server:

```
192.168.111.211 PapercutPrintService
192.168.111.211 PapercutPrintService.lbtestdom.com
```

2. Add the following Registry Key to each Server:

> **Note** In the example presented here, **PapercutPrintService** is the hostname that will be used to access the load balanced Servers via the virtual service (VIP) created on the load balancer. This can be set to be any appropriate name, although whatever name is used, it must be the **same name** that is used for the DNS entry described in the "Configure DNS Name Resolution" section below.

```
Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters
Value: OptionalNames
Type: REG_MULTI_SZ
Data: PapercutPrintService
```

Windows 2012 & 2016

For Windows 2012 & 2016, the following Registry Keys must be added to each Server:

> **Note** In the example presented here, **PapercutPrintService** is the hostname that will be used to access the load balanced Servers via the virtual service (VIP) created on the load balancer. This can be set to be any appropriate name, although whatever name is used, it must be the **same**

name that is used for the DNS entry described in the "Configure DNS Name Resolution" section below.

```
Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa
Value: DisableLoopbackCheck
Type: REG_DWORD
Data: 1
```

```
Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters
Value: DisableStrictNameChecking
Type: REG_DWORD
Data: 1
```

```
Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters
Value: OptionalNames
Type: REG_MULTI_SZ
Data: PapercutPrintService
```

12.1.3. Configure DNS Name Resolution

1. Create a DNS Host (A) record that points to the VIP address. The hostname used must match the value set for the REG_MULTI_SZ **OptionalNames** registry entry, in this example: **PapercutPrintService → 192.168.111.180**.

12.1.4. Disable NetBIOS over TCP/IP

1. On each Server, disable NetBIOS over TCP/IP on **all** interfaces:

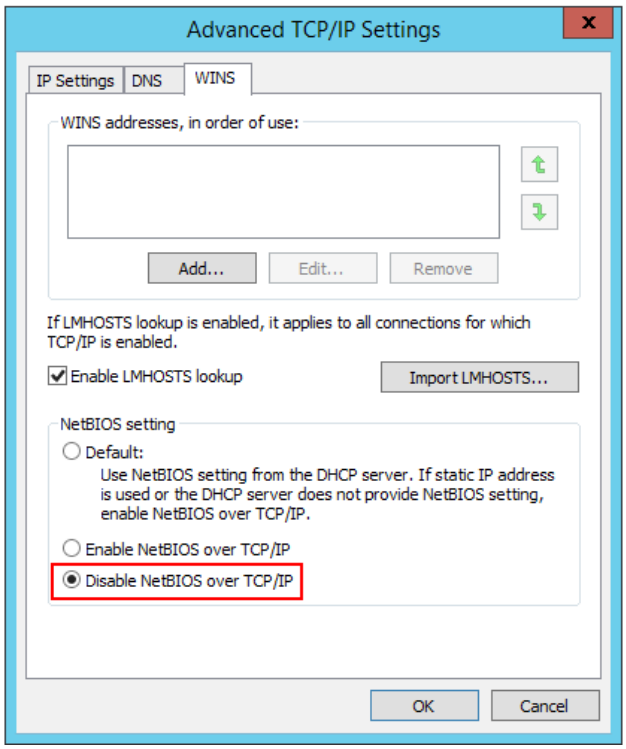

12.1.5. Server Reboot

To apply the changes, reboot each Server.

13. Testing & Verification

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

You should now be able to access your printers by browsing using either the Virtual Service IP address, or the share name. In this example:

```
\\192.168.111.180
```

or

```
\\PapercutPrintService
```

13.1. Using System Overview

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

> **Note** The Application_Servers VIP actively health checks both application servers and will only display the active server in the pool with a green upward arrow. The passive application server will be presented with a red downward arrow until application-server failover occurs on the backend. Servers that a marked with a red arrow will not receive any connections from the load balancer until marked as healthy (green) and online.

The example below shows a Layer 4 DR mode configuration.

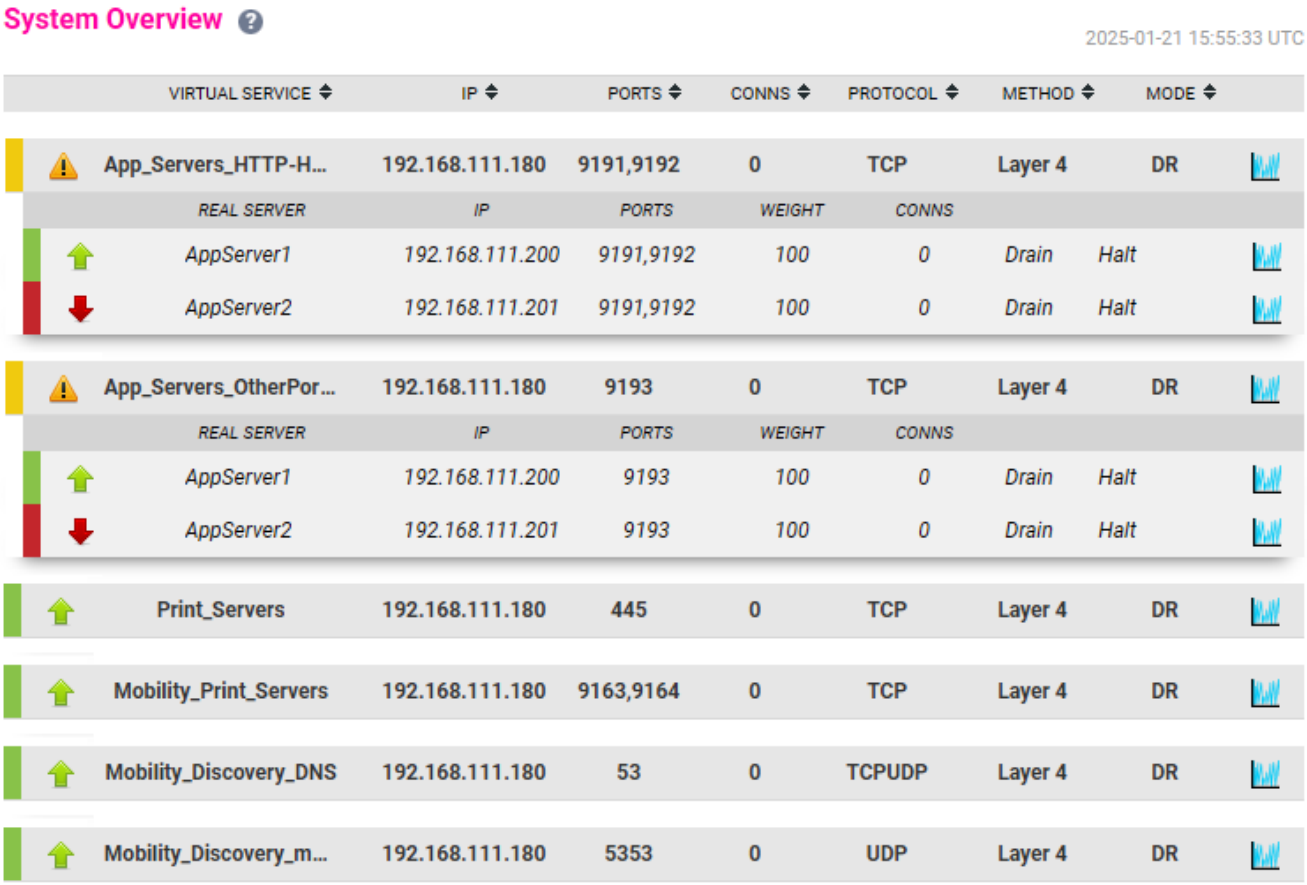

13.2. Client Connection Tests

Ensure that clients are able to print via the load balancer. Make sure that any DNS records are modified to point at the VIPs on the load balancer rather than individual servers.

13.3. Testing PaperCut Application Server Failover

Test	How
Test if the active server is handling traffic	Using a web browser, enter the IP address of the active server (not the Network Load Balancer IP). If the server is in the active state, you will see the PaperCut login page.
Test if the passive server is ready to pick up the load	Using a web browser, enter the IP address of the passive server (not the Network Load Balancer IP). You should see a web page displaying **High Availability activated Server in passive monitoring mode**.
Perform a failover	Trigger a failure on the active Application Server and confirm that the passive server has become active and is working as expected.

> **Note** For more information, please refer to Application Server Failover FAQs.

14. 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.

15. Further Documentation

For additional information, please refer to the Administration Manual.

16. Appendix

16.1. 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 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.

16.1.1. Non-Replicated Settings

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

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	Interface IP addresses, bonding configuration and VLANs
Local Configuration	Routing	Default gateways and static routes
Local Configuration	System Date & time	Time and date related settings
Local Configuration	Physical – Advanced Configuration	Various appliance settings
Local Configuration	Portal Management	Portal management settings
Local Configuration	Security	Security settings
Local Configuration	SNMP Configuration	SNMP settings
Local Configuration	Graphing	Graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Backup & Restore	Local XML backups
Maintenance	Software Updates	Appliance software updates
Maintenance	Firewall Script	Firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

> **Important** Make sure that where any of the above have been configured on the Primary appliance, they're also configured on the Secondary.

16.1.2. 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*.

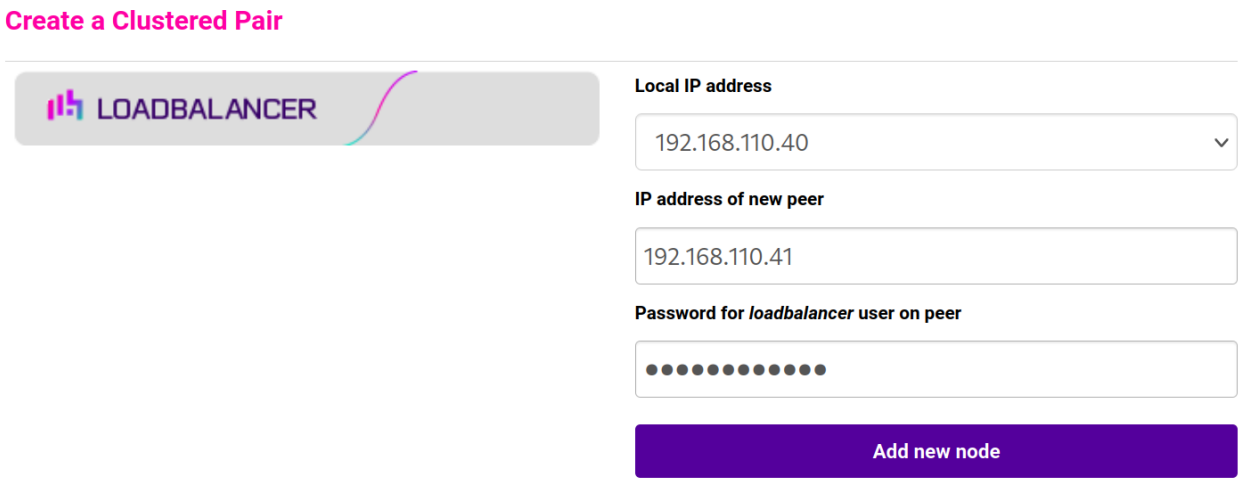

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

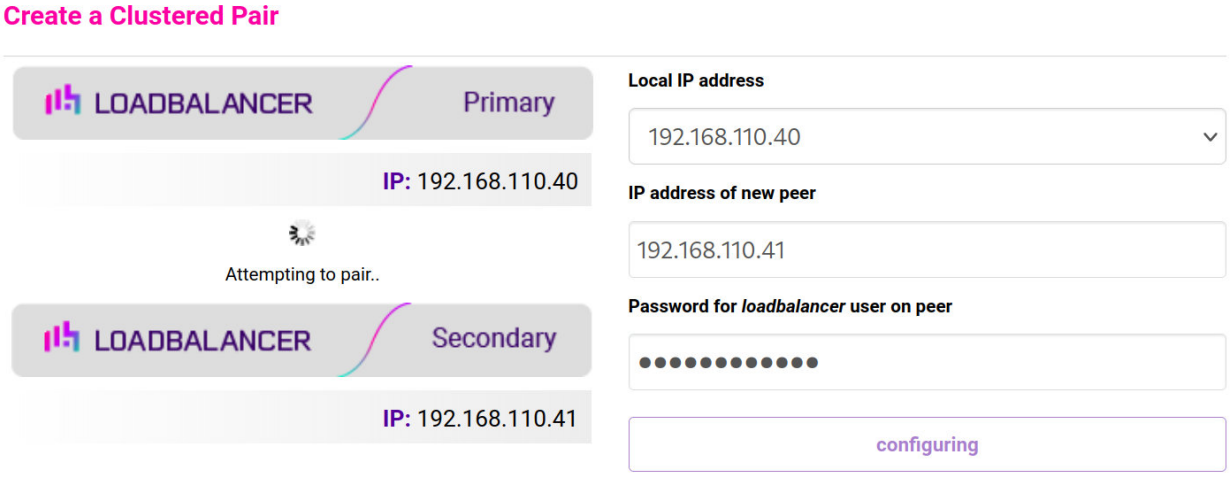

6. Once complete, the following will be displayed on the Primary appliance:

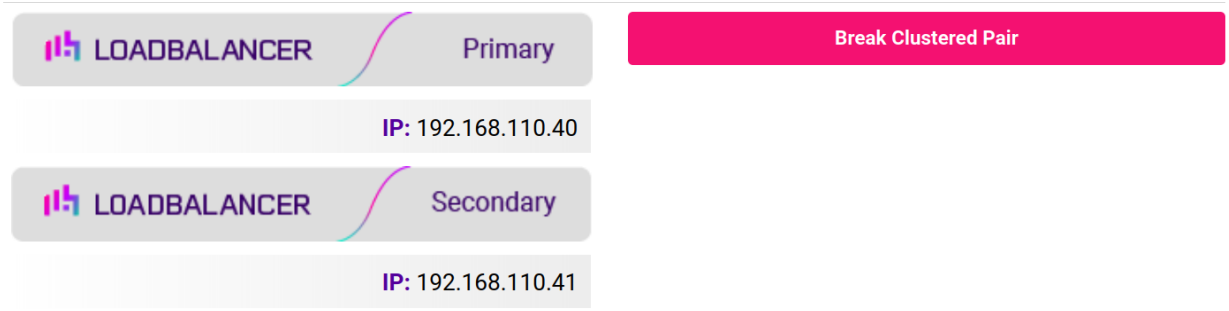

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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.

16.2. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	1 June 2020	Initial version		IBG
1.0.1	15 June 2020	Configuration updates, Papercut hyperlinks added	Required content updates	IBG
1.0.2	19 June 2020	Updated screenshots and hyperlinks Added additional ports for the Papercut Web User Interface service	Required content updates	IBG
1.0.3	26 June 2020	Removed fallback server configuration Replaced system overview image Added note for papercut_wui vip in testing and verification	Required content updates	IBG
1.0.4	29 June 2020	Updated Papercut product information Document title and filename change	Required content updates Differentiating the "Version 19 and earlier" document from the new "Version 20" PaperCut document	IBG, AH
1.0.5	10 August 2020	Updated loopback adaptor settings	Incorrect loopback adaptor configuration	IBG
1.0.6	16 October 2020	Added Layer 7 SNAT configuration Added Fallback Server configuration	Required for multi-site configuration	IBG
1.1.0	1 January 2022	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.1.1	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH

Version	Date	Change	Reason for Change	Changed By
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
2.0.0	31 January 2025	Major document overhaul	Various additions, corrections and improvements to the document's content and structure	RJC

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.