

# Load Balancing PaperCut

Version 1.2.0



# Table of Contents

1. About this Guide	4
2. Loadbalancer.org Appliances Supported	4
3. Software Versions Supported	4
3.1. Loadbalancer.org Appliance	4
3.2. PaperCut NG, MF and Mobility Print	4
4. PaperCut	4
4.1. PaperCut Print Server Components	5
Application Server	5
Secondary Server (Print Provider, Mobility Print)	5
Site Server	5
5. Load Balancing PaperCut	5
5.1. Load Balancing & HA Requirements	5
5.2. Persistence (aka Server Affinity)	6
5.3. Virtual Service (VIP) Requirements	6
5.4. Load Balanced Ports	6
6. Deployment Concept	6
7. Load Balancer Deployment Methods	7
7.1. Layer 4 DR Mode	8
7.2. Layer 4 SNAT Mode	8
7.3. Layer 7 SNAT Mode	9
7.4. Our Recommendation	10
8. Configuring Microsoft Print Servers using PaperCut for Load Balancing	11
8.1. Registry Modifications	11
8.2. Configuring Name Resolution	11
8.3. Layer 4 DR Mode – Solving the ARP Problem	11
9. Loadbalancer.org Appliance – the Basics	11
9.1. Virtual Appliance	11
9.2. Initial Network Configuration	12
9.3. Accessing the Appliance WebUI	12
Main Menu Options	13
9.4. Appliance Software Update	14
Determining the Current Software Version	14
Checking for Updates using Online Update	14
Using Offline Update	14
9.5. Ports Used by the Appliance	15
9.6. HA Clustered Pair Configuration	16
10. Appliance Configuration for PaperCut Print Servers – Using Layer 4 DR Mode	16
10.1. Configuring VIP 1 – PaperCut Application Servers	16
Configuring The Virtual Service (VIP)	16
Define the Real (Active Application Server) Server	17
10.2. Configuring VIP 2 – PaperCut Secondary Server (PaperCut Print Provider)	18
Configuring The Virtual Service (VIP)	18
Define the Real (Print Server) Servers	18
10.3. Configuring VIP 3 – PaperCut Mobility Print	19
Configuring The Virtual Service (VIP)	19
10.4. Finalizing the Layer 4 DR mode Configuration	20
The ARP Problem	20
10.5. Configuring the Print Servers	21

11. Appliance Configuration for PaperCut Print Servers – Using Layer 4 SNAT Mode	21
11.1. Configuring VIP 1 – PaperCut Application Servers	21
Configuring The Virtual Service (VIP)	21
Define the Real (Active Application Server) Server	22
11.2. Configuring VIP 2 – PaperCut Secondary Server (PaperCut Print Provider)	23
Configuring The Virtual Service (VIP)	23
Define The Real (Print Server) Servers	24
11.3. Configuring VIP 3 – PaperCut Mobility Print	25
Configuring The Virtual Service (VIP)	25
11.4. Configuring the Print Servers	26
12. Appliance Configuration for PaperCut Print Servers – Using Layer 7 SNAT Mode	26
12.1. Configuring VIP 1 – PaperCut Application Servers	26
Configuring The Virtual Service (VIP)	26
Define the Real (Active Application Server) Server	27
12.2. Configuring VIP 2 – PaperCut Secondary Server (PaperCut Print Provider)	28
Configuring The Virtual Service (VIP)	28
Define the Real (Print Server) Servers	29
12.3. Configuring VIP 3 – PaperCut Mobility Print	30
Configuring The Virtual Service (VIP)	30
12.4. Configuring the Print Servers	30
13. PaperCut Microsoft Print Server Configuration	31
13.1. Step 1 - Initial Configuration	31
13.2. Step 2 – Registry Modifications	31
13.3. Step 3 – Configure Name Resolution	31
13.4. Step 4 – Server Reboot	32
13.5. Deploying Printers via Group Policy	32
14. Testing & Verification	32
14.1. Using System Overview	32
14.2. Client Connection Tests	33
14.3. Testing PaperCut Application Server failover	33
15. Technical Support	34
16. Additional Documentation	34
17. Appendix	35
17.1. Configuring HA - Adding a Secondary Appliance	35
Non-Replicated Settings	35
Adding a Secondary Appliance - Create an HA Clustered Pair	36
17.2. DR Mode Server Configuration	37
17.3. The ARP Problem - Detecting It and Solving It	37
Detecting the ARP Problem	37
Solving the ARP Problem for Linux	38
Solving the ARP Problem for Mac OS X/BSD	43
Windows Server 2012 & Later	44
17.4. Fallback Server Settings	49
17.5. Local Fallback Server	50
17.6. Using a Separate Dedicated Server	50
17.7. Using a Layer 7 VIP	50
17.8. Configuring A real Server as the Fallback Server	51
17.9. Configuring Primary / Secondary Real Servers	51
17.10. Document Revision History	52

# 1. About this Guide

This guide details the steps required to configure a load balanced PaperCut Application and Secondary print server utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers, Microsoft printer servers and Papercut application 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>.

Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

## 3. Software Versions Supported

### 3.1. Loadbalancer.org Appliance

- V8.4.3 and later

#### Note

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

### 3.2. Papercut NG, MF and Mobility Print

- Papercut NG/MF 20 and later

## 4. PaperCut

PaperCut is a print management solutions provider which delivers this via three applications:

- **PaperCut NG** for easy print management that lets you hit the ground running with full tracking, visibility. It comes with detailed print job tracking and reporting to truly rein in costly, wasteful printing. Boasting eco-friendly policies to help you use less paper, save on toner, and make sustainable habits the status quo.
- **PaperCut MF** lets you cut costs and waste in your workplace by managing print, scan, copy, and fax. It has powerful exclusive features including Secure Print Release, Integrated Scanning, Scan to Cloud and Job Ticketing.
- **PaperCut Mobility Print** keeps users printing when they're outside your network, or on an untrusted guest network. It keeps jobs local to keep printing quick, and only uses the Internet when necessary – and cloud jobs compress and encrypt to save space and keep your data safe.



## 4.1. PaperCut Print Server Components

### Application Server

This is the main application, where you can administer reports, printing costs and print quotas, as well as other print-related actions.

#### Note

When deploying an Application server in failover mode it is recommend to move print queues over to the PaperCut secondary servers. It is also recommended to move the Web print, print deploy service on to share storage. For details on PaperCut's recommended deployment methods for the application servers, see their documentation for Application Server Failover.

### Secondary Server (Print Provider, Mobility Print)

This reports to the application server, updating user and print information that the secondary print server has handled.

#### Note

Mobility Print can be installed on the secondary print servers and can be made highly available when placed behind a load balancer. Mobility Print allows users to print from their mobile devices via network print services and can be load balanced using TCP/UDP 53, 9163, and 9164.

### Site Server

The PaperCut Site Server component ensures continuous availability of printing resources to support key business functions over unreliable network links or during unplanned network disruptions, in remote offices. It is ready to take over the role of a Primary Application Server in the event of a WAN outage. Key roles taken over include authentication, copy and print tracking and Find-Me printing, leaving a remote office with the ability to still be able to print.

## 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, please see the [PaperCut Help Center](#) for further details.

### 5.1. Load Balancing & HA Requirements

This guide details the configuration of a load balanced Microsoft print server deployment using the PaperCut application.

For load balancing print servers, the preferred and default load balancer configuration uses Layer 4 DR Mode (Direct Routing, aka DSR / Direct Server Return). This is a very high performance solution that requires little change to your existing infrastructure. It is necessary to solve "the ARP problem" on the real print servers. This is a straightforward process, and is covered in [DR Mode Server Configuration](#).

It is also possible to load balance a PaperCut Secondary print server using Layer 4 SNAT Mode. This mode might be preferable if making changes to the real print servers is not possible, although some Windows Registry keys need to be added. Please note that load balanced connections using layer 4 SNAT mode are not source IP transparent, which is not usually an issue when load balancing print servers but should still be considered.



The PaperCut Application servers can only be deployed in an active-passive configuration when placed behind the load balancer. This is known as a 'failover' configuration and purely provides high availability, not load balancing.

In order to configure the Application servers in an active-passive configuration, the environment must be deployed with:

#### Note

- A persistent network drive accessible by all servers
- A highly available external database (RDBMS), such as:
  - Microsoft SQL Server
  - PostgreSQL
  - MySQL
  - Oracle

Configuring the external database is beyond the scope of this guide.

## 5.2. Persistence (aka Server Affinity)

Neither Microsoft print servers or the PaperCut application require session affinity at the load balancing layer.

## 5.3. Virtual Service (VIP) Requirements

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

- PaperCut Application Servers
- PaperCut Mobility Print
- PaperCut Print Provider

## 5.4. Load Balanced Ports

The following table shows the ports that are load balanced:

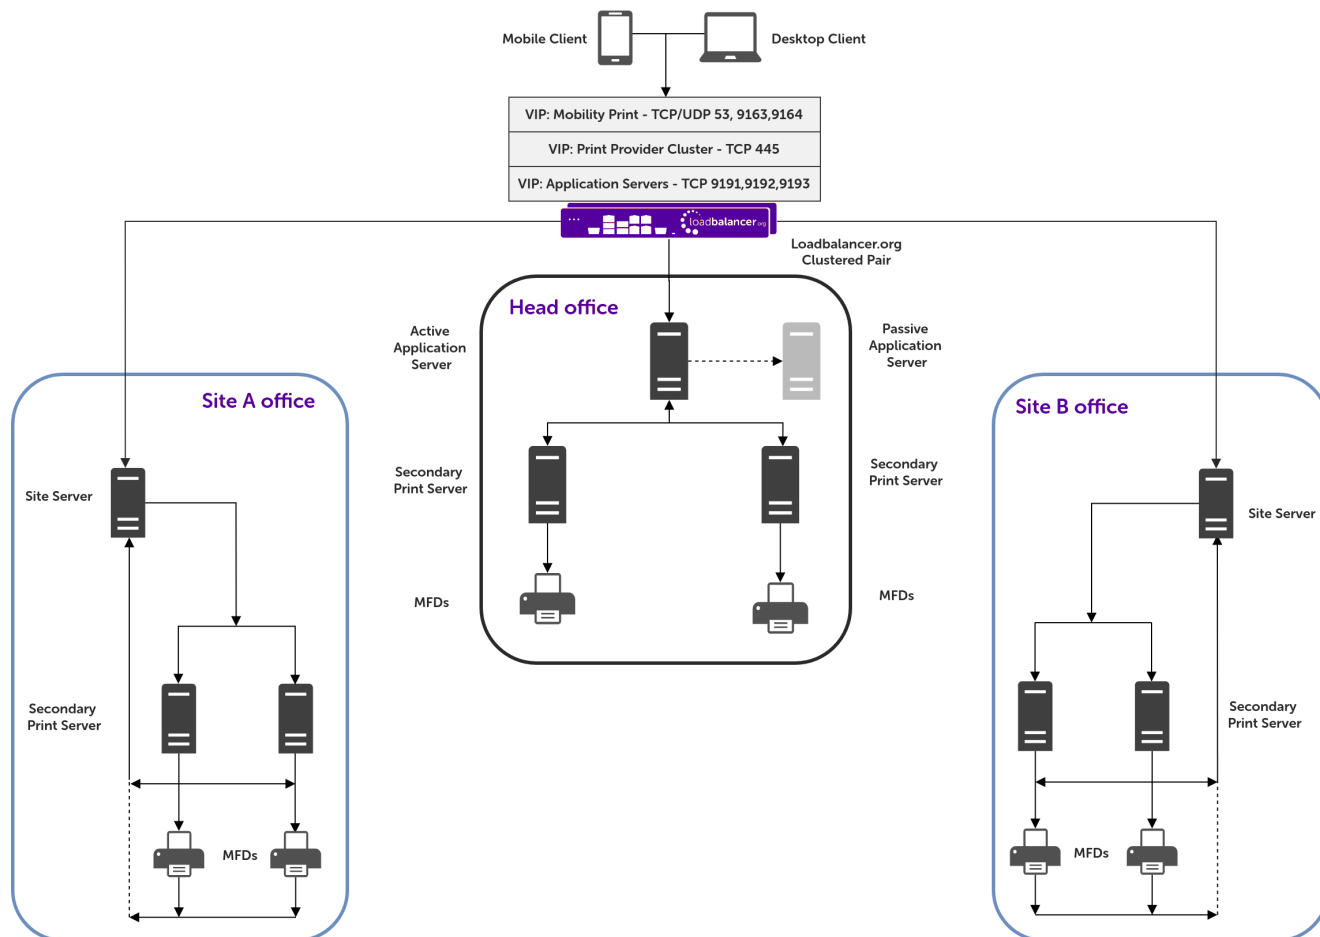
Ports	Use	Transport Layer Protocols
445	Papercut Print Provider	TCP
53, 9163, 9164	Papercut Mobility Print	TCP/UDP
9191,9192,9193	Papercut Web User Interface	TCP
9173, 9175	Papercut Print Deploy	TCP

#### Note

A list of additional ports required to configure high availability to work with your required Multifunctional Device vendor can be found on the [PaperCut Help Center](#).

# 6. Deployment Concept





VIPs = **V**irtual **I**P Addresses

### 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:

1. Layer 4 DR mode
2. Layer 4 NAT mode
3. Layer 4 SNAT mode
4. Layer 7 SNAT mode

For Microsoft Print Servers using PaperCut, layer 4 DR mode and layer 4 & 7 SNAT modes are recommended. These modes are described below and are used for the configurations presented in this guide. For configuring using DR mode refer to [Appliance Configuration for PaperCut Print Servers – Using Layer 4 DR Mode](#) and for configuring using layer 4 SNAT mode refer to [Appliance Configuration for PaperCut Print Servers – Using Layer 4 SNAT Mode](#).

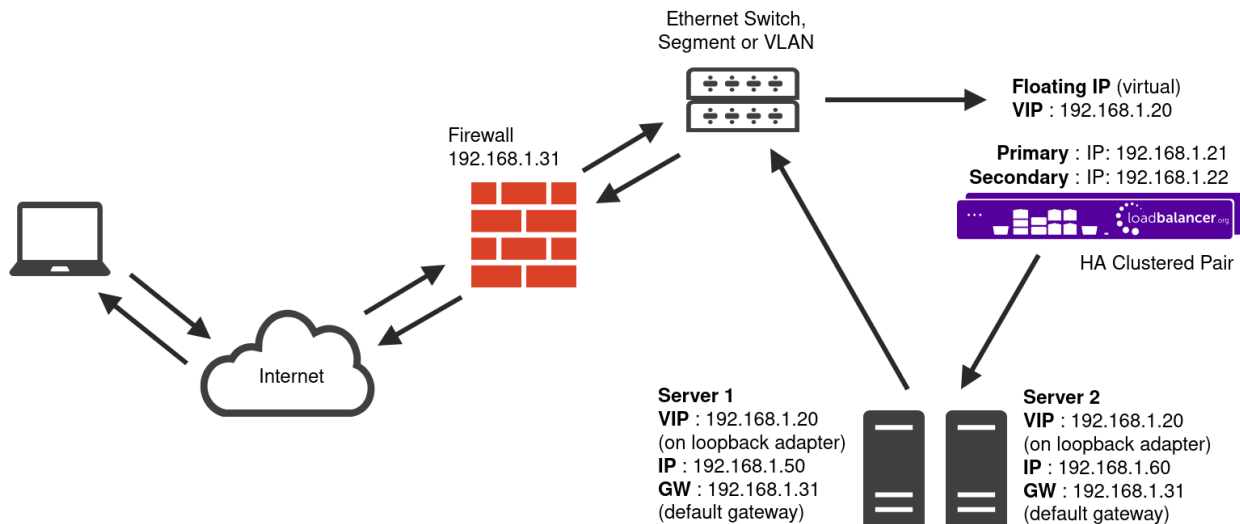


## 7.1. Layer 4 DR Mode

One-arm direct routing (DR) mode is a very high performance solution that requires little change to your existing infrastructure.

### Note

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

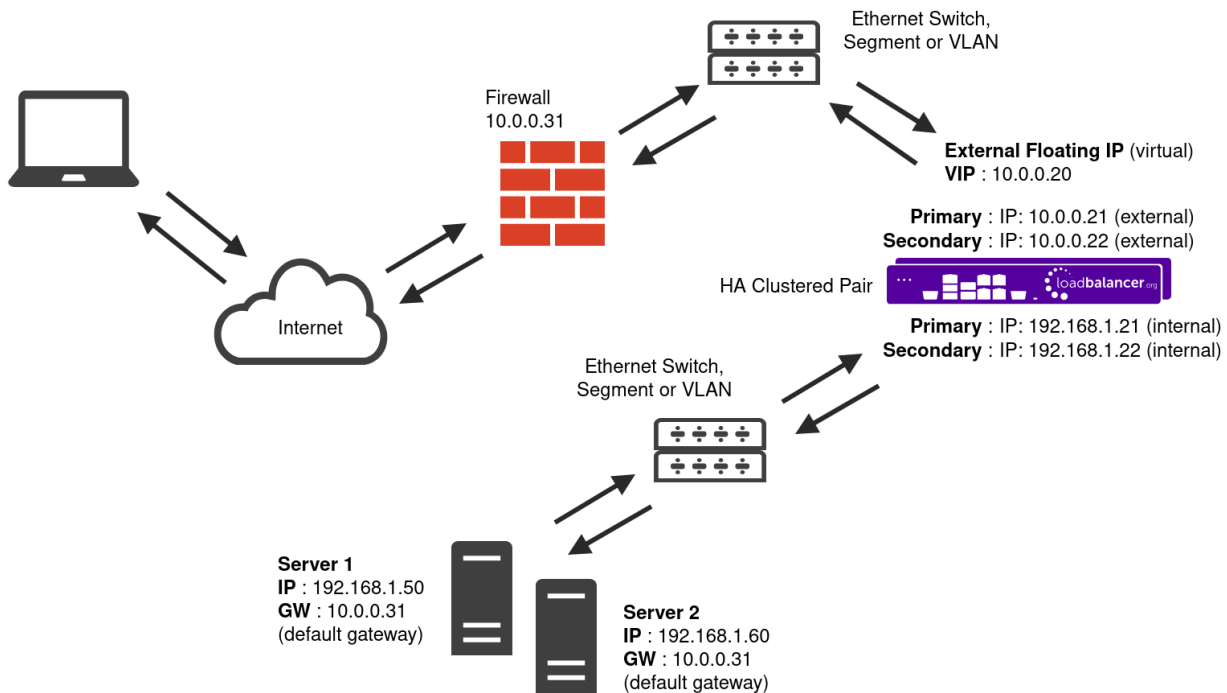


- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that you need to ensure that the Real Server (and the load balanced application) respond to both the Real Server's own IP address and the VIP.
- The Real Servers should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Servers 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 for HTTP, 50 times quicker for Terminal Services and much, much faster for streaming media or FTP.
- The load balancer must have an Interface in the same subnet as the Real Servers to ensure layer 2 connectivity required for DR mode to work.
- 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 4 SNAT Mode

Layer 4 SNAT mode is a high performance solution, although not as fast as Layer 4 NAT mode or Layer 4 DR mode.






- The load balancer translates all requests from the external Virtual Service to the internal Real Servers in the same way as NAT mode - please refer to [Layer 4 NAT Mode](#) for more information.
- Layer 4 SNAT mode is not transparent, an iptables SNAT rule translates the source IP address to be the load balancer rather than the original client IP address.
- Layer 4 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, **eth0** is normally used for the internal network and **eth1** is used for the external network although this is not mandatory.
- If the Real Servers require Internet access, Autonat should be enabled using the WebUI option: **Cluster Configuration > Layer 4 - Advanced Configuration**, the external interface should be selected.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 4 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 4 SNAT mode VIPs and layer 7 SNAT mode VIPs because the required firewall rules conflict.

### 7.3. 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.



- Because layer 7 SNAT mode is a full proxy, any server in the cluster can be on any accessible subnet 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, **eth0** is normally used for the internal network and **eth1** is used for the external network although this is not mandatory.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

## 7.4. Our Recommendation

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 Real 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.

If DR mode cannot be used, for example if it is not possible to make changes to the real servers, or if the real servers are located in remote routed networks, then layer 4 SNAT mode is recommended.

## 8. Configuring Microsoft Print Servers using PaperCut for Load Balancing

To configure Microsoft print servers for load balancing the following settings need to be applied:

### 8.1. Registry Modifications

For the print servers that are going to be load balanced, to enable them to be accessed via a shared name.

### 8.2. Configuring Name Resolution

For printer load balancing to work, DNS name resolution should be configured. A host name and corresponding "Host (A)" record for the virtual service should be created, and should match the virtual IP (VIP) address defined on the load balancer.

#### Note

For details of the required changes for registry modifications and configuring name resolution please refer to [PaperCut Microsoft Print Server Configuration](#).

### 8.3. Layer 4 DR Mode – Solving the ARP Problem

If using layer 4 DR mode, the 'ARP problem' must be solved on each real server for DR mode to work. For detailed steps on solving the ARP problem, please refer to [The ARP Problem - Detecting It and Solving It](#) for more information. For a detailed explanation of DR mode and the nature of the ARP problem, please refer to the section that covers [Layer 4 DR Mode](#).

## 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 Server and other network 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](#).

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

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


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

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


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

**Username:** loadbalancer

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

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

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



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



### Note

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

## Main Menu Options

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

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

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

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

**View Configuration** - Display the saved appliance configuration settings

**Reports** - View various appliance reports & graphs

**Logs** - View various appliance logs

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

## 9.4. Appliance Software Update

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

### Determining the Current Software Version

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

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

English ▼

### Checking for Updates using Online Update

#### Note

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

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

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

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

#### Note

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

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

**Information:** Update completed successfully.

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

### Using Offline Update

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





## Note

Please contact [support@loadbalancer.org](mailto:support@loadbalancer.org) to check if an update is available and obtain the latest offline update files.

To perform an offline update:

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

## Software Update

### Offline Update

The following steps will lead you through offline update.

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

Archive:  No file chosen

Checksum:  No file chosen

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

## 9.5. Ports Used by the Appliance

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

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



## 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. Appliance Configuration for PaperCut Print Servers – Using Layer 4 DR Mode






When deploying PaperCut, three virtual services must be configured: a virtual service for the PaperCut Application Server, the Print Provider, and a virtual service for the PaperCut Mobility Print.

## 10.1. Configuring VIP 1 – PaperCut Application Servers

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the required *Label* (name) for the VIP, e.g. **Papercut\_WUI**.
3. Set the *Virtual Service IP address* field to the required IP address, e.g. **172.24.11.38**.
4. Set the *Virtual Service Ports* field to **9191,9192,9193**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the virtual service.

**Layer 4 - Add a new Virtual Service**

Virtual Service		
Label	<input type="text" value="Papercut_WUI"/>	
IP Address	<input type="text" value="172.24.11.38"/>	
Ports	<input type="text" value="9191,9192,9193"/>	
Protocol		
Protocol	<input type="text" value="TCP"/>	
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	

8. Now click **Modify** next to the newly created Virtual Service.
9. Disable *Persistence* by unchecking the **Enable** check box.



- Under **Health Checks** set the **Check Type** to **Negotiate**.
- Leave the **Check Port** field empty.
- In the **Request to send** field put the application server health monitoring authorization key which 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 in the Application server under **Web user interface > Options > Advanced**

#### HTTP header

```
Authorization:JjtxY8ztIIZhAOKtGHs2swxw7Q3eyVXH
```

- Click **Update**.

#### Note

In some cases other ports may need to be forwarded such as port 9192, 9193 and additional ports depending on a customers multifunctional devices. For a list of PaperCut ports please refer to the [PaperCut Help Centre](#).

## Define the Real (Active Application Server) Server

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

#### Layer 4 Add a new Real Server - Papercut\_WUI

Label	<input type="text" value="App_Svr1"/>	?
Real Server IP Address	<input type="text" value="172.24.11.36"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

Cancel

Update

- Enter an appropriate label for the Real Server, e.g. **App\_Svr1**.
- Change the **Real Server IP Address** field to the required address, e.g. **172.24.11.36**.
- Click **Update**.

6. Repeat the above steps for the remaining application server.

## 10.2. Configuring VIP 2 – PaperCut Secondary Server (PaperCut Print Provider)

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the required *Label* (name) for the VIP, e.g. **Print\_Provider**.
3. Set the *Virtual Service IP address* field to the required IP address, e.g. **172.24.11.38**.
4. Set the *Virtual Service 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.
8. Now click **Modify** next to the newly created Virtual Service.
9. Disable *Persistence* by unchecking the **Enable** check box.
10. Click **Update**.

**Layer 4 - Add a new Virtual Service**

Virtual Service		
Label	<input type="text" value="Print_Provider"/>	<a href="#">?</a>
IP Address	<input type="text" value="172.24.11.38"/>	<a href="#">?</a>
Ports	<input type="text" value="445"/>	<a href="#">?</a>
Protocol		
Protocol	<input type="text" value="TCP"/>	<a href="#">?</a>
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	<a href="#">?</a>

### Define the Real (Print Server) Servers

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:

#### Layer 4 Add a new Real Server - Print\_Provider

Label	<input type="text" value="PS1"/>	?
Real Server IP Address	<input type="text" value="172.24.11.39"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

CancelUpdate

3. Enter an appropriate label for the Real Server, e.g. **PS1**.
4. Change the *Real Server IP Address* field to the required address, e.g. **172.24.11.39**.
5. Click **Update**.
6. Repeat the above steps to add your other Print Server(s).

#### Note

In the next section, "Configuring VIP 3 – PaperCut Mobility Print" we will make use of the **Duplicate Service** button to retain the configuration including the added real servers. We will then need to amend the configuration with a new label and IP address accordingly, while other configuration items, such as added real servers, will be retained.

**Duplicate Service**



## 10.3. Configuring VIP 3 – PaperCut Mobility Print

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Modify** on the **PrintProviderVIP** virtual service.
2. Click the **Duplicate Service** located in the top right of the menu.
3. Define the required *Label* (name) for the VIP, e.g. **MobilityPrint**.
4. Set the *Virtual Service Ports* field to **53,9163,9164**.
5. Leave the *Protocol* set to **TCP/UDP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Under the **Health Checks** set the *Check Port* to **9163**.
8. Click **Update**.

**WARNING:** Label Print\_Provider already exists. A temporary label of Print\_Provider-1 has been assigned.

Duplicate Service ?

#### Layer 4 - Modify Virtual Service

Virtual Service		
Label	<input type="text" value="MobilityPrint"/>	?
IP Address	<input type="text" value="172.24.11.38"/>	?
Ports	<input type="text" value="53,9163,9164"/>	?
IP Protocol		
Protocol	<input type="text" value="TCP/UDP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	?
Connection Distribution Method		
Balance Mode	<input type="text" value="Weighted Least Connection"/>	?
Persistence		
Enable	<input type="checkbox"/>	?
Health Checks		
Check Type	<input type="text" value="Connect to port"/>	?
Check Port	<input type="text" value="9163"/>	?

#### Note

Please be aware that Mobility Print will need to be installed on the same Secondary print servers within the cluster. However it is recommended that you segregate these services at a VIP level as this allows for the more granular control and health checking of those services. For example if you create a singular VIP with multiple services being load balanced by that VIP, should one of those services were to fail the VIP and corresponding real server will also be marked as down. Having a singular multi service VIP is fine during testing but not recommended for production.

## 10.4. Finalizing the Layer 4 DR mode Configuration

When using a layer 4 DR mode configuration, all real servers need to be configured to solve the "ARP problem."

### The ARP Problem

DR mode works by changing the MAC address of the inbound packets to match the Real Server selected by the load balancing algorithm. To enable DR mode to operate:

- Each Real Server must be configured to accept packets destined for both the VIP address **and** the Real Server's IP address (RIP). This is because in DR mode the destination address of load balanced packets is the VIP address, whilst for other traffic such as health checks, administration traffic etc. it's the Real Server's



own IP address (the RIP). The service/process (e.g. IIS) must also respond to both addresses.

- Each Real Server must be configured so that it does not respond to ARP requests for the VIP address – only the load balancer should do this.

Configuring the Real Servers in this way is referred to as "Solving the ARP problem". The steps required depend on the OS used as detailed in [The ARP Problem - Detecting It and Solving It](#).

## 10.5. Configuring the Print Servers

Now follow the steps in [PaperCut Microsoft Print Server Configuration](#).

# 11. Appliance Configuration for PaperCut Print Servers – Using Layer 4 SNAT Mode

When deploying PaperCut, three virtual services must be configured: a virtual service for the PaperCut Application Server, the Print Provider, and a virtual service for the PaperCut Mobility Print.






## 11.1. Configuring VIP 1 – PaperCut Application Servers

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the required *Label* (name) for the VIP, e.g. **Papercut\_WUI**.
3. Set the *Virtual Service IP address* field to the required IP address, e.g. **172.24.11.38**.
4. Set the *Virtual Service Ports* field to **9191,9192,9193**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **SNAT**.
7. Click **Update** to create the virtual service.



#### Layer 4 - Add a new Virtual Service

Virtual Service		
Label	<input type="text" value="Papercut_WUI"/>	
IP Address	<input type="text" value="172.24.11.38"/>	
Ports	<input type="text" value="9191,9192,9193"/>	
Protocol		
Protocol	<input type="text" value="TCP"/>	
Forwarding		
Forwarding Method	<input type="text" value="SNAT"/>	
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

- Now click **Modify** next to the newly created Virtual Service.
- Disable **Persistence** by unchecking the **Enable** check box.
- Under **Health Checks**, set the **\_Check Type** to **Negotiate**.
- Leave the **Check Port** field empty.
- In the **Request to send** field, input the application server health monitoring authorization key which 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 in the Application server under **Web user interface > Options > Advanced**.

#### HTTP header

```
Authorization:JjtxY8ztIIZhAOKtGHs2swxw7Q3eyVXH
```

- Click **Update**.

## Define the Real (Active Application Server) Server

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

#### Layer 4 Add a new Real Server - PaperCut\_WUI

Label	<input type="text" value="App_Svr1"/>	?
Real Server IP Address	<input type="text" value="172.24.11.36"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

CancelUpdate

3. Enter an appropriate *Label* for the Real Server, e.g. **App\_Svr1**.
4. Change the *Real Server IP Address* field to the required address, e.g. **172.24.11.36**.
5. Click **Update**.
6. Repeat the above steps for the remaining application server.

## 11.2. Configuring VIP 2 – PaperCut Secondary Server (PaperCut Print Provider)

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the required *Label* (name) for the VIP, e.g. **Print\_Provider**.
3. Set the *Virtual Service IP address* field to the required IP address, e.g. **172.24.11.38**.
4. Set the *Virtual Service Ports* field to **445**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **SNAT**.
7. Click **Update** to create the virtual service.
8. Now click **Modify** next to the newly created Virtual Service.
9. Disable *Persistence* by unchecking the **Enable** check box.
10. Click **Update**.

#### Layer 4 - Add a new Virtual Service

Virtual Service		
Label	<input type="text" value="Print_Provider"/>	?
IP Address	<input type="text" value="172.24.11.38"/>	?
Ports	<input type="text" value="445"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="SNAT"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

### Define The Real (Print Server) Servers

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:

#### Layer 4 Add a new Real Server - Print\_Provider

Label	<input type="text" value="PS1"/>	?
Real Server IP Address	<input type="text" value="172.24.11.39"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate *Label* for the Real Server, e.g. **PS1**.
4. Change the *Real Server IP Address* field to the required address, e.g. **172.24.11.39**.
5. Click **Update**.
6. Repeat the above steps for the remaining application server.

#### Note

In the next section, "Configuring VIP 3 – PaperCut Mobility Print", we will make use of the **Duplicate Service** button to retain the configuration including the added real servers. We will then need to amend the configuration with a new label and IP address accordingly, while other configuration items, such as added real servers, will be retained.





Duplicate Service ?

## 11.3. Configuring VIP 3 – PaperCut Mobility Print

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Modify** on the **PrintProviderVIP** virtual service.
2. Click the **Duplicate Service** located in the top right of the menu.
3. Define the required *Label* (name) for the VIP, e.g. **MobilityPrint**.
4. Set the *Virtual Service Ports* field to **53,9163,9164**.
5. Leave the *Protocol* set to **TCP/UDP**.
6. Leave the *Forwarding Method* set to **SNAT**.
7. Under **Health Checks** set the *Check Port* to **9163**.
8. Click **Update**.

**WARNING:** Label Print\_Provider already exists. A temporary label of Print\_Provider-1 has been assigned.

Duplicate Service ?

#### Layer 4 - Modify Virtual Service

Virtual Service		
Label	<input type="text" value="MobiltyPrint"/>	?
IP Address	<input type="text" value="172.24.11.38"/>	?
Ports	<input type="text" value="53,9163,9164"/>	?
IP Protocol		
Protocol	<input type="text" value="TCP/UDP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="SNAT"/>	?
Connection Distribution Method		
Balance Mode	<input type="text" value="Weighted Least Connection"/>	?
Persistence		
Enable	<input type="checkbox"/>	?
Health Checks		
Check Type	<input type="text" value="Connect to port"/>	?
Check Port	<input type="text" value="9163"/>	?

## 11.4. Configuring the Print Servers

Now follow the steps in [PaperCut Microsoft Print Server Configuration](#).

## 12. Appliance Configuration for PaperCut Print Servers – Using Layer 7 SNAT Mode

When deploying PaperCut, three virtual services must be configured: a virtual service for the PaperCut Application Server, the Print Provider, and a virtual service for the PaperCut Mobility Print.

### 12.1. Configuring VIP 1 – PaperCut Application Servers

#### Configuring The Virtual Service (VIP)

1. Using the web user interface, 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. **Papercut\_WUI**.



3. Set the *Virtual Service IP address* field to the required IP address, e.g. **172.24.11.38**.
4. Set the *Virtual Service Ports* field to **9191,9192,9193**.
5. Set the *Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

### Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="Papercut_WUI"/>	?
IP Address	<input type="text" value="172.24.11.38"/>	?
Ports	<input type="text" value="9191,9192,9193"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

Cancel Update

7. Click **Modify** next to the newly created Virtual Service.
8. Under *Persistence* select **None**.
9. Under *Health Checks* set the *Check Type* to **Negotiate**.
10. Leave the *Check Port* field empty.
11. In the *Request to send* field put the application server health monitoring authorization key which 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 in the Application server under *Web user interface > Options > Advanced*.

#### HTTP header

```
Authorization:JjtxY8ztIIZhA0KtGHs2swxw7Q3eyVXH
```

12. Click **Update**.

## Define the Real (Active Application Server) Server

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:



### Layer 7 Add a new Real Server - PaperCut\_WUI

Label	<input type="text" value="App_Svr1"/>	<a href="#">?</a>
Real Server IP Address	<input type="text" value="172.24.11.36"/>	<a href="#">?</a>
Real Server Port	<input type="text"/>	<a href="#">?</a>
Re-Encrypt to Backend	<input type="checkbox"/>	<a href="#">?</a>
Weight	<input type="text" value="100"/>	<a href="#">?</a>

CancelUpdate

3. Enter an appropriate *Label* for the Real Server, e.g. **App\_Svr1**.
4. Change the *Real Server IP Address* field to the required address, e.g. **172.24.11.36**.
5. Click **Update**.
6. Repeat the above steps for the remaining application server.

## 12.2. Configuring VIP 2 – PaperCut Secondary Server (PaperCut Print Provider)

### Configuring The Virtual Service (VIP)

1. Using the web user interface, 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\_Provider**.
3. Set the *Virtual Service IP address* field to the required IP address, e.g. **172.24.11.38**.
4. Set the *Virtual Service Ports* field to **445**.
5. Set the *Protocol* to **TCP Mode**.
6. Click **Update** to create the virtual service.

## Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="Print_Provider"/>	?
IP Address	<input type="text" value="172.24.11.38"/>	?
Ports	<input type="text" value="445"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

- Now click **Modify** next to the newly created Virtual Service.
- Under *Persistence* select **None**.
- Click **Update**.

## Define the Real (Print Server) Servers

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

### Layer 7 Add a new Real Server - Print\_Provider

Label	<input type="text" value="PS1"/>	?
Real Server IP Address	<input type="text" value="172.24.11.39"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

#### Note

In the next section, "Configuring VIP 3 – PaperCut Mobility Print", we will make use of the **Duplicate Service** button to retain the configuration including the added real servers. We will then need to amend the configuration with a new label and IP address accordingly, while other configuration items, such as added real servers, will be retained.

 Duplicate Service

## 12.3. Configuring VIP 3 – PaperCut Mobility Print

### Configuring The Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Modify** on the **PrintProviderVIP** virtual service.
2. Click **Duplicate Service**, located in the top right of the menu.
3. Define the required *Label* (name) for the VIP, e.g. **MobilityPrint**.
4. Set the *Virtual Service Ports* field to **53,9163,9164**.
5. Set the *Protocol* to **TCP Mode**.
6. Under **Health Checks** click *Advanced* and set *Check Port* to **9163**.
7. Click **Update**.
8. Click **Reload Haproxy** to commit the configuration.

**WARNING:** Label Print\_Provider already exists. A temporary label of Print\_Provider-1 has been assigned.

Duplicate Service ?

### Layer 7 - Modify Virtual Service

Virtual Service		
Manual Configuration	<input type="checkbox"/>	?
Label	<input type="text" value="Mobility_Print"/>	?
IP Address	<input type="text" value="172.24.11.38"/>	?
Ports	<input type="text" value="53,9163,9164"/>	?
Protocol [Advanced]		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
Connection Distribution Method		
Balance Mode	<input type="text" value="Weighted Least Connections"/>	?
Persistence [Advanced]		
Persistence Mode	<input type="text" value="None"/>	?
Health Checks [Advanced]		
Health Checks	<input type="text" value="Connect to port"/>	?
Check Port	<input type="text" value="9163"/>	?

## 12.4. Configuring the Print Servers



Now follow the steps in [PaperCut Microsoft Print Server Configuration](#).

## 13. PaperCut Microsoft Print Server Configuration

### 13.1. Step 1 - Initial Configuration

Complete the following steps on each print server:

1. Join the server to the same domain as the client PCs.
2. Install the **Print and Document Service** role / **Print Server** service.
3. Install & share the printers (use the same share names and permissions across all servers).

### 13.2. Step 2 – Registry Modifications

To enable the print servers to be accessed via a shared name (**PapercutPrintService** in this guide), add the following registry entries to each print server:

Ref.	Registry Key Requirements
1	Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa  Value: <b>DisableLoopbackCheck</b>  Type: REG_DWORD  Data: 1
2	Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters  Value: <b>DisableStrictNameChecking</b>  Type: REG_DWORD  Data: 1
3	Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters  Value: <b>OptionalNames</b>  Type: REG_MULTI_SZ  Data: PapercutPrintService

### 13.3. Step 3 – Configure Name Resolution



To configure DNS name resolution complete the following steps:

1. Disable NetBIOS over TCP/IP on *all* interfaces of each print server.
2. Create a DNS record for the share name, in this example: **PapercutPrintService** → **172.24.11.38**.

## 13.4. Step 4 – Server Reboot

To apply all the changes, reboot each print server.

## 13.5. Deploying Printers via Group Policy

- Ensure that the load balanced print server name (e.g. **PapercutPrintService**) is resolvable by DNS as explained in [Step 3 – Configure Name Resolution](#).
- On your print server, open: **Administrative Tools > Printer Management**.
  - Right-click Print Servers and enter the name for your load balanced print server (e.g. **PapercutPrintService**) and click OK.
  - Expand the **Printers** section.
  - Right click the printer you want to deploy, and click **Deploy with Group Policy**.
  - Select the relevant GPO and configure the remaining settings according to your requirements.

### Note

PaperCut NG and MF have a fantastic feature called Print Deploy which makes deployment of print queues out to end users workstations super simple. For further details please see the [Papercut Help Center](#).

## 14. 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:

```
\\172.24.11.38
```

or

```
\\PapercutPrintService
```

### 14.1. 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 secondary 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 Papercut\_WUI 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 are marked with a red arrow will not receive any connections from the load balancer until marked as healthy (green) and online.

## System Overview ?

2020-06-26 12:38:36 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	Papercut_WUI	172.24.11.38	9191,9192..	1	TCP	Layer 4	DR	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	App_Svr1	172.24.11.36	9191,9192,..	100	1	Drain	Halt	
	App_Svr2	172.24.11.37	9191,9192,..	100	0	Drain	Halt	
	Print_Provider	172.24.11.38	445	0	TCP	Layer 4	DR	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	PS1	172.24.11.39	445	100	0	Drain	Halt	
	PS2	172.24.11.40	445	100	0	Drain	Halt	
	MobilityPrint	172.24.11.38	53,9163,9..	0	TCPUDP	Layer 4	DR	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	PS1	172.24.11.39	53,9163,9164	100	0	Drain	Halt	
	PS2	172.24.11.40	53,9163,9164	100	0	Drain	Halt	

## 14.2. Client Connection Tests

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

## Note

For more details on testing & diagnosing load balanced services please refer to [Testing Load Balanced Services](#).

## 14.3. Testing PaperCut Application Server failover

Test	How
Test if the active server is handling traffic.	Attempt to load the PaperCut NG/MF admin web interface using the IP or hostname of the server that you want to test—not the IP/hostname of the NLB. If the server is in the active state, you will see the PaperCut login page.

Test	How
Test if the passive server is ready to pick up the load.	<p>Attempt to load the Admin web interface using the IP or hostname of the server that you want to test—not the IP/hostname of the NLB. If the server is in the passive state, you will see a page that looks like this:</p> <div> <p><b>High availability activated</b></p> <p><b>Server in passive monitoring mode</b></p> </div>
Test if a device is connected via the NLB.	Change the IP/hostname that the device is configured with to be the IP of the Network Load Balancer and restart the device. If the device connects, the NLB is correctly handling the traffic.
Test if secondary components (user client, secondary server, etc.) are connected to the NLB.	Change the configured IP and restart (same process as above).
Perform a failover	Trigger a failure on the active Application Server and confirm that traffic is routed to and operation continues automatically on another server in the pool. We recommend performing this multiple times for each server in the pool.

## 15. 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](mailto:support@loadbalancer.org).

## 16. Additional Documentation

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



# 17. Appendix

## 17.1. Configuring HA - Adding a Secondary Appliance

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

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

### Note

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

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

### Non-Replicated Settings

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

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



### ⚠ Important

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


## Adding a Secondary Appliance - Create an HA Clustered Pair

### 📌 Note

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

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

### Create a Clustered Pair

 **LOADBALANCER**

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41


Password for *loadbalancer* user on peer

••••••••••

Add new node

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


### Create a Clustered Pair

 **LOADBALANCER**

Primary

IP: 192.168.110.40

Attempting to pair..

 **LOADBALANCER**

Secondary

IP: 192.168.110.41

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41


Password for *loadbalancer* user on peer

••••••••••


configuring

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

## High Availability Configuration - primary

 **LOADBALANCER** Primary

IP: 192.168.110.40

 **LOADBALANCER** Secondary

IP: 192.168.110.41

**Break Clustered Pair**

**Make Active**

- 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](#).

## 17.2. DR Mode Server Configuration

When using Layer 4 DR mode the ARP problem must be solved. This involves configuring each PaperCut Secondary Print Server to accept traffic destined for the VIP in addition to its own IP address and ensuring that each server does not respond to ARP requests for the VIP address – only the load balancer should do this.

## 17.3. The ARP Problem - Detecting It and Solving It

### Detecting the ARP Problem

Attempt to connect to the VIP and then use **Reports > Layer 4 Current Connections** to check whether the connection state is SYN\_RECV as shown below.

### LAYER 4 CURRENT CONNECTIONS

**Check Status**

#### IPVS connection entries

pro	expire	state	source	virtual	destination
TCP	00:26	SYN_RECV	192.168.64.7:20415	192.168.111.232:80	192.168.110.240:80
TCP	00:26	SYN_RECV	192.168.64.7:20414	192.168.111.232:80	192.168.110.240:80
TCP	04:18	NONE	192.168.64.7:0	192.168.111.232:80	192.168.110.240:80

If it is, this is normally a good indication that the ARP problem has not been correctly solved.



## Solving the ARP Problem for Linux

There are two different approaches on how to configure a Linux server for correct operation when DR mode load balancing is in use:

- Modifying the server's ARP behaviour and adding the relevant VIP addresses to the loopback interface
- Using NAT to convince the server to accept and reply to packets addressed to the relevant VIP addresses

Four independent methods are described below along with instructions. Each method follows one of the two approaches above. The specific method chosen will depend on technical requirements, the Linux distribution in use, and personal preferences.

The first method involves setting kernel parameters to alter the server's ARP behaviour and adding IP addresses to the loopback interface. This method should be universally applicable to any Linux server **making this the preferred method**.

If setting kernel parameters and adding IP addresses is not possible for some reason, the remaining three methods describe setting up a server for DR mode operation by using NAT via the **redirect** target/statement. The specific instructions depend on the packet filtering framework and tooling in use, which varies between Linux distributions. Methods are presented for iptables, nftables, and the `firewall-cmd` tool.

### Method 1: ARP Behaviour and Loopback Interface Changes

This is the preferred method as it should be applicable to any Linux server and doesn't require any additional packet filtering or NAT considerations.

Each real server needs the loopback interface to be configured with the virtual IP addresses (VIPs) of the relevant load balanced services. This is often just a single VIP address, but the logic described below can be extended to cover multiple VIPs on a server. Having the VIPs on the loopback interface allows the server to accept inbound load balanced packets that are addressed to a VIP.

The server **must not** respond to ARP requests for the VIP addresses. The server also **must not** use ARP to announce the fact that it owns the VIP addresses. This is necessary to prevent IP address conflicts, as **all** of the real servers **and** the load balancer will own the VIP addresses. Only the load balancer should announce ownership of the VIPs.

To configure the behaviour described above, follow all of the steps below on each real server.

#### Step 1 of 4: Re-configuring ARP behaviour

**This step is only applicable if IPv4-based virtual services are in use.**

Add the following lines to the file `/etc/sysctl.conf` (create this file if it does not already exist):

```
net.ipv4.conf.all.arp_ignore=1
net.ipv4.conf.eth0.arp_ignore=1
net.ipv4.conf.eth1.arp_ignore=1
net.ipv4.conf.all.arp_announce=2
net.ipv4.conf.eth0.arp_announce=2
net.ipv4.conf.eth1.arp_announce=2
```



Adjust the commands shown above to suit the server's network configuration, e.g. a different number of network interfaces or a different interface naming convention.

For reference, the effect of these kernel parameter changes on the server is as follows:

#### Note

- **arp\_ignore=1**: This configures the server to only reply to an ARP request if the request's target IP address is local to the incoming interface. This can never be true for VIP addresses on the loopback interface, as the loopback interface can never be an incoming interface for ARP requests from other devices. Hence, ARP requests for VIP addresses are always ignored.
- **arp\_announce=2**: This prevents the server from sending an ARP request out of an interface **A** where the ARP request's sender/source address is stated to be an IP address that is local to some other interface **B**. For example, this prevents the server from sending an ARP request *from* a VIP address (which is local to the loopback interface) out of **eth0**, which would announce that the server owns the VIP address.

### Step 2 of 4: Re-configuring duplicate address detection (DAD) behaviour

This step is only applicable if IPv6-based virtual services are in use.

Add the following lines to the file `/etc/sysctl.conf` (create this file if it does not already exist):

```
net.ipv6.conf.lo.dad_transmits=0
net.ipv6.conf.lo.accept_dad=0
```

For reference, the effect of these kernel parameter changes on the server is as follows:

#### Note

- **dad\_transmits=0**: This prevents a given interface from sending out duplicate address detection probes in order to test the uniqueness of unicast IPv6 addresses. Any IPv6 VIP addresses will *not* be unique, so this mechanism is disabled.
- **accept\_dad=0**: This prevents a given interface from accepting duplicate address detection messages. This prevents any IPv6 VIP addresses from being marked as duplicate addresses.

### Step 3 of 4: Applying the new settings

To apply the new settings, either reboot the real server or execute the following command to immediately apply the changes:

```
/sbin/sysctl -p
```

#### Note

Steps 1, 2, and 3 can be replaced by instead modifying the necessary kernel variables by writing directly to their corresponding files under `/proc/sys/`. Note that changes made in this way *will not persist across reboots*.

Execute the following commands (as root) to implement these temporary changes (adapting the number of interfaces and interface names as needed):



```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_ignore
echo 1 > /proc/sys/net/ipv4/conf/eth0/arp_ignore
echo 1 > /proc/sys/net/ipv4/conf/eth1/arp_ignore
echo 2 > /proc/sys/net/ipv4/conf/all/arp_announce
echo 2 > /proc/sys/net/ipv4/conf/eth0/arp_announce
echo 2 > /proc/sys/net/ipv4/conf/eth1/arp_announce
echo 0 > /proc/sys/net/ipv6/conf/lo/dad_transmits
echo 0 > /proc/sys/net/ipv6/conf/lo/accept_dad
```

#### Step 4 of 4: Adding the virtual IP addresses (VIPs) to the loopback interface

Each of the VIP addresses must be permanently added to the loopback interface. VIPs must be added with a network prefix of /32 for IPv4 addresses or /128 for IPv6 addresses. The IP addresses can be added using the usual configuration files and tools for modifying network interfaces, which vary between different Linux distributions.

**As an alternative**, the `ip` command can be used as a universal way to add IP addresses to any Linux server. Note that addresses added in this way **will not persist across reboots**. To make these addresses permanent, add the `ip` commands to an appropriate startup script such as `/etc/rc.local`.

Execute the following `ip` command for each IPv4 VIP:

```
ip addr add dev lo <IPv4-VIP>/32
```

Execute the following `ip` command for each IPv6 VIP:

```
ip addr add dev lo <IPv6-VIP>/128
```

To check that the VIPs have been successfully added, execute the command:

```
ip addr ls
```

To remove an IPv4 VIP from the loopback adapter, execute the command:

```
ip addr del dev lo <IPv4-VIP>/32
```

To remove an IPv6 VIP from the loopback adapter, execute the command:

```
ip addr del dev lo <IPv6-VIP>/128
```

#### Method 2: NAT "redirect" via iptables

iptables can be used on each real server to identify incoming packets that are addressed to a virtual IP address (VIP) and redirect those packets to the server itself. This is achieved using the **REDIRECT** target in iptables, which performs the necessary NAT to make this possible. This allows a real server to accept packets addressed to a



VIP without the server owning the VIP.

Execute the following command to put the necessary iptables rule in place to redirect traffic for a single IPv4 VIP address. Note that iptables rules added in this way **will not persist across reboots**. To make such a rule permanent, either add the rule to an iptables firewall script, if one is provided with the Linux distribution in question, or add the command to an appropriate startup script such as `/etc/rc.local` on each real server.

```
iptables -t nat -A PREROUTING -d <IPv4-VIP> -j REDIRECT
```

The VIP address should be changed to match the virtual service in question, for example:

```
iptables -t nat -A PREROUTING -d 10.0.0.21 -j REDIRECT
```

The example above will redirect any incoming packets destined for 10.0.0.21 (the virtual service) locally, i.e. to the primary address of the incoming interface on the real server.

If a real server is responsible for serving **multiple** VIPs then additional iptables rules should be added to cover each VIP.

For an IPv6 VIP address, a command like the following should be used:

```
ip6tables -t nat -A PREROUTING -d <IPv6-VIP> -j REDIRECT
```

The VIP address should be changed to match the virtual service in question, for example:

```
ip6tables -t nat -A PREROUTING -d 2001:db8::10 -j REDIRECT
```

#### Note

Method 2 may not be appropriate when using IP-based virtual hosting on a web server. This is because an iptables **REDIRECT** rule will redirect incoming packets to the **primary address** of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 1 instead.

### Method 3: NAT "redirect" via nftables

nftables is the modern Linux kernel packet filtering framework. It is supported on all major Linux distributions and has replaced iptables as the default framework on most major distributions.

nftables can be used on each real server to identify incoming packets that are addressed to a virtual IP address (VIP) and redirect those packets to the server itself. This is achieved using the **redirect** statement in nftables, which performs the necessary NAT to make this possible. This allows a real server to accept packets addressed to a VIP without the server owning the VIP.

Use a script like the following to put the necessary nftables structures in place to redirect traffic for both IPv4 and IPv6 VIP addresses. To make such a configuration permanent, either add the `inet nat` table to an nftables firewall script, if one is provided with the Linux distribution in question, or configure a script like the following to

execute as a startup script on each real server.

```
#!/usr/sbin/nft -f

table inet nat {
    chain prerouting {
        comment "Allow server to accept packets destined for VIP addresses";
        type nat hook prerouting priority -100; policy accept;
        ip daddr <IPv4-VIP> redirect comment "Description"
        ip6 daddr <IPv6-VIP> redirect comment "Description"
    }
}
```

The VIP addresses and comments should be changed to match the virtual services in question, for example:

```
#!/usr/sbin/nft -f

table inet nat {
    chain prerouting {
        comment "Allow server to accept packets destined for VIP addresses";
        type nat hook prerouting priority -100; policy accept;
        ip daddr 10.0.0.21 redirect comment "VIP 1: HTTP"
        ip6 daddr 2001:db8::10 redirect comment "VIP 2: HTTPS"
    }
}
```

The example above will redirect any incoming packets destined for 10.0.0.21 or 2001:db8::10 (the virtual services) locally, i.e. to the primary address of the incoming interface (for each IP version) on the real server.

Note that **Linux kernels prior to 5.2** may not support performing NAT (which is required for the **redirect** statement) in an inet family table. In this scenario, use either an ip or an ip6 family table instead, or both if a mixture of IPv4 and IPv6 VIPs are in use on the same server. Also note that older kernels may not support the use of comments in chains.

Note that **Linux kernels prior to 4.18** require explicitly registering both prerouting and postrouting chains in order for the implicit NAT of the **redirect** statement to be correctly performed in both the inbound and outbound directions.

A legacy-friendly setup may look like the following:

```
#!/usr/sbin/nft -f

table ip nat {
    chain prerouting {
        type nat hook prerouting priority -100; policy accept;
        ip daddr 10.0.0.21 counter redirect comment "VIP 1: HTTP"
    }

    chain postrouting {
        type nat hook postrouting priority 100; policy accept;
    }
}
```

```

}

table ip6 nat {
    chain prerouting {
        type nat hook prerouting priority -100; policy accept;
        ip6 daddr 2001:db8::10 counter redirect comment "VIP 2: HTTPS"
    }

    chain postrouting {
        type nat hook postrouting priority 100; policy accept;
    }
}

```

#### Note

Method 3 may not be appropriate when using IP-based virtual hosting on a web server. This is because an nftables **redirect** statement will redirect incoming packets to the *primary address* of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 1 instead.

### Method 4: NAT "redirect" via firewall-cmd

Some recent versions of Linux distributions make use of firewalld as a high-level firewall configuration framework. In this case, while it may actually be iptables performing the work at a lower level, it may be preferred to implement the iptables NAT solution described in [method 2](#) in firewalld, as opposed to directly manipulating iptables. This is achieved by using the **firewall-cmd** tool provided by firewalld and executing a command like the following on each real server:

```
firewall-cmd --permanent --direct --add-rule ipv4 nat PREROUTING 0 -d <IPv4-VIP> -j REDIRECT
```

The VIP address should be changed to match the virtual service in question, for example:

```
firewall-cmd --permanent --direct --add-rule ipv4 nat PREROUTING 0 -d 10.0.0.50 -j REDIRECT
```

To apply the new configuration, reload the firewall rules like so:

```
firewall-cmd --reload
```

Configuration applied in this way will be permanent and will persist across reboots.

#### Note

Method 4 may not be appropriate when using IP-based virtual hosting on a web server. This is because an iptables **REDIRECT** rule will redirect incoming packets to the *primary address* of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 1 instead.

### Solving the ARP Problem for Mac OS X/BSD

OS X is BSDish, so you need to use BSDish syntax:



```
ifconfig lo0 alias <VIP> netmask 255.255.255.255 -arp up
```

You'll need to add this to the startup scripts on all of your Real Servers.

#### Note

Don't forget that the service on the Real Servers needs to listen on both the RIP address and VIP address as mentioned previously. Failure to correctly configure the Real Servers to handle the ARP problem is the most common mistake in DR mode configurations.

## Windows Server 2012 & Later

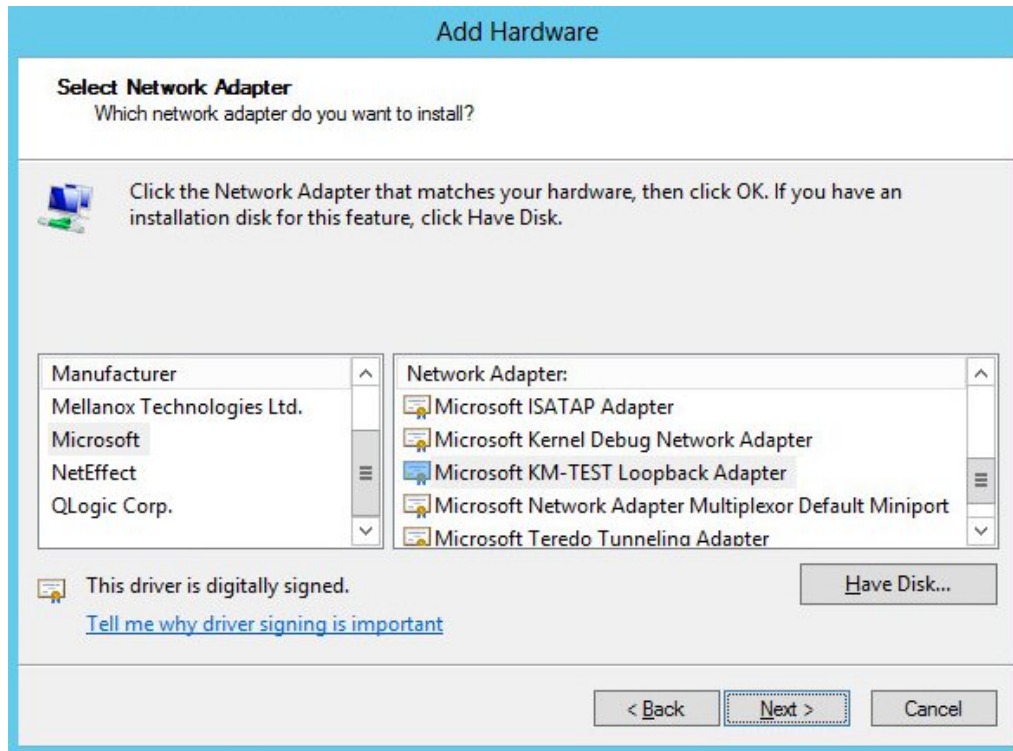
Windows Server 2012 and later support Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter that must be installed and configured on each load balanced (Real) Server. The IP address configured on the Loopback Adapter must be the same as the Virtual Service (VIP) address. This enables the server to receive packets that have their destination set as the VIP address. If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

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

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

### Step 1 of 3: Install the Microsoft Loopback Adapter

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



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

### Step 2 of 3: Configure the Loopback Adapter

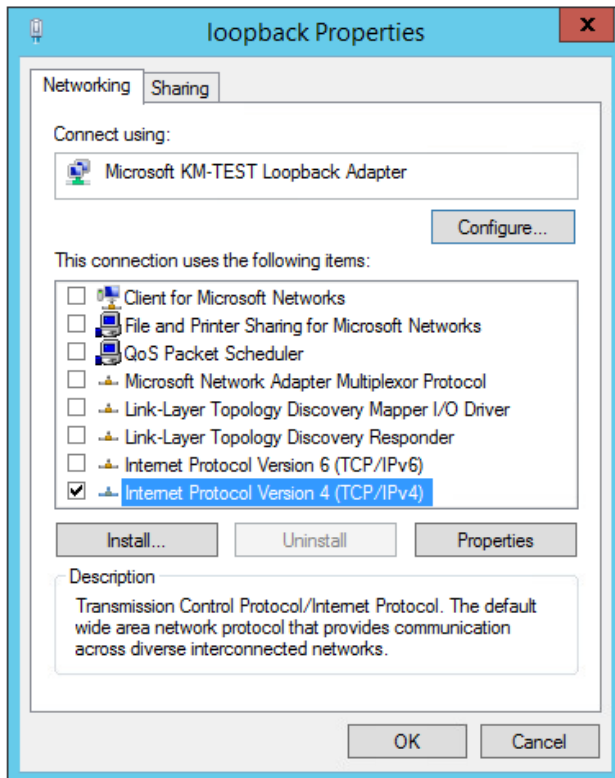
1. Open Control Panel and click **Network and Sharing Center**.
2. Click **Change adapter settings**.
3. Right-click the new Loopback Adapter and select **Properties**.

#### Note

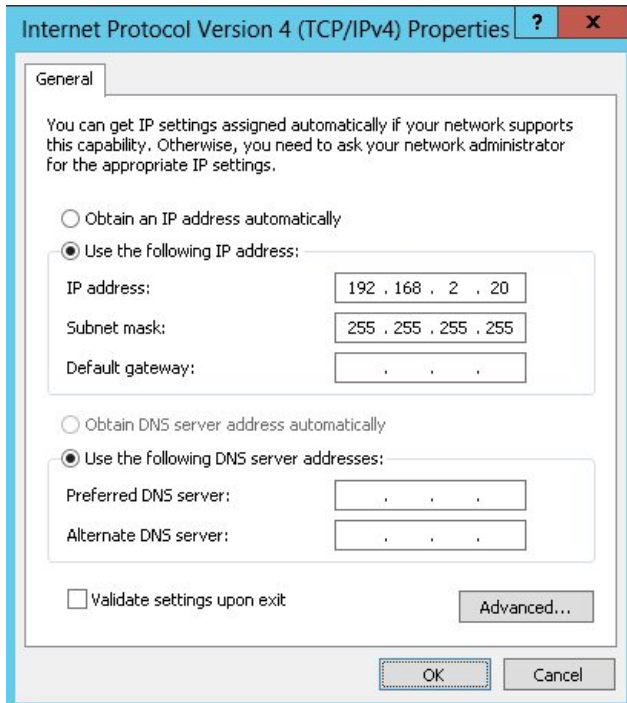
You can configure IPv4 or IPv6 addresses or both depending on your requirements.

### IPv4 Addresses

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



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



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

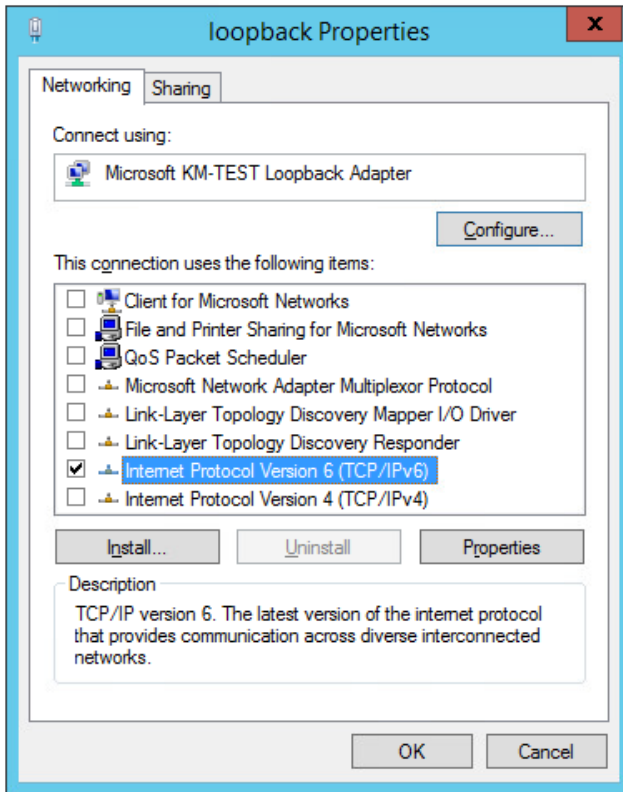
**Note** If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be

added to the Loopback Adapter.

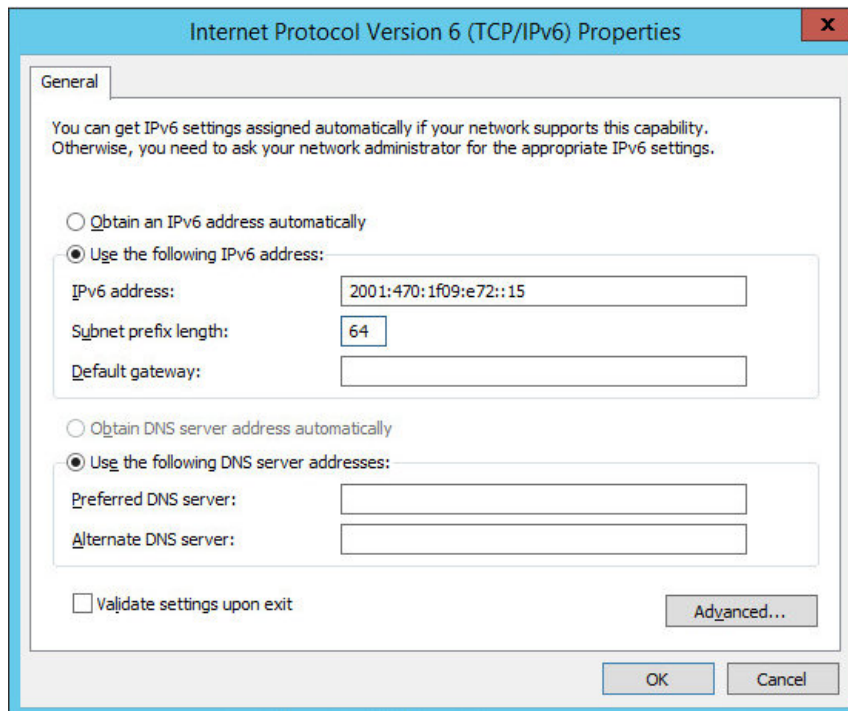
3. Click **OK** then click **Close** to save and apply the new settings.

## 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
```

## 17.4. Fallback Server Settings

The fallback server is activated under the following conditions for both Layer 4 & Layer 7 Virtual Services:

- When all associated Real Servers have failed their health check
- When all associated Real Servers have been taken offline via the WebUI

The fallback page can be provided in the following ways:



- Using the load balancer's built in NGINX fallback page
- Using a separate server to host the fallback page
- Using a Layer 7 VIP

## 17.5. Local Fallback Server

The appliance has a built in fallback server that uses NGINX. The local fallback page can be modified using the WebUI menu option: *Maintenance > Fallback Page*

### Fallback Page

```

1  <html>
2  <head>
3  <title>The page is temporarily unavailable</title>
4  <style>
5  body { font-family: Tahoma, Verdana, Arial, sans-serif; }
6  </style>
7  </head>
8  <body bgcolor="white" text="black">
9  <table width="100%" height="100%">
10 <tr>
11 <td align="center" valign="middle">
12 The page you are looking for is temporarily unavailable.<br/>
13 Please try again later.<br/>
14 (WUI port reminder 9080)
15 </td>
16 </tr>
17 </table>
18 </body>
19 </html>
20
21
22
23

```

The local fallback server is an NGINX instance that by default listens on port 9081.

If a layer 4 VIP is added that listens on port 80, NGINX is automatically configured to listen on ports **9081 & 80**.

#### Note

You can use any valid HTML for the default page, simply copy and paste the required HTML into the Fallback Page.

If you are using the load balancer for your holding page and your Real Servers are all offline then the local NGINX server is exposed to hacking attempts. If you are concerned about this you can change the fallback server to be one of your internal servers.

## 17.6. Using a Separate Dedicated Server

For DR mode the fallback server must be listening on the same port as the VIP (port re-mapping is not possible with DR mode). Also, don't forget to solve the ARP problem for the dedicated fallback server.

## 17.7. Using a Layer 7 VIP

It's possible to set the fallback server to be a layer 7 VIP. This is especially useful in WAN/DR site environments. It also enables an external fallback server to be easily configured for Layer 4 VIPs without having to comply with the requirements mentioned in the previous section. To do this, create a layer 7 fallback VIP and configure your fallback server as an associated RIP. Then enable the MASQ option for the Layer 4 VIP and set the fallback VIP as

its fallback server. If all servers are down, requests will then be routed via the Layer 7 VIP to your fallback server. If the layer 4 VIP is multi-port, specify 0 as the port for the fallback server. Requests will then be forwarded to the correct port.

## 17.8. Configuring A real Server as the Fallback Server

It's possible to configure one of the Real Servers as the fallback server. This can be useful for example when all servers are very busy and health checks start to fail simply because the response is taking longer than the configuration allows. In this case, traffic will still be sent to one of the Real Servers rather than to a separate fallback page.

## 17.9. Configuring Primary / Secondary Real Servers

If you want to setup a VIP that sends all traffic to a primary server and only sends traffic to a secondary server if the primary server fails, configure the VIP with the primary server as a RIP, and the secondary server as the fallback server.



## 17.10. 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	<p>Combined software version information into one section</p> <p>Added one level of section numbering</p> <p>Added software update instructions</p> <p>Added table of ports used by the appliance</p> <p>Reworded 'Further Documentation' section</p>	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	<p>New document theme</p> <p>Modified diagram colours</p>	Branding update	AH



**Visit us:** [www.loadbalancer.org](http://www.loadbalancer.org)

**Phone us:** +44 (0)330 380 1064

**Phone us:** +1 833 274 2566

**Email us:** [info@loadbalancer.org](mailto:info@loadbalancer.org)

**Follow us:** [@loadbalancer.org](https://twitter.com/loadbalancer.org)

## About Loadbalancer.org

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

