

Load Balancing Smoothwall Secure Web Gateway

Version 1.6.1



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. Smoothwall Web Gateway	4
4. Loadbalancer.org & Smoothwall	4
5. Benefits of Implementing a Load Balancer	4
6. Load Balancer Configuration Options	5
6.1. Deployment Modes	5
6.1.1. Layer 4 DR Mode	5
6.1.2. Layer 4 NAT Mode	6
6.2. Persistence / Server Affinity	8
6.2.1. Source IP Address (Recommended)	8
6.2.2. Destination Hash	8
7. Web Gateway Deployment Modes	8
7.1. 1 – Explicit Proxy Mode (Recommended)	9
7.2. 2 – Transparent Routed Proxy Mode	9
8. Summary of Deployment Options	9
9. Loadbalancer.org Appliance – the Basics	10
9.1. Virtual Appliance	10
9.2. Initial Network Configuration	10
9.3. Accessing the Appliance WebUI	10
9.3.1. Main Menu Options	12
9.4. Appliance Software Update	13
9.4.1. Online Update	13
9.4.2. Offline Update	13
9.5. Ports Used by the Appliance	14
9.6. HA Clustered Pair Configuration	15
10. Option 1 – Explicit Proxy Mode (Recommended)	15
10.1. Option 1A – Using DR (Direct Return) Mode (Recommended)	15
10.1.1. Deployment Architecture	15
10.1.2. Load Balancer Configuration	16
10.1.3. Web Gateway Configuration	17
10.1.4. Finalize Settings	18
10.2. Option 1B – Using NAT Mode	18
10.2.1. Deployment Architecture	18
10.2.2. Load Balancer Configuration	20
10.2.3. Web Gateway Configuration	23
10.2.4. Finalize Settings	23
10.3. Option 1C – Using NAT Mode (Preferred NAT Topology)	23
10.3.1. Deployment Architecture	23
10.3.2. Load Balancer Configuration	25
10.3.3. Web Gateway Configuration	27
10.3.4. Finalize Settings	28
10.4. Configuration Settings Common to Options 1A, 1B & 1C	28
10.4.1. Web Gateway Operating Mode	28
10.4.2. Proxy Port Configuration	28
10.4.3. Client Configuration	29

11. Option 2 - Transparent Routed Proxy Mode	30
11.1. Deployment Architecture	30
11.2. Load Balancer Configuration	31
11.2.1. Create the Virtual Service (VIP)	32
11.2.2. Add the Floating IP	33
11.2.3. Configure Firewall Rules	33
11.2.4. Define the Real Servers (RIPs)	34
11.3. Web Gateway Configuration	35
11.3.1. Web Gateway Operating Mode	35
11.4. Router/Default Gateway Configuration	35
11.5. Client Configuration	36
12. Testing & Verification	36
12.1. Layer 4 – Current Connections	36
12.1.1. Explicit Proxy Mode	36
12.1.2. Transparent Mode	37
13. Technical Support	37
14. Further Documentation	37
15. Appendix	38
15.1. Configuring HA - Adding a Secondary Appliance	38
15.1.1. Non-Replicated Settings	38
15.1.2. Configuring the HA Clustered Pair	39
15.2. Modified Transparent Mode Firewall Rules	40
15.3. Configuring Email Alerts for Virtual Services	41
15.3.1. Layer 4	41
15.4. Configuring Email Alerts for Heartbeat	42
15.5. Configuring a Smart Host (SMTP relay)	43
16. Document Revision History	44

1. About this Guide

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

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

2. Loadbalancer.org Appliances Supported

All our products can be used with Smoothwall Secure Web Gateway. 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. Smoothwall Web Gateway

- All versions

4. Loadbalancer.org & Smoothwall

Loadbalancer.org and Smoothwall have partnered to provide high performance, robust and highly available Web Filtering solutions that enable customers to deploy with confidence.

5. Benefits of Implementing a Load Balancer

Implementing Loadbalancer.org appliances enables multiple Smoothwall Web Gateways to be deployed in a cluster. This provides the following key benefits:

- **High-Availability** – If a Web Gateway fails, service is not interrupted
- **Maintenance** – Web Gateways can easily be taken out of the cluster for maintenance
- **Performance** – For additional performance simply add more Web Gateways to the cluster



6. Load Balancer Configuration Options

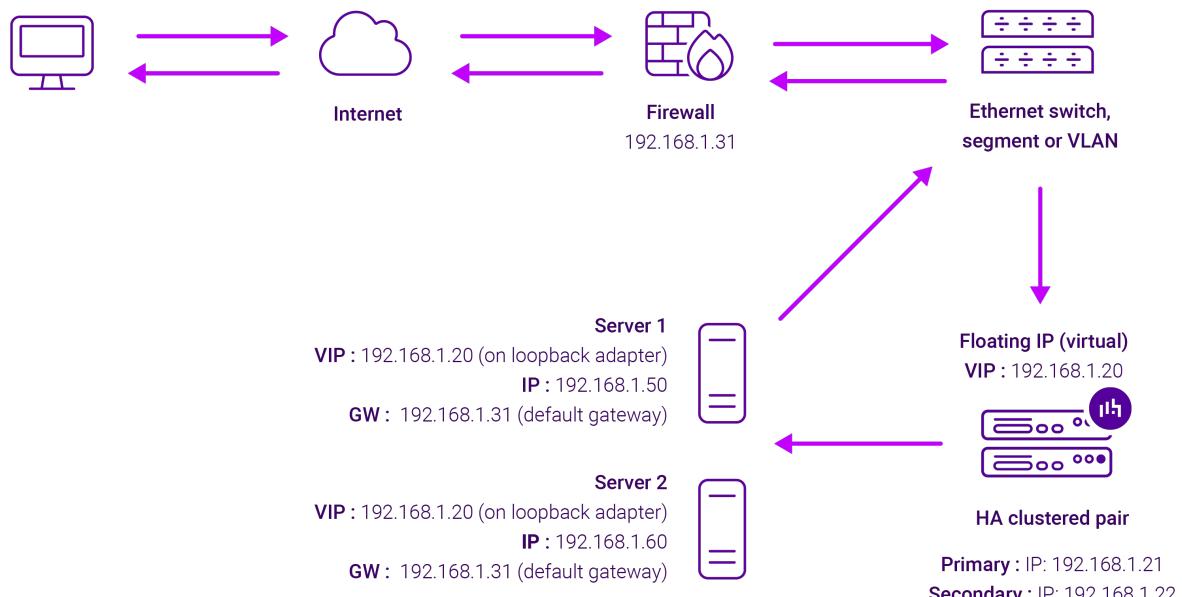
The following sections describe the various load balancer deployment modes and persistence options that are used when load balancing Smoothwall Web Gateways.

6.1. Deployment Modes

6.1.1. 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*.

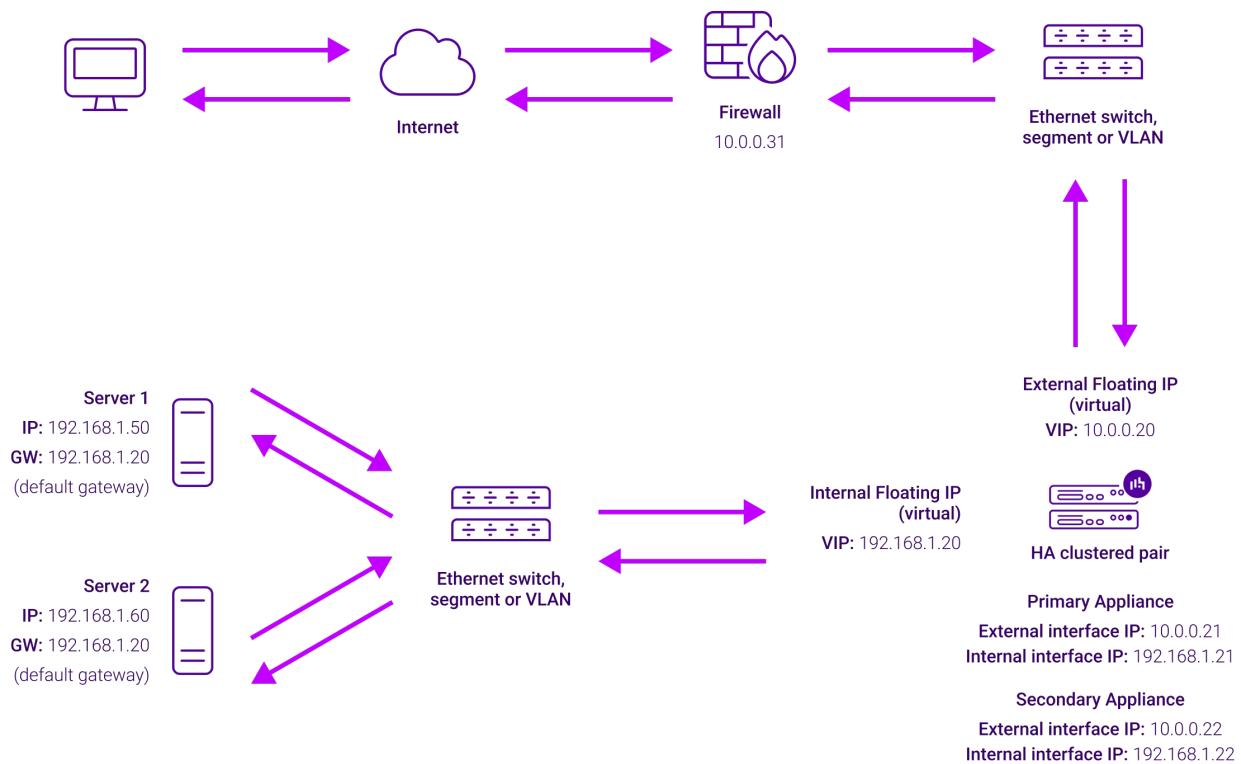


- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that each Real Server (and the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

6.1.2. Layer 4 NAT Mode

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



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

Note

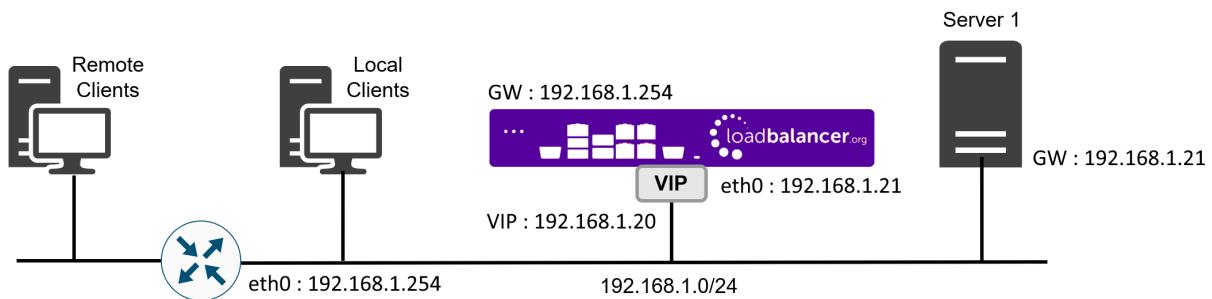
This can be achieved by using two network adapters, or by creating VLANs on a single adapter.

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

Note

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

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



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

Note

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

- To support local clients, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer. For more information please refer to [One-Arm \(Single Subnet\) NAT Mode](#).
- If you want Real Servers to be accessible on their own IP address for non-load balanced services, e.g. RDP, you will need to setup individual SNAT and DNAT firewall script rules for each Real Server or add additional VIPs for this.
- Port translation is possible with Layer 4 NAT mode, e.g. VIP:80 → RIP:8080 is supported.
- NAT mode is transparent, i.e. the Real Server will see the source IP address of the client.

NAT Mode Packet re-Writing

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

The following table shows an example NAT mode setup:

Protocol	VIP	Port	RIP	Port
TCP	10.0.0.20	80	192.168.1.50	80



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

Packet rewriting works as follows:

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

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

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

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

3) Replies return to the load balancer as:

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

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

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

6.2. Persistence / Server Affinity

Persistence may or may not be required and depends on the specific Web Gateway being used. Two possible methods are described in the following sections.

6.2.1. Source IP Address (Recommended)

Source IP persistence is the default option for Layer 4 services. When set, clients connecting from the same source IP address within the persistence timeout period (the default is 5 minutes) will always be sent to the same Web Gateway. It's recommended that this should be set to 1 hour minimum.

6.2.2. Destination Hash

Another option at Layer 4 is to change the load balancing algorithm (i.e. the "scheduler") to destination hash (DH). This causes the load balancer to select the Web Gateway based on a hash of the destination IP address. This causes session requests to be directed at the same server based solely on the destination IP address of a packet which therefore makes client connections persistent for a particular Internet host.

Since this setting is a scheduler, the way connections are load balanced will also change. However it should still provide a well balanced distribution of client sessions between the Web Gateways.

7. Web Gateway Deployment Modes

There are two implementation methods that are typically used – Non Transparent Explicit Proxy Mode & Transparent Routed Proxy Mode.



7.1. 1 – Explicit Proxy Mode (Recommended)

This mode requires the load balancer's VIP address to be defined in users' browsers. This means that the load balancer will receive client requests and distribute these requests across the back-end Web Gateways. Smoothwall refer to this as "Non-Transparent Mode". Please refer to the section [Option 1 – Explicit Proxy Mode \(Recommended\)](#) for configuration details.

7.2. 2 – Transparent Routed Proxy Mode

With this mode, client requests must be routed to the load balancer/Web Gateway cluster. This can be achieved by either setting the default gateway on the client PCs to be the load balancer, or by adding rules to the default gateway device. Rules would typically be configured for HTTP & HTTPS traffic on ports 80 and 443. Smoothwall refer to this as "Transparent Mode". Please refer to the section [Option 2 - Transparent Routed Proxy Mode](#) for configuration details.

8. Summary of Deployment Options

Option	Web Gateway Mode	Load Balancer Mode	Notes
Option 1A (Recommended)	Non-Transparent Mode	DR Mode	<p>The Web Gateways must be configured to accept traffic for the VIP.</p> <p>Please refer to Option 1 – Explicit Proxy Mode (Recommended) for configuration details.</p>
Option 1B	Non-Transparent Mode	NAT Mode	<p>The load balancer must be set as the default gateway for the Web Gateways.</p> <p>Please refer to Option 1B – Using NAT Mode for configuration details.</p>
Option 1C	Non-Transparent Mode	NAT Mode	<p>A static route must be configured on the Web Gateways to send client return traffic back via the load balancer.</p> <p>Please refer to Option 1C – Using NAT Mode (Preferred NAT Topology) for configuration details.</p>



Option	Web Gateway Mode	Load Balancer Mode	Notes
Option 2	Transparent Mode	DR Mode	<p>Firewall rules must be added to the load balancer to transparently send traffic to the Web Gateways.</p> <p>Please refer to Option 2 - Transparent Routed Proxy Mode for configuration details.</p>

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](#).



i 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/

i 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](#).

i 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>

i Note

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

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



Primary | Secondary Active | Passive Link 8 Seconds

System Overview

Local Configuration

Cluster Configuration

Maintenance

View Configuration

Reports

Logs

Support

Live Chat

WARNING: YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS.
Buy with confidence. All purchases come with a 90 day money back guarantee.
Already bought? Enter your license key [here](#)

Buy Now

System Overview 2025-05-08 12:37:21 UTC

Would you like to run the Setup Wizard?

Accept **Dismiss**

VIRTUAL SERVICE **IP** **PORTS** **CONN** **PROTOCOL** **METHOD** **MODE**

No Virtual Services configured.

Network Bandwidth

Bytes/s

200 k, 150 k, 100 k, 50 k, 0

Wed 18:00, Thu 00:00, Thu 06:00, Thu 12:00

RX 28 Min, 2713 Avg, 27344772 Total, TX 0 Min, 13777 Avg, 138872181 Total

System Load Average

System Load

1.0, 0.8, 0.6, 0.4, 0.2, 0.0

Wed 18:00, Thu 00:00, Thu 06:00, Thu 12:00

1m average 0.00 Min, 0.08 Avg, 0.68 Max, 5m average 0.00 Min, 0.04 Avg, 0.30 Max, 15m average 0.00 Min, 0.02 Avg, 0.12 Max

Memory Usage

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.



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, RIPv 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 & RIPv

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:

Information: Update 8.13.2 is now available for this appliance.

Online Update

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:

Information: Update completed successfully. Return to system overview.

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:

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

Upload and Install

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
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](#)



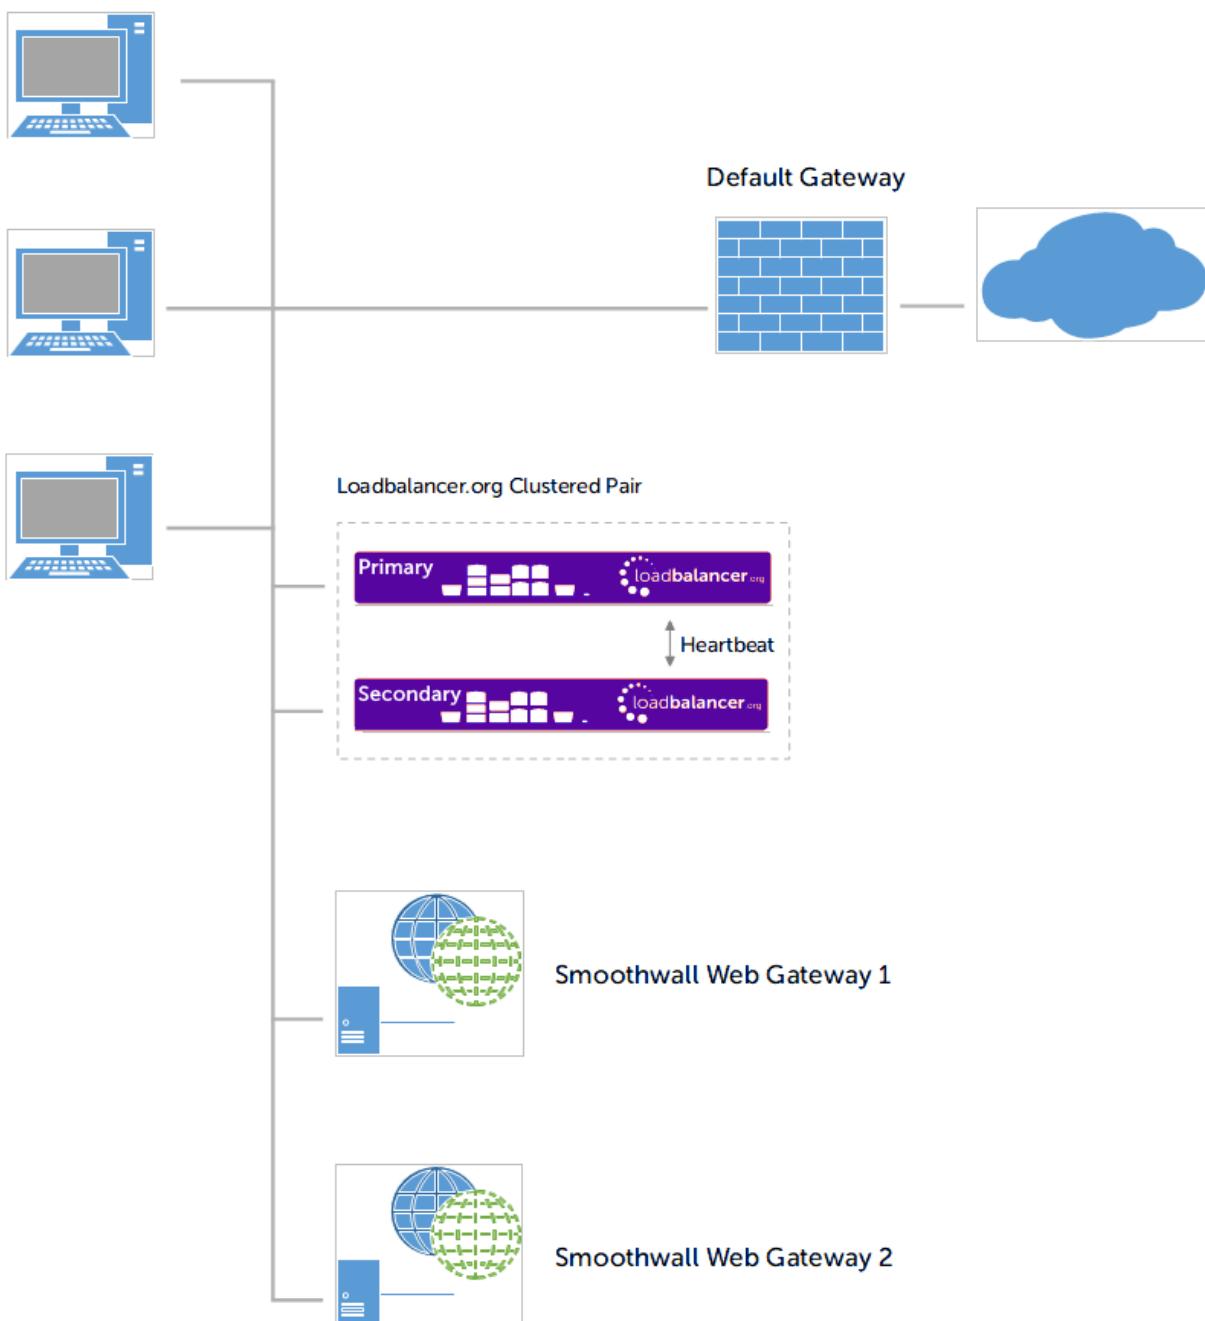
9.6. HA Clustered Pair Configuration

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. Option 1 – Explicit Proxy Mode (Recommended)

10.1. Option 1A – Using DR (Direct Return) Mode (Recommended)

10.1.1. Deployment Architecture



Notes

- Browser settings on client PC's must be changed to point at the Virtual Service (VIP) on the load balancer (see [Client Configuration](#))
- The load balancers are configured in one-arm Layer 4 DR mode
- The Smoothwall Web Gateways must be configured to accept traffic for the VIP (see [Web Gateway 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](#)

10.1.2. Load Balancer Configuration

Create the Virtual Service (VIP)

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

Label	Proxy	?	
Virtual Service	IP Address	192.168.2.202	?
	Ports	800	?
Protocol	TCP	?	
Forwarding Method	Direct Routing	?	
Cancel Update			

4. Enter an appropriate label (name) for the VIP, e.g. **Proxy**.
5. Set the **Virtual Service IP address** field to the required IP address, e.g. **192.168.2.202**.
6. Set the **Virtual Service Ports** field to the required port, e.g. **800**.

It's possible to use * in the **Ports** field instead of a specific port or list/range of ports.

Note

Using * allows the virtual service to load balance traffic on all ports, which can be useful if the service uses many different ports. It also removes the need to make changes on the load balancer if additional ports are added to the Smoothwall servers in the future.

7. Ensure that **Protocol** is set to **TCP**.
8. Ensure that **Forwarding Method** is set to **Direct Routing**.
9. Click **Update**.
10. Now click **Modify** next to the newly created VIP.



11. Ensure **Persistence** is enabled and set the **Persistence Timeout** to **3600** (i.e. 1 hour).

It's **optionally** possible to define the parent node of the Smoothwall cluster as a fallback server for the new virtual service.

Note

Smoothwall parent nodes do not ordinarily process client traffic and are designed to handle logging and cluster configuration only. It may be desirable, however, for the load balancer to direct client traffic to the parent node in the event that **all** real servers, i.e. all Smoothwall child nodes, fail health checking and are marked as offline.

To implement this, enter the IP address of the parent node in the **IP Address** field under the **Fallback Server** section, and leave the **Port** field under the **Fallback Server** section empty.

12. Click **Update**.

Define the Real Servers (RIPs)

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

Label	Proxy1	?
Real Server IP Address	192.168.2.210	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?
		Cancel Update

4. Enter an appropriate label (name) for the first Web Gateway, e.g. **Proxy1**.
5. Change the **Real Server IP Address** field to the required IP address, e.g. **192.168.2.210**.
6. Click **Update**.
7. Repeat the above steps to add your other Web Gateway(s).

Note

If you'd like to configure email alerts for the Virtual Service, please refer to [Configuring Email Alerts for Virtual Services](#) in the appendix.

10.1.3. Web Gateway Configuration

Modify the Web Gateways to accept traffic for the VIP

Concept

As mentioned previously, DR mode is our recommended load balancer operating mode. To use this mode, changes are required to the real servers, i.e. the Web Gateways. The real servers must accept traffic for the VIP,



but they must not respond to any ARP requests for that IP, only the VIP should do this.

Using the Smoothwall WebUI

To configure the Smoothwall appliance for load balancing use the WebUI option: *Web Proxy > Settings > Advanced*, then enter the required Virtual Service (VIP) IP address as shown below:

Load balancing

Virtual IPs must be used with an external load balancing device.

Direct Server Return Virtual IP:

Search / Add entries
192.168.2.202



Edit in full-text mode

1 entries

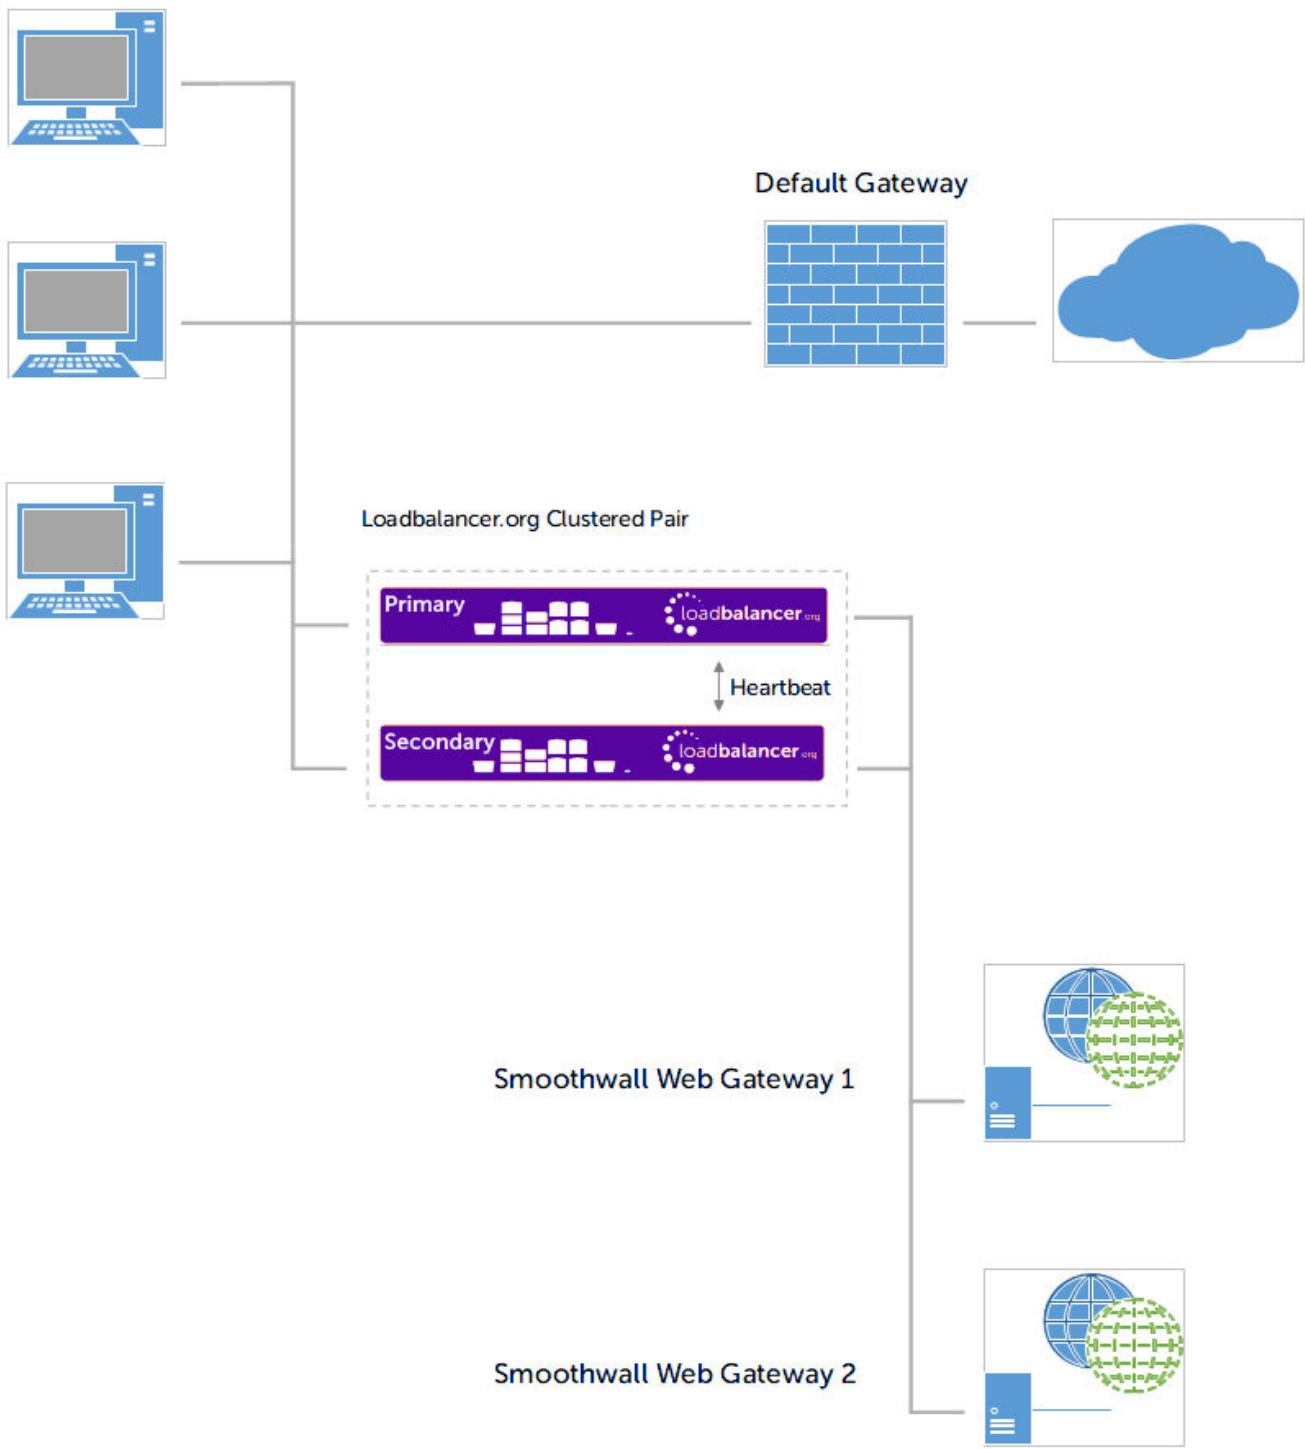
10.1.4. Finalize Settings

Now refer to the section [Configuration Settings Common to Options 1A, 1B & 1C](#) to finalize Web Gateway settings and configure client browser settings.

10.2. Option 1B – Using NAT Mode

10.2.1. Deployment Architecture





Notes

- Browser settings on client PC's must be changed to point at the Virtual Service (VIP) on the load balancer (see [Client Configuration](#))
- The load balancer is configured in two-arm Layer 4 NAT mode
- Return traffic MUST pass back via the load balancer. To enable this, the default gateway for the Web Gateways is configured to be the load balancer. For an HA pair, a floating IP address must be configured to allow the gateway to move between Primary and Secondary in the event of a failover (see [Define a Floating IP to be used as the Default Gateway for the Web Gateways](#))

- 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](#)

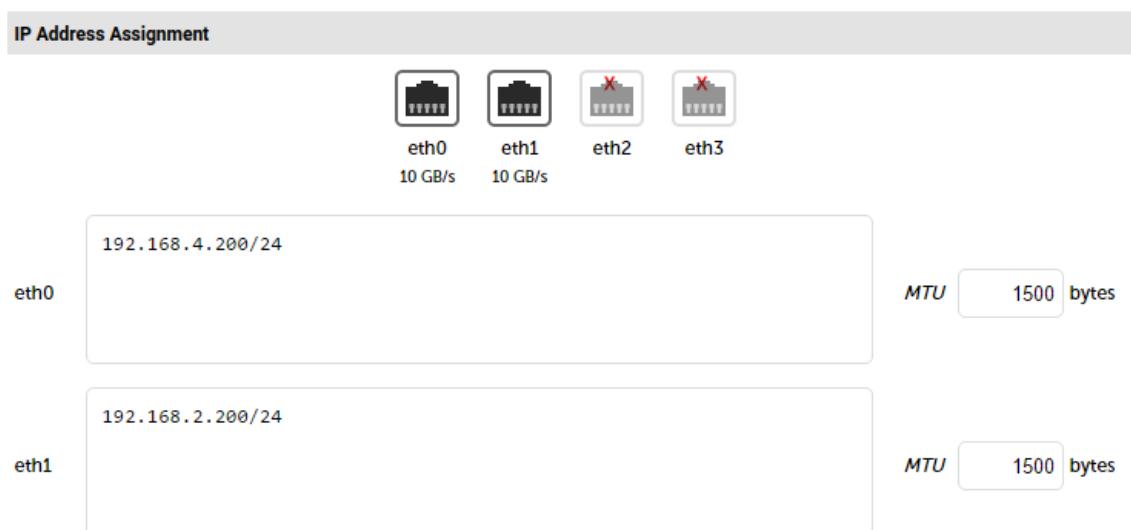
10.2.2. Load Balancer Configuration

Configure Network Settings

Two interfaces are required. Typically **eth0** is used for the internal (Web Gateway) subnet and **eth1** is used for the external (client & VIP) subnet, although this is not mandatory since interfaces can be used as required / preferred.

To configure network settings on the load balancer:

1. Ensure that the required cables are plugged in (hardware) or virtual NICs are connected (virtual).
2. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.
3. Define the required IP addresses and subnet mask:



4. Configure the required IP address for **eth0**, e.g. **192.168.4.200/24**.
5. Configure the required IP address for **eth1**, e.g. **192.168.2.200/24**.
6. Click **Configure Interfaces**.

Define a Floating IP to be used as the Default Gateway for the Web Gateways

As mentioned, when using a clustered pair of load balancers for HA (our recommended configuration), a floating IP must be used as the default gateway for the Web Gateways. This will "float" between the Primary and Secondary units in the event of a failover or failback. This ensures that the Web Gateways always have a consistent return path via the load balancer – whether the Primary or Secondary is active.

To configure a Floating IP:

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IP's*.



New Floating IP Add Floating IP

2. Define a suitable IP address for the default gateway , e.g. **192.168.4.205**.
3. Click **Add Floating IP**.

Create the Virtual Service (VIP)

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

Label	<input type="text" value="Proxy"/>	?
Virtual Service	IP Address <input type="text" value="192.168.2.202"/>	?
	Ports <input type="text" value="8080"/>	?
Protocol	<input type="text" value="TCP"/>	?
Forwarding Method	<input type="text" value="NAT"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

4. Enter an appropriate label (name) for the VIP, e.g. **Proxy**.
5. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.2.202**.
6. Set the *Virtual Service Ports* field to the required port, e.g. **8080**.

It's possible to use * in the *Ports* field instead of a specific port or list/range of ports.

Note Using * allows the virtual service to load balance traffic on all ports, which can be useful if the service uses many different ports. It also removes the need to make changes on the load balancer if additional ports are added to the Smoothwall servers in the future.

7. Ensure that *Protocol* is set to **TCP**.
8. Ensure that *Forwarding Method* is set to **NAT**.
9. Click **Update**.
10. Now click **Modify** next to the newly created VIP.
11. Ensure *Persistence* is enabled and set the *Persistence Timeout* to **3600** (i.e. 1 hour).

Note It's **optionally** possible to define the parent node of the Smoothwall cluster as a fallback server for the new virtual service.



Smoothwall parent nodes do not ordinarily process client traffic and are designed to handle logging and cluster configuration only. It may be desirable, however, for the load balancer to direct client traffic to the parent node in the event that **all** real servers, i.e. all Smoothwall child nodes, fail health checking and are marked as offline.

To implement this, enter the IP address of the parent node in the **IP Address** field under the **Fallback Server** section, and set the **Port** field under the **Fallback Server** section to the required port (if **multiple** ports must be supported, set the **Port** field value to 0).

12. Click **Update**.

Define the Real Servers (RIPs)

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

Label	Proxy1	?
Real Server IP Address	192.168.4.210	?
Real Server Port	8080	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?
		Cancel Update

4. Enter an appropriate label (name) for the first Web Gateway, e.g. **Proxy1**.
5. Set the **Real Server IP Address** field to the required IP address, e.g. **192.168.4.210**.
6. Set the **Real Server Port** field to the required port, e.g. **8080**.
7. Click **Update**.
8. Repeat the above steps to add your other Web Gateway(s).

Enable Auto-NAT

By default, servers behind the load balancer in a NAT configuration will not have access to the outside network. By enabling Auto-NAT, servers (i.e. the Web Gateways) will have their requests automatically mapped to the load balancer's external IP address. The default configuration is to map all requests originating from internal network **eth0** to the external IP on **eth1**. A different interface can be selected if required.

To enable Auto-NAT on the load balancer:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Advanced configuration*.



Email Alert Destination Address

Auto-NAT

Multi-threaded

Update

- Set the Auto-NAT field to the external interface. As mentioned the default configuration is to use **eth1** and the external interface and **eth0** as the internal interface, but can be set to suit your needs.
- Click **Update**.

Note

If you'd like to configure email alerts for the Virtual Service, please refer to [Configuring Email Alerts for Virtual Services](#) in the appendix.

10.2.3. Web Gateway Configuration

Configure the Default Gateway

As mentioned, Option 1B requires the default gateway on the Web Gateway to be the load balancer. When using an HA pair of load balancers, the gateway on the load balancer must be a Floating IP to provide a consistent return path via the load balancer – whether the Primary or Secondary is active. [Define a Floating IP to be used as the Default Gateway for the Web Gateways](#) details how to create the Floating IP.

Note

Please refer to the Smoothwall Web Gateway documentation for instructions on setting the default gateway. This should be done on all Web Gateways.

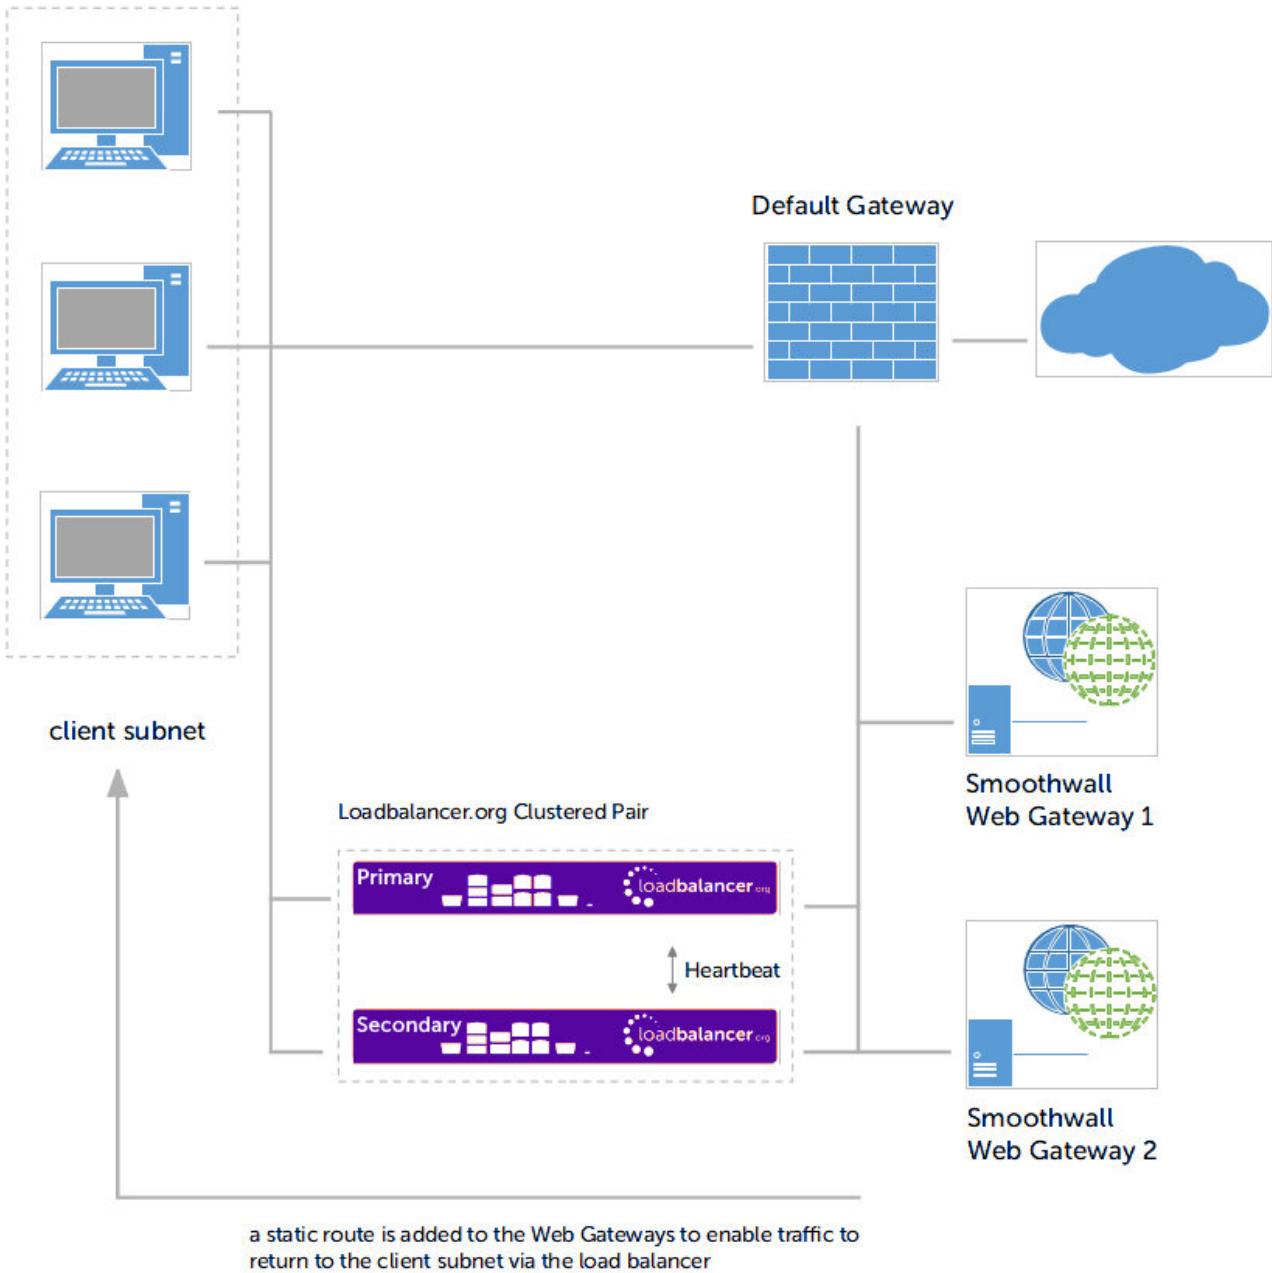
10.2.4. Finalize Settings

Now refer to the section [Configuration Settings Common to Options 1A, 1B & 1C](#) to finalize Web Gateway settings and configure client browser settings.

10.3. Option 1C – Using NAT Mode (Preferred NAT Topology)

10.3.1. Deployment Architecture





Notes

- Browser settings on client PC's must be changed to point at the Virtual Service (VIP) on the load balancer (see [Client Configuration](#))
- The load balancers are configured in two-arm Layer 4 NAT mode
- Return traffic MUST pass back via the load balancer. To enable this, a static route is configured on the Web Gateways to send return traffic back via the load balancer. For an HA pair, a floating IP address must be configured to allow the gateway to move between Primary and Secondary in the event of a failover (see [Define a Floating IP to be used as the gateway for the Static Route on the Web Gateways](#))
- This method is more efficient & faster than Option 1B since the Web Gateways can access the Internet directly rather than going via the load balancer
- 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](#)



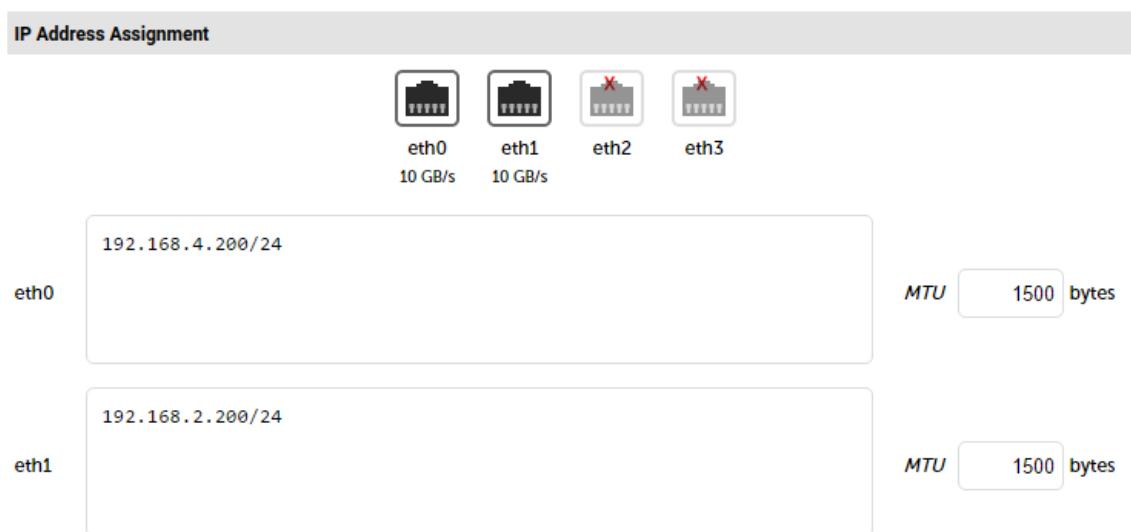
10.3.2. Load Balancer Configuration

Configure Network Settings

Two interfaces are required. Typically **eth0** is used for the internal (Web Gateway) subnet and **eth1** is used for the external (client & VIP) subnet, although this is not mandatory since interfaces can be used as required / preferred.

To configure network settings on the load balancer:

1. Ensure that the required cables are plugged in (hardware) or virtual NICs are connected (virtual).
2. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.
3. Define the required IP addresses and subnet mask:



4. Configure the required IP address for **eth0**, e.g. **192.168.4.200/24**.
5. Configure the required IP address for **eth1**, e.g. **192.168.2.200/24**.
6. Click **Configure Interfaces**.

Define a Floating IP to be used as the gateway for the Static Route on the Web Gateways

As mentioned, when using a clustered pair of load balancers for HA (our recommended configuration), a floating IP must be used as the gateway for the static route on the Web Gateways. This will "float" between the Primary and Secondary units in the event of a failover or failback. This ensures that the Web Gateways always have a consistent return path via the load balancer – whether the Primary or Secondary is active.

To configure a Floating IP:

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IP's*.

New Floating IP	192.168.4.205
Add Floating IP	



2. Define a suitable IP address for the default gateway, e.g. **192.168.4.205**.

3. Click **Add Floating IP**.

Create the Virtual Service (VIP)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Services*.

2. Click **Add a New Virtual Service**.

3. Enter the following details:

Label	Proxy	?	
Virtual Service	IP Address	192.168.2.202	?
	Ports	8080	?
Protocol	TCP	?	
Forwarding Method	NAT	?	
Cancel Update			

4. Enter an appropriate label (name) for the VIP, e.g. **Proxy**.

5. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.2.202**.

6. Set the *Virtual Service Ports* field to the required port, e.g. **8080**.

It's possible to use * in the *Ports* field instead of a specific port or list/range of ports.

Note

Using * allows the virtual service to load balance traffic on all ports, which can be useful if the service uses many different ports. It also removes the need to make changes on the load balancer if additional ports are added to the Smoothwall servers in the future.

7. Ensure that *Protocol* is set to **TCP**.

8. Ensure that *Forwarding Method* is set to **NAT**.

9. Click **Update**.

10. Now click **Modify** next to the newly created VIP.

11. Ensure *Persistence* is enabled and set the *Persistence Timeout* to **3600** (i.e. 1 hour).

It's **optionally** possible to define the parent node of the Smoothwall cluster as a fallback server for the new virtual service.

Note

Smoothwall parent nodes do not ordinarily process client traffic and are designed to handle logging and cluster configuration only. It may be desirable, however, for the load balancer to direct client traffic to the parent node in the event that **all** real servers, i.e. all Smoothwall child nodes, fail health checking and are marked as offline.

To implement this, enter the IP address of the parent node in the *IP Address* field under the



Fallback Server section, and set the *Port* field under the *Fallback Server* section to the required port (if **multiple** ports must be supported, set the *Port* field value to 0).

12. Click **Update**.

Define the Real Servers (RIPs)

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

Label	Proxy1	?
Real Server IP Address	192.168.4.210	?
Real Server Port	8080	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?
		Cancel Update

4. Enter an appropriate label (name) for the first Web Gateway, e.g. **Proxy1**.
5. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.4.210**.
6. Set the *Real Server Port* field to the required port, e.g. **8080**.
7. Click **Update**.
8. Repeat the above steps to add your other Web Gateway(s).

Note

If you'd like to configure email alerts for the Virtual Service, please refer to [Configuring Email Alerts for Virtual Services](#) in the appendix.

10.3.3. Web Gateway Configuration

Configure a Static Route

As mentioned, Option 1C requires a Static Route to be defined on the Web Gateway that forces client return traffic to pass back via the load balancer. When using an HA pair of load balancers, the gateway for the static route must be a Floating IP to provide a consistent return path via the load balancer – whether the Primary or Secondary is active. Define a Floating IP to be used as the gateway for the Static Route on the Web Gateways details how to create the Floating IP.

Note

Please refer to the Smoothwall Web Gateway documentation for instructions on configuring a Static Route. This should be done on all Web Gateways.



10.3.4. Finalize Settings

Now refer to the section [Configuration Settings Common to Options 1A, 1B & 1C](#) to finalize Web Gateway and client browser settings.

10.4. Configuration Settings Common to Options 1A, 1B & 1C

The steps in the following 3 sub sections must be followed for options 1A, 1B & 1C.

10.4.1. Web Gateway Operating Mode

The Smoothwall Web Gateway can easily be configured for client configured Explicit Proxy Mode using the policy wizard. Use the WebUI option: *Web Proxy > Authentication > Policy Wizard* and select "Non-transparent" as shown below:

Step 1: What

Type:	<input checked="" type="radio"/> Non-transparent
	<input type="radio"/> Transparent
Method:	No authentication
Interface:	Port 1 (192.168.66.40)
Port:	800

Next

Now click **Next** to run through the wizard and configure the remaining settings and apply the policy.

10.4.2. Proxy Port Configuration

The required proxy port can be set using the WebUI option: *Web Proxy > Authentication > Policy Wizard* as shown below:

Step 1: What

Type:	<input checked="" type="radio"/> Non-transparent
	<input type="radio"/> Transparent
Method:	No authentication
Interface:	Port 1 (192.168.66.40)
Port:	800

Step 2: Where

Step 3: Options for unauthenticated requests

Where unauthenticated requests are allowed, such as for categories in the exceptions page, they are assigned to a specific group. Available groups: 8000, 8080, 8081, 8888, 9000, 9090, 9091, 10080

Included groups:

Now click **Next** to run through the wizard and configure the remaining settings and apply the policy.



Note

The default proxy port for Smoothwall Web Gateway is 800.

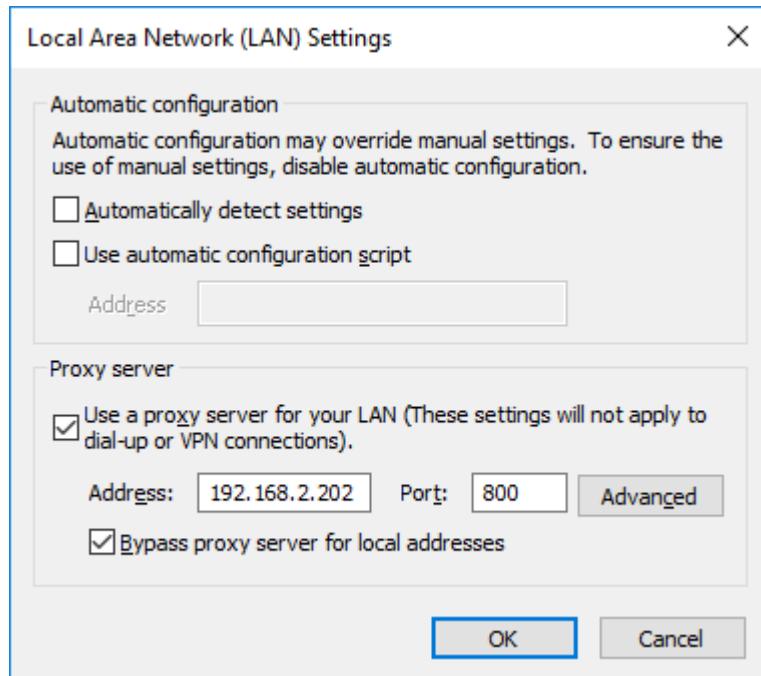
10.4.3. Client Configuration

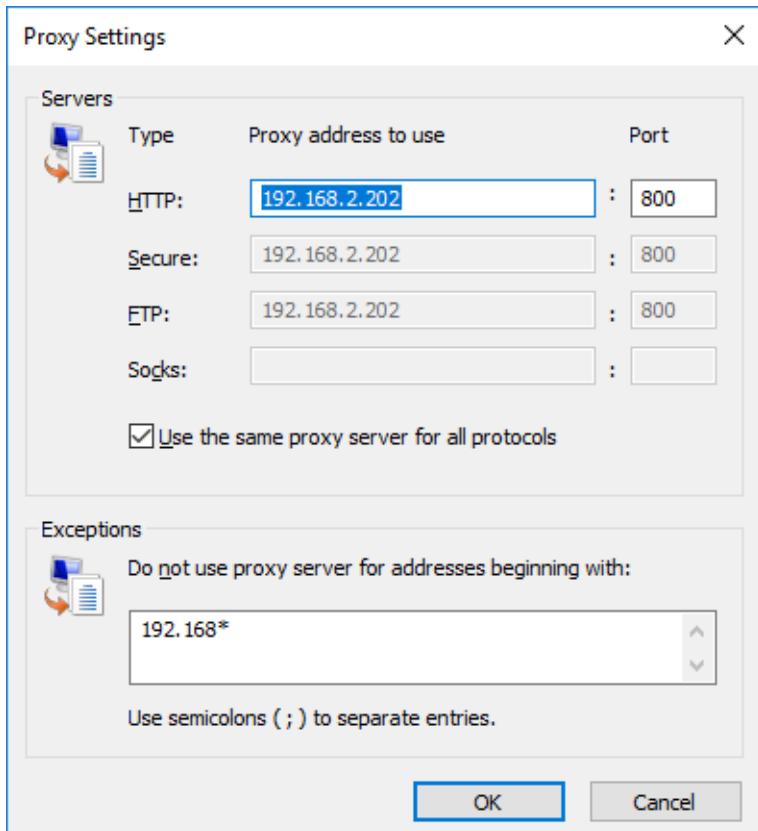
Client browser settings must be set so that browsers connect via the VIP. In a Microsoft based LAN environment, this is typically achieved using AD group policy.

Note

Depending on your requirements, it may be necessary to use an FQDN rather than an IP address for the Proxy server address. If you use an FQDN, make sure you have a valid DNS configuration that correctly resolves the hostname.

Browser Network Settings:

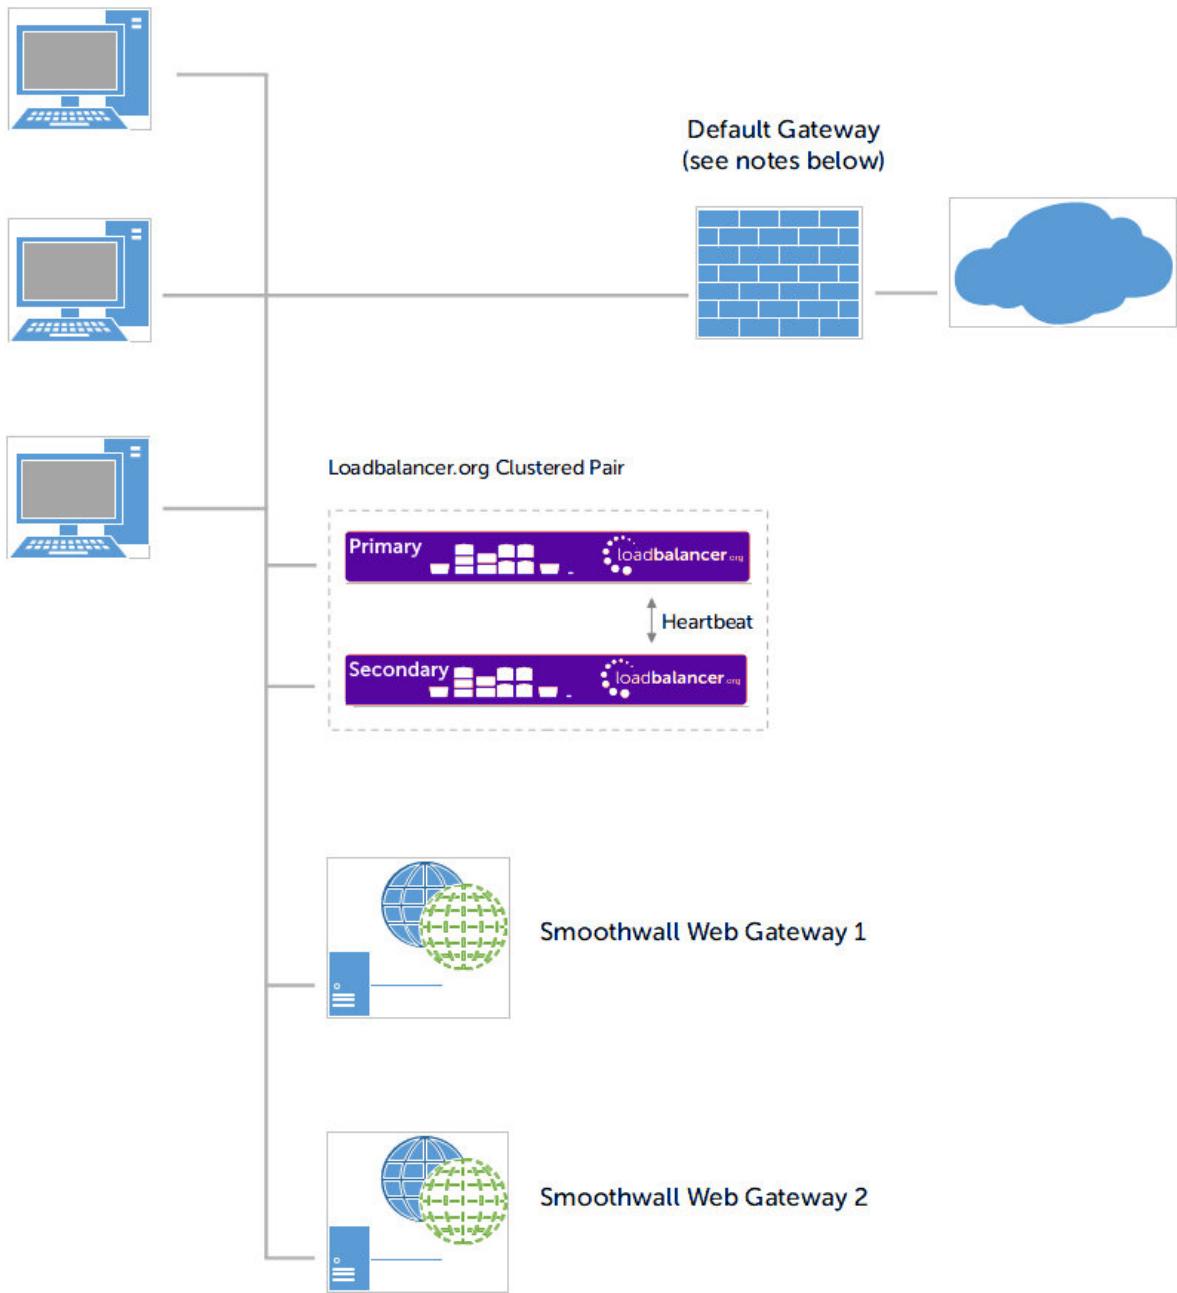




11. Option 2 - Transparent Routed Proxy Mode

11.1. Deployment Architecture





Notes

- Rules must be added to the router/firewall so that the required traffic (typically HTTP & HTTPS on port 80 & 443) is sent transparently to the load balancer, please see [Router/Default Gateway Configuration](#) for example rules for a Linux router
- As with Explicit Proxy Mode, the load balancer is configured in Layer 4 DR mode
- Firewall rules must be added to the load balancer to transparently send traffic to the Web Gateways (see [Configure Firewall Rules](#))
- 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](#)

11.2. Load Balancer Configuration



11.2.1. Create the Virtual Service (VIP)

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

Label	Proxy	?	
Virtual Service	IP Address	1	?
	Ports		?
Protocol	Firewall Marks	?	
Forwarding Method	Direct Routing	?	
Cancel Update			

4. Enter an appropriate label (name) for the VIP, e.g. **Proxy**.
5. Change the *Virtual Service IP address* field to **1**.

Note

This is the reference number for the 'Firewall Mark'. The same reference number is used when configuring the firewall rules – please see [Configure Firewall Rules](#) for more details.

6. Clear the *Virtual Service Ports* field, the ports are defined in the firewall rules in [Configure Firewall Rules](#).
7. Ensure that *Protocol* is set to **Firewall Marks**.
8. Ensure that *Forwarding Method* is set to **Direct Routing**.
9. Click **Update**.
10. Now click **Modify** next to the newly created VIP.
11. Ensure **Persistence** is enabled and set the **Persistence Timeout** to **3600** (i.e. 1 hour).

Note

It's **optionally** possible to define the parent node of the Smoothwall cluster as a fallback server for the new virtual service.

Smoothwall parent nodes do not ordinarily process client traffic and are designed to handle logging and cluster configuration only. It may be desirable, however, for the load balancer to direct client traffic to the parent node in the event that **all** real servers, i.e. all Smoothwall child nodes, fail health checking and are marked as offline.

To implement this, enter the IP address of the parent node in the **IP Address** field under the **Fallback Server** section, and leave the **Port** field under the **Fallback Server** section empty.

12. Under the *Health Checks* section change *Check Type* to **Ping Server**.

13. Click **Update**.

11.2.2. Add the Floating IP

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IPs*.



2. Enter an appropriate IP address for the Virtual Service, e.g. **192.168.2.202**.

3. Click **Add Floating IP**.

11.2.3. Configure Firewall Rules

The *Firewall Script* page is *locked* by default on newer Loadbalancer.org appliances as part of "Secure Mode", which makes applying the changes described below impossible.

Note

To enable editing of the firewall script, navigate to *Local Configuration > Security*, set *Appliance Security Mode* to **Custom**, and click the **Update** button to apply the change. Editing the *Firewall Script* page will then be possible.

1. Using the WebUI, navigate to: *Maintenance > Firewall Script*.

2. Scroll down to the Firewall Marks section.

3. Add the following lines to this section as shown in the screen shot below:

```
iptables -t mangle -A PREROUTING -p tcp --dport 80 -j MARK --set-mark 1
iptables -t mangle -A PREROUTING -p tcp --dport 443 -j MARK --set-mark 1
ip rule add prio 100 fwmark 1 table 100
ip route add local 0/0 dev lo table 100
```

Note

Please see [Modified Transparent Mode Firewall Rules](#) if you intend to forward ALL traffic to the Web Gateways.



Firewall Script

```
1  #!/bin/sh
2  # $Id$
3
4  #
5  # User firewall script for Loadbalancer.org appliance.
6  #
7  #
8  #
9  #
10 #
11 # Please note:
12 #           Most configurations will not require any changes to be made to
13 #           this script.
14 #
15 #
16 #           Administrators will only need to modify this script if their
17 #           needs are not met by the lock-down wizard, auto-NAT, and
18 #           automatic firewall mark functions of the web interface.
19 #
20 #
21 #
22 #
23 ##### One-arm NAT Mode #####
24 # For one-arm NAT, ICMP re-directs will need to be disabled.
25 # (1 = on, 0 = off)
26 #echo "0" >/proc/sys/net/ipv4/conf/all/send_redirects
27 #echo "0" >/proc/sys/net/ipv4/conf/default/send_redirects
28 #
29 #
30 #
31 #
32 #####

```

Update

4. Click **Update**.

11.2.4. Define the Real Servers (RIPs)

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

Label	<input type="text" value="Proxy1"/>	?
Real Server IP Address	<input type="text" value="192.168.2.210"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?
		Cancel Update

4. Enter an appropriate label (name) for the first Web Gateway, e.g. **Proxy1**.
5. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.2.210**.
6. Click **Update**.
7. Repeat the above steps to add your other Web Gateway(s).



Note

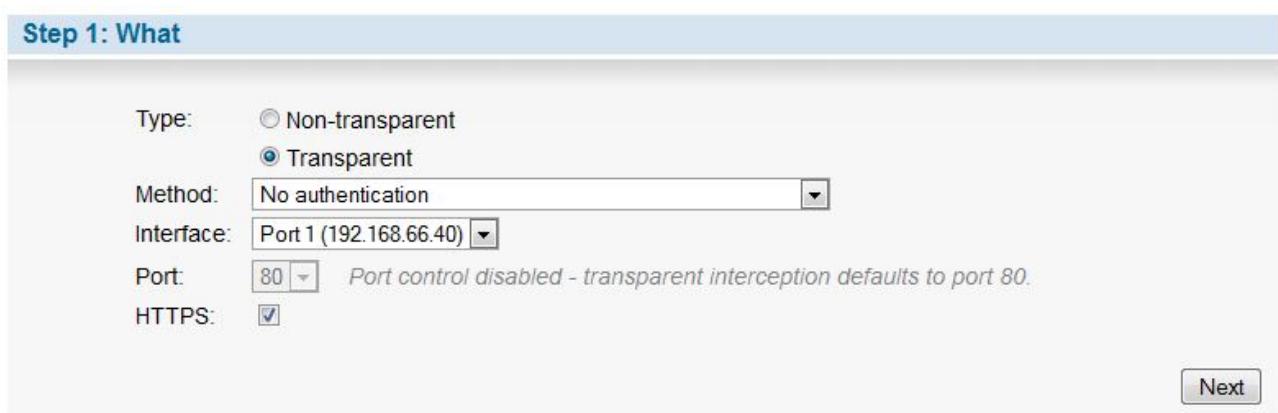
If you'd like to configure email alerts for the Virtual Service, please refer to [Configuring Email Alerts for Virtual Services](#) in the appendix.

11.3. Web Gateway Configuration

11.3.1. Web Gateway Operating Mode

The Smoothwall Web Gateway can easily be configured for transparent mode using the policy wizard.

Use the WebUI option: **Web Proxy > Authentication > Policy Wizard** and select "Transparent" as shown below:



Now click **Next** to run through the wizard and configure the remaining settings and apply the policy.

Note

When using Transparent Routed Mode, it's not necessary to modify the Web Gateway to accept traffic destined for the VIP, this is only required when using Explicit Proxy Mode. However, it's still recommended to configure this so both modes are catered for from the start.

11.4. Router/Default Gateway Configuration

Depending on your network configuration, rules must be added to the router/default gateway so that all required traffic (typically HTTP & HTTPS on port 80 & 443) is sent to the floating IP address on the load balancer. The load balancer then distributes this traffic between the Web Gateways. The example shown below is for a Linux based router:

Example iptables rules for a Linux based router:

```
SUBNET="192.168.2.0/24"
FWMARK="5"
TABLE="10"
LOADBALANCER ="192.168.2.202"
iptables -t mangle -A PREROUTING -s $SUBNET -p tcp --dport 80 -j MARK --set-mark $FWMARK
iptables -t mangle -A PREROUTING -s $SUBNET -p tcp --dport 443 -j MARK --set-mark $FWMARK
ip route add default via $LOADBALANCER dev eth3 table $TABLE
ip rule add fwmark $FWMARK table $TABLE
```

This example uses policy routing via firewall marks. This works by first selecting and marking the packets we want to be sent to the Web Gateway, i.e. all packets on port 80 & 443. Then, when the kernel goes to make a



routing decision, the marked packets aren't routed using the normal routing table, instead via table 10 in this case. Table 10 has only one entry: route packets to the Web Gateway.

Note

This is required when no changes have been made to the clients default gateway settings.

11.5. Client Configuration

If rules are configured on the router as described in the section above, no client change are required. If such rules are not configured, then the default gateway on the client PCs must be modified to be the load balancer.

12. Testing & Verification

Note

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

To verify that the traffic is passing through the load balancer correctly the following reporting options can be used:

1. System Overview
2. Reports > Layer 4 Status
3. Reports > Layer 4 Current Connections

Several reporting and dashboard options are also available on the Web Gateway, for more details please refer to the Smoothwall Web Gateway documentation.

12.1. Layer 4 – Current Connections

12.1.1. Explicit Proxy Mode

The example screen shot below illustrates that the test client (192.168.64.7) sends requests to the VIP (192.168.111.88), the load balancer then forwards the request onto the Web Gateway (192.168.64.60).



Layer 4 Current Connections

```
Check Status

IPVS connection entries
pro expire state      source          virtual          destination
TCP 18:07 ESTABLISHED 192.168.64.7:3565 192.168.111.88:8080 192.168.64.60:8080
TCP 18:05 ESTABLISHED 192.168.64.7:3563 192.168.111.88:8080 192.168.64.60:8080
TCP 18:03 ESTABLISHED 192.168.64.7:3541 192.168.111.88:8080 192.168.64.60:8080
TCP 18:05 ESTABLISHED 192.168.64.7:3570 192.168.111.88:8080 192.168.64.60:8080
TCP 18:07 ESTABLISHED 192.168.64.7:3567 192.168.111.88:8080 192.168.64.60:8080
TCP 18:03 ESTABLISHED 192.168.64.7:3572 192.168.111.88:8080 192.168.64.60:8080
TCP 18:05 ESTABLISHED 192.168.64.7:3560 192.168.111.88:8080 192.168.64.60:8080
TCP 18:03 ESTABLISHED 192.168.64.7:3561 192.168.111.88:8080 192.168.64.60:8080
TCP 18:03 ESTABLISHED 192.168.64.7:3571 192.168.111.88:8080 192.168.64.60:8080
TCP 18:07 ESTABLISHED 192.168.64.7:3566 192.168.111.88:8080 192.168.64.60:8080
TCP 02:58 NONE        192.168.64.7:0   192.168.111.88:8080 192.168.64.60:8080
TCP 18:03 ESTABLISHED 192.168.64.7:3564 192.168.111.88:8080 192.168.64.60:8080
TCP 18:03 ESTABLISHED 192.168.64.7:3568 192.168.111.88:8080 192.168.64.60:8080
```

12.1.2. Transparent Mode

The example screen shot below illustrates the difference when running in transparent mode.

Layer 4 Current Connections

```
Check Status

IPVS connection entries
pro expire state      source          virtual          destination
TCP 00:41 FIN_WAIT    192.168.64.7:5774 70.42.56.98:80      192.168.64.60:80
TCP 00:15 FIN_WAIT    192.168.64.7:5758 74.208.104.65:80    192.168.64.60:80
TCP 14:19 ESTABLISHED 192.168.64.7:5681 93.188.129.145:80  192.168.64.60:80
TCP 00:50 FIN_WAIT    192.168.64.7:5779 70.42.56.98:80      192.168.64.60:80
TCP 00:47 FIN_WAIT    192.168.64.7:5778 70.42.56.98:80      192.168.64.60:80
TCP 14:35 ESTABLISHED 192.168.64.7:5679 176.34.178.134:80  192.168.64.60:80
TCP 14:35 ESTABLISHED 192.168.64.7:5691 178.236.5.70:80    192.168.64.60:80
```

Many reporting and dashboard options are also available in the Smoothwall Web Gateway user interface. For more details please refer to the appropriate Smoothwall documentation.

13. Technical Support

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

14. Further Documentation

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



15. Appendix

15.1. 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.

15.1.1. 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



① Important

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

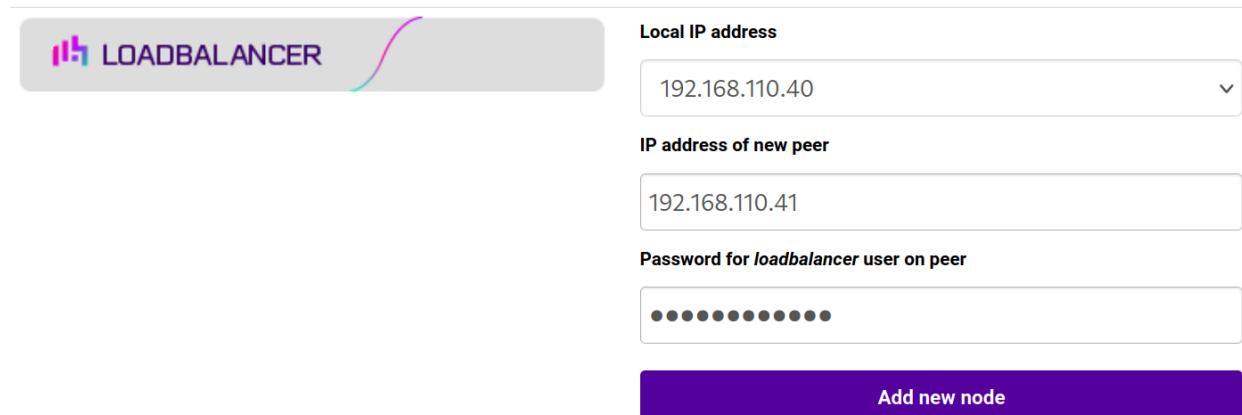
15.1.2. Configuring the HA Clustered Pair

ℹ 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



Local IP address
192.168.110.40

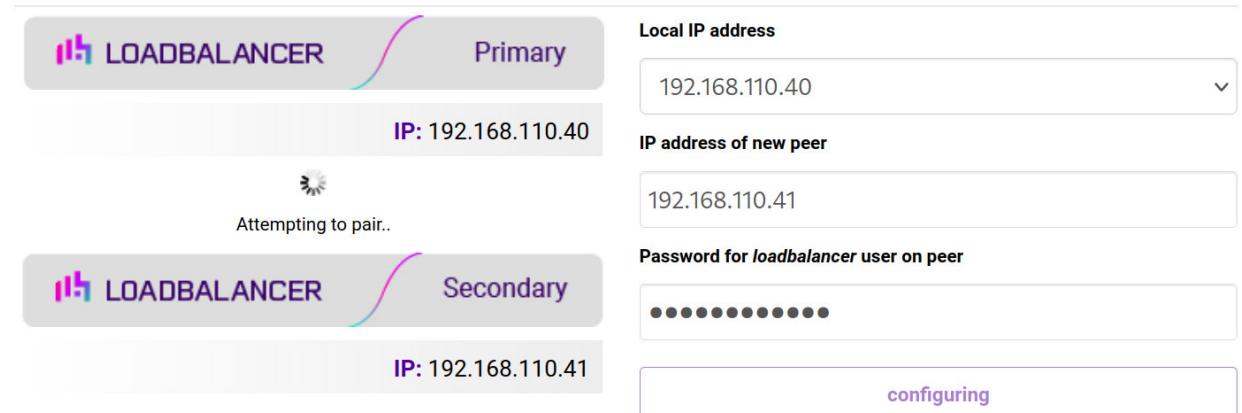
IP address of new peer
192.168.110.41

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



Primary

IP: 192.168.110.40

Attempting to pair..

Secondary

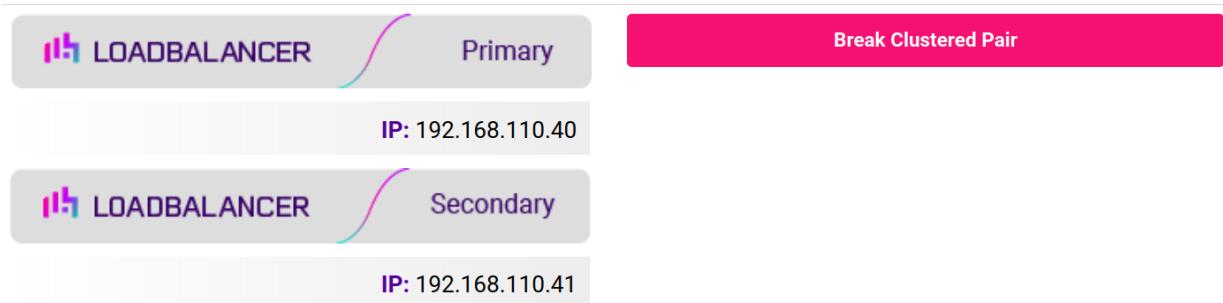
IP: 192.168.110.41

configuring

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



High Availability Configuration - primary



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

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](#).

Note If you'd like to configure email alerts for heartbeat, please refer to [Configuring Email Alerts for Heartbeat](#) in the appendix.

15.2. Modified Transparent Mode Firewall Rules

If ALL traffic is to be forwarded to the Web Gateways, the firewall rules below should be used rather than the rules in [Configure Firewall Rules](#):

Replace:

```
iptables -t mangle -A PREROUTING -p tcp --dport 80 -j MARK --set-mark 1
iptables -t mangle -A PREROUTING -p tcp --dport 443 -j MARK --set-mark 1
ip rule add prio 100 fwmark 1 table 100
ip route add local 0/0 dev lo table 100
```

With:

```
iptables -t mangle -A PREROUTING -p tcp -j MARK --set-mark 1
iptables -t mangle -A PREROUTING -p udp -j MARK --set-mark 1
iptables -t mangle -A PREROUTING -p tcp -d <LB-IP> -j MARK --set-mark 2
iptables -t mangle -A PREROUTING -p udp -d <LB-IP> -j MARK --set-mark 2
ip rule add prio 100 fwmark 1 table 100
ip route add local 0/0 dev lo table 100
```



Notes

- <LB-IP> should be replaced with the base IP address of the load balancer (typically **eth0**), this is the address used by heartbeat and for administration purpose
- If these modified firewall rules are used, then either the default gateway for client PC's should be changed to be the load balancer, or the rules on the router should be changed to forward all traffic to the load balancer
- This will only work for TCP and UDP traffic. So for example, ICMP and some VPN technologies will not work because the load balancer only supports TCP and UDP.

Don't hesitate to contact our support team if you need further assistance: support@loadbalancer.org.

15.3. Configuring Email Alerts for Virtual Services

Email alerts can be configured for layer 4 and layer 7 Virtual Services. This enables emails to be sent when one or more of the associated Real Servers fail their health check and also when they subsequently start to pass their health check.

15.3.1. Layer 4

For layer 4 Virtual Services, settings can be configured globally for all VIPs or individually per VIP.

Global Layer 4 Email Settings

Once configured, these settings apply to all layer 4 VIPs by default.

To configure global email alert settings for layer 4 services:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 Advanced Configuration*.

Email Alert Source Address	lb1@loadbalancer.org	?
Email Alert Destination Address	alerts@loadbalancer.org	?
Auto-NAT	off ▾	?
Multi-threaded	yes ▾	?
		Update

2. Enter an appropriate email address in the *Email Alert Source Address* field.

e.g. `lb1@loadbalancer.org`

3. Enter an appropriate email address in the *Email Alert Destination Address* field.

e.g. `alerts@loadbalancer.org`



4. Click **Update**.

VIP Level Settings

Note VIP level settings override the global settings.

Once configured, these settings apply to the individual VIP.

To configure VIP level email alerts:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 Virtual Service* and click **Modify** next to the VIP to be configured.
2. Scroll down to the *Fallback Server* section.

Email Alert Destination Address	alerts@loadbalancer.org	?
		Cancel Update

3. Enter an appropriate email address in the *Email Alert Destination Address* field.

e.g. alerts@loadbalancer.org

4. Click **Update**.

Note You can set the *Email Alert Source Address* field as explained above if required to configure a default source address.

15.4. Configuring Email Alerts for Heartbeat

Email alerts can be setup for heartbeat once a clustered pair has been configured. This enables alerts to be sent when the primary/secondary communication state has changed. This can occur when the secondary appliance takes over from the primary, when the primary takes over from the secondary and also when there is a communication issue between the 2 appliances.

To configure email alert settings for Heartbeat:

1. Using the WebUI, navigate to: *Cluster Configuration > Heartbeat Configuration*.
2. Scroll down to the **Email Alerts** section.

Email Alerts

Email Alert Destination Address	alerts@loadbalancer.org	?
Email Alert Source Address	lb1@loadbalancer.org	?



3. Enter an appropriate email address in the *Email Alert Destination Address* field.
4. Enter an appropriate email address in the *Email Alert Source Address* field.
5. Click **Modify Heartbeat Configuration**.

15.5. Configuring a Smart Host (SMTP relay)

For Heartbeat (and layer 4 services), email alerts are sent from the load balancer directly to the mail server defined in the destination domain's DNS MX record by default. Alternatively, a custom smart host (mail relay server) can be specified. A smart host is an email server through which approved devices can send emails. Where possible, we recommend that you use a smart host for email alerts as this often helps improve the deliverability of emails.

To configure a Smart Host:

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration*.
2. Scroll down to the **SMTP Relay** section.
3. Specify the FQDN or IP address of the *Smart Host*.
4. Click **Update**.

 **Note**

By default the *Smart Host* is set as the destination email domain's DNS MX record when the *Email Alert Destination Address* is configured. It must either be left at its default setting or a custom smart host must be configured to enable email alerts to be sent.



16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.4.0	6 August 2019	<p>Styling and layout</p> <p>Added notes explaining how * can be used in virtual service port fields</p> <p>Added note explaining how to disable "Secure Mode" to unlock the firewall script page</p>	<p>General styling updates</p> <p>Required updates</p>	AH
1.4.1	29 August 2019	Added note sections on how to optionally define a parent Smoothwall node as a fallback server	Additional option documented based on discussions with Smoothwall	AH
1.4.2	5 June 2020	<p>New title page</p> <p>Updated Canadian contact details</p>	<p>Branding update</p> <p>Change to Canadian contact details</p>	AH
1.5.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH,RJC,ZAC
1.5.1	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
1.5.2	2 February 2023	Updated screenshots	Branding update	AH
1.5.3	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.6.0	24 March 2023	<p>New document theme</p> <p>Modified diagram colours</p>	Branding update	AH



Version	Date	Change	Reason for Change	Changed By
1.6.1	12 August 2025	<p>Corrected terminology on diagrams</p> <p>Added section on configuring email alerts</p> <p>Corrected spelling errors</p>	<p>Technical accuracy</p> <p>Optional configuration</p>	RJC





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](https://www.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.

