



Load Balancing Clearswift Secure Web Gateway

Deployment Guide
v1.3.5

Table of Contents

1. About this Guide.....	4
2. Loadbalancer.org Appliances Supported.....	4
3. Loadbalancer.org Software Versions Supported.....	4
4. Clearswift Web Gateway Appliances Supported.....	4
5. Benefits of Implementing a Load Balancer.....	5
6. Load Balancer Configuration Options.....	5
Deployment Modes.....	5
Layer 4 (Recommended).....	5
Layer 7.....	5
Persistence / Server Affinity.....	6
Source IP Address (Recommended).....	6
Destination Hash.....	6
7. Web Gateway Deployment Modes.....	6
1 – Explicit Proxy Mode (Recommended).....	6
2 – Transparent Routed Proxy Mode.....	6
8. Summary of Deployment Options.....	6
9. Loadbalancer.org Appliance – the Basics.....	7
Virtual Appliance Download & Deployment.....	7
Initial Network Configuration.....	7
Accessing the Web User Interface (WebUI).....	8
HA Clustered Pair Configuration.....	9
10. Option 1 - Explicit Proxy Mode (Recommended).....	10
Option 1A – Using DR (Direct Return) Mode (Recommended).....	10
Deployment Architecture.....	10
Load Balancer Configuration.....	11
Web Gateway Configuration.....	12
Finalize Settings.....	13
Option 1B – Using NAT Mode.....	14
Deployment Architecture.....	14
Load Balancer Configuration.....	15
Web Gateway Configuration.....	17
Finalize Settings.....	17
Option 1C – Using NAT Mode (Preferred NAT Topology).....	18
Deployment Architecture.....	18
Load Balancer Configuration.....	19
Web Gateway Configuration.....	21
Finalize Settings.....	21
Configuration Settings Common to Options 1A, 1B & 1C.....	21
Web Gateway Operating Mode.....	21
Proxy Port Configuration.....	21
Client Configuration.....	22
11. Option 2 - Transparent Routed Proxy Mode.....	23
Deployment Architecture.....	23

Load Balancer Configuration.....	24
Create the Virtual Service (VIP).....	24
Add the Floating IP.....	24
Configure Firewall Rules.....	25
Define the Real Servers (RIPs).....	26
Web Gateway Configuration.....	26
Web Gateway Operating Mode.....	26
Router/Default Gateway Configuration.....	26
Client Configuration.....	27
12. Testing & Validation.....	27
Layer 4 – Current Connections.....	27
13. Technical Support.....	28
14. Further Documentation.....	28
15. Conclusion.....	28
16. Appendix.....	29
1 – Clustered Pair Configuration – Adding a Slave Unit.....	29
2 – Modified Transparent Mode Firewall Rules.....	31
3 - Company Contact Information.....	32

1. About this Guide

This guide details the steps required to configure a load balanced Clearswift Secure Web Gateway environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Clearswift 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 relevant Administration Manual:

- [v7 Administration Manual](#)
- [v8 Administration Manual](#)

2. Loadbalancer.org Appliances Supported

All our products can be used for load balancing Clearswift Web Gateways. The complete list of models is shown below:

Discontinued Models	Current Models *
Enterprise R16	Enterprise R20
Enterprise VA R16	Enterprise MAX
Enterprise VA	Enterprise 10G
Enterprise R320	Enterprise Ultra
	Enterprise VA R20
	Enterprise VA MAX
	Enterprise AWS **
	Enterprise AZURE **

* For full specifications of these models please refer to: <http://www.loadbalancer.org/products/hardware>

** Some features may not be supported, please check with Loadbalancer.org support

3. Loadbalancer.org Software Versions Supported

- v7.6.4 and later

4. Clearswift Web Gateway Appliances Supported

- Clearswift Web Gateway - all versions

5. Benefits of Implementing a Load Balancer

Implementing Loadbalancer.org appliances enables multiple Clearswift Secure 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 Web Gateways.

DEPLOYMENT MODES

LAYER 4 (RECOMMENDED)

DR Mode - Direct Server Return Mode (Recommended)

In this mode, traffic from the client to the Web Gateway passes via the load balancer, return traffic passes directly back to the client which maximizes performance. Direct routing works by changing the destination MAC address of the incoming packet on the fly which is very fast. This mode is transparent by default meaning that the Web Gateway sees the real client IP address and not the IP address of the load balancer.

Due to its speed, overall simplicity and effectiveness, Direct Routing (DR) mode with source IP persistence is our recommended method and can be used in both Explicit Proxy Mode & Transparent Routed Proxy Mode.

NAT Mode - Network Address Translation Mode

This mode requires the implementation of a two-arm infrastructure with an internal and external subnet to carry out the translation (the same way a firewall works). Return traffic MUST pass back via the load balancer. This can be achieved by either setting the default gateway on the Web Gateways to be the load balancer or by configuring a static route on the Web Gateways that forces client return traffic to pass back via the load balancer. This mode offers high performance and like DR mode is transparent by default.

LAYER 7

SNAT Mode - Source Network Address Translation

Using HAProxy in SNAT mode means that the load balancer is acting as a full proxy and therefore it doesn't have the same raw throughput as the layer 4 methods. Also, this method is not transparent by default so the real servers (i.e. the Web Gateways) will see the source address of each request as the load balancers IP address. This is generally not desirable, although this can be resolved in two ways: either by reading the X-Forwarded-For header that's included by default when using HAProxy, or by enabling TProxy on the load balancer. The issue with using TProxy is that the default gateway on the real servers must be changed to be the load balancer and it also requires a two-arm infrastructure with two subnets which complicates the deployment. The same requirements apply when using layer 4 NAT mode as mentioned above. SNAT mode does not have the raw throughput of the layer 4 solutions and is therefore not normally used for Web Gateway load balancing deployments.

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.

SOURCE IP ADDRESS (RECOMMENDED)

Source IP persistence is the default option for Layer 4 services and can easily be selected for Layer 7 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.

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 proxy 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 Web Gateway servers.

7. Web Gateway Deployment Modes

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

1 – EXPLICIT PROXY MODE (RECOMMENDED)

This mode requires the load balancers VIP address to be defined in users browsers. This means that the load balancer will receive client requests and then distribute these requests across the back-end Web Gateways. Clearswift refer to this as “Standard Proxy Mode”. Please refer to the section starting on page [10](#) for configuration details.

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. Clearswift refer to this as “Transparent Proxy Mode”. Please refer to the section starting on page [23](#) for configuration details.

8. Summary of Deployment Options

Option	Web Gateway Mode	Load Balancer Mode	Notes
Option 1A (Recommended)	Standard Proxy Mode	DR Mode	The Web Gateways must be configured to accept traffic for the VIP. Please refer to page 10 for

			configuration details.
Option 1B	Standard Proxy Mode	NAT Mode	The load balancer must be set as the default gateway for the Web Gateways. Please refer to page 14 for configuration details.
Option 1C	Standard Proxy Mode	NAT Mode	A static route must be configured on the Web Gateways to send client return traffic back via the load balancer. Please refer to page 18 for configuration details.
Option 2	Transparent Proxy Mode	DR Mode	Firewall rules must be added to the load balancer to transparently send traffic to the Web Gateways. Please refer to page 23 for configuration details.

9. Loadbalancer.org Appliance – the Basics

VIRTUAL APPLIANCE DOWNLOAD & DEPLOYMENT

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 and XEN and has been optimized for each Hypervisor. By default, the VA is allocated 1 CPU, 2GB of RAM and has an 8GB 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 the [Administration Manual](#) and the ReadMe.txt text file included in the VA download for more detailed information on deploying the VA using various Hypervisors.

INITIAL NETWORK CONFIGURATION

The IP address, subnet mask, default gateway and DNS settings can be configured in several ways as detailed below:

Method 1 - Using the Network Setup Wizard at the console

After boot up, follow the instructions on the console to configure the IP address, subnet mask, default gateway and DNS settings.

Method 2 - Using the WebUI

Using a browser, connect to the WebUI on the default IP address/port: **http://192.168.2.21:9080**

To set the IP address & subnet mask, use: *Local Configuration > Network Interface Configuration*

To set the default gateway, use: *Local Configuration > Routing*

To configure DNS settings, use: *Local Configuration > Hostname & DNS*

Method 3 - Using Linux commands

At the console, set the initial IP address using the following command:

```
ip addr add <IP address>/<mask> dev eth0
```

At the console, set the initial default gateway using the following command:

```
route add default gw <IP address> <interface>
```

At the console, set the DNS server using the following command:

```
echo nameserver <IP address> >> /etc/resolv.conf
```

Note:

If method 3 is used, you must also configure these settings using the WebUI, otherwise the settings will be lost after a reboot.

ACCESSING THE WEB USER INTERFACE (WEBUI)

The WebUI can be accessed via HTTP at the following URL: **http://192.168.2.21:9080/lbadmin**

* *Note the port number → 9080*

The WebUI can be accessed via HTTPS at the following URL: **https://192.168.2.21:9443/lbadmin**

* *Note the port number → 9443*

(replace 192.168.2.21 with the IP address of your load balancer if it's been changed from the default)

Login using the following credentials:

Username: loadbalancer

Password: loadbalancer

Note:

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

Once logged in, the WebUI will be displayed as shown on the following page:

The screenshot displays the 'Enterprise VA MAX' interface of Loadbalancer.org. The top navigation bar shows 'Master | Slave', 'Active | Passive', and 'Link' status, along with a refresh button and '5 Seconds' interval. The left sidebar contains a menu with items like 'System Overview', 'Local Configuration', 'Cluster Configuration', 'Maintenance', 'View Configuration', 'Reports', 'Logs', and 'Support'. The main content area is titled 'SYSTEM OVERVIEW' and includes a modal dialog asking 'Would you like to run the Setup Wizard?' with 'Accept' and 'Dismiss' buttons. Below the dialog is a filter bar for 'VIRTUAL SERVICE', 'IP', 'PORTS', 'CONNS', 'PROTOCOL', 'METHOD', and 'MODE', with a message stating 'No Virtual Services configured.' Three charts are displayed: 'Network Bandwidth' showing RX and TX traffic, 'System Load Average' showing 1m, 5m, and 15m averages, and 'Memory Usage' showing Used, Page, Buffer, and Free memory.

(shows v8.2.x)


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 slave unit is covered in section 1 of the Appendix on page [29](#).

10. Option 1 - Explicit Proxy Mode (Recommended)

OPTION 1A – USING DR (DIRECT RETURN) MODE (RECOMMENDED)

DEPLOYMENT ARCHITECTURE



Notes:

- Browser settings on client PC's must be changed to point at the Virtual Service (VIP) on the load balancer (see page [22](#))
- The load balancer is configured in one-arm Layer 4 DR mode
- The Clearswift Web Gateways must be configured to accept traffic for the VIP (see page [12](#))
- 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 slave unit is covered in section 1 of the Appendix on page [29](#)
- For more information on Clearswift Web Gateway deployment options please refer to [this URL](#)

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	<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="Direct Routing"/>	?	
		<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**
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 *Persistence Timeout* to **3600** (i.e. 1 hour)
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	<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"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="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)

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.

To configure a Linux based Web Gateway to accept traffic for the VIP, the iptables command below must be added to an appropriate startup script such as `/etc/rc.local` so that it is automatically executed each time the Web Gateway boots. It can also be executed immediately by running the command at the command prompt, but the setting will be lost after a reboot unless the command has been added to a startup script.

```
iptables -t nat -A PREROUTING -p tcp -d <VIP address> -j REDIRECT
```

e.g.

```
iptables -t nat -A PREROUTING -p tcp -d 192.168.2.202 -j REDIRECT
```


i.e. Redirect any incoming packets destined for the VIP to the local address

Note:

For more information please refer to the [Administration Manual](#) and search for 'ARP Problem'.

Configuring the Clearswift Web Gateway Appliance

At the console or via an SSH session, login as user 'console', the following menu will be displayed:



1. Select the **Login** option
2. Enter the password for the 'system' user, a command prompt will then be displayed
3. Now run through the following commands to configure the Web Gateway:

```
$ sudo su
```

Note:

When prompted, enter the password for the 'system' user.

```
# mkdir /opt/nfs  
# echo '#!/bin/sh' > /opt/nfs/iptablescript.nfs  
# echo '/sbin/iptables -t nat -A PREROUTING -p tcp -d 192.168.2.202 -j REDIRECT' >>  
/opt/nfs/iptablescript.nfs
```

Note:

The IP address specified in the command above should match your VIP address.

```
# touch /opt/nfs/nfsinfo  
# touch /opt/nfs/appmountchk  
# chmod +x /opt/nfs/appmountchk  
# chmod +x /opt/nfs/iptablescript.nfs  
# /etc/init.d/rc.firewall stop  
# /etc/init.d/rc.firewall start
```

4. Reboot the Web Gateway to apply the new setting

Note:


These steps must be followed on all Gateways.

FINALIZE SETTINGS

Now refer to the section "*Configuration Settings Common to Options 1A, 1B & 1C*" on page [21](#) to finalize Web Gateway settings and configure client browser settings.

OPTION 1B – USING NAT MODE

DEPLOYMENT ARCHITECTURE



Notes:

- Browser settings on client PC's must be changed to point at the Virtual Service (VIP) on the load balancer (see page [22](#))
- 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 master and slave in the event of a failover (see page [15](#))
- 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 slave unit is covered in section 1 of the Appendix on page [29](#)
- For more information on Clearswift Web Gateway deployment options please refer to [this URL](#)

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:

The screenshot displays the 'IP Address Assignment' configuration page. At the top, there are four network interface icons: eth0 (10 GB/s), eth1 (10 GB/s), eth2 (disabled), and eth3 (disabled). Below the icons, the configuration for eth0 is shown with an IP address of 192.168.4.200/24 and an MTU of 1500 bytes. The configuration for eth1 is shown with an IP address of 192.168.2.200/24 and an MTU of 1500 bytes.

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 master and slave units in the event of a failover or fallback. This ensures that the Web Gateways always have a consistent return path via the load balancer – whether the master or slave is active.

To configure a Floating IP:

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

The screenshot displays the 'Floating IP's' configuration page. It shows a 'New Floating IP' field with the value 192.168.4.205 and an 'Add Floating IP' button.

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**
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 *Persistence Timeout* to **3600** (i.e. 1 hour)
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	<input type="text" value="Proxy1"/>	?
Real Server IP Address	<input type="text" value="192.168.4.210"/>	?
Real Server Port	<input type="text" value="8080"/>	?
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"/>


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	<input type="text"/>	?
Auto-NAT	eth1 (Default) ▼	?
Multi-threaded	yes ▼	?
<input type="button" value="Update"/>		

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

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 master or slave is active. Page [15](#) details how to create the Floating IP.

Note:

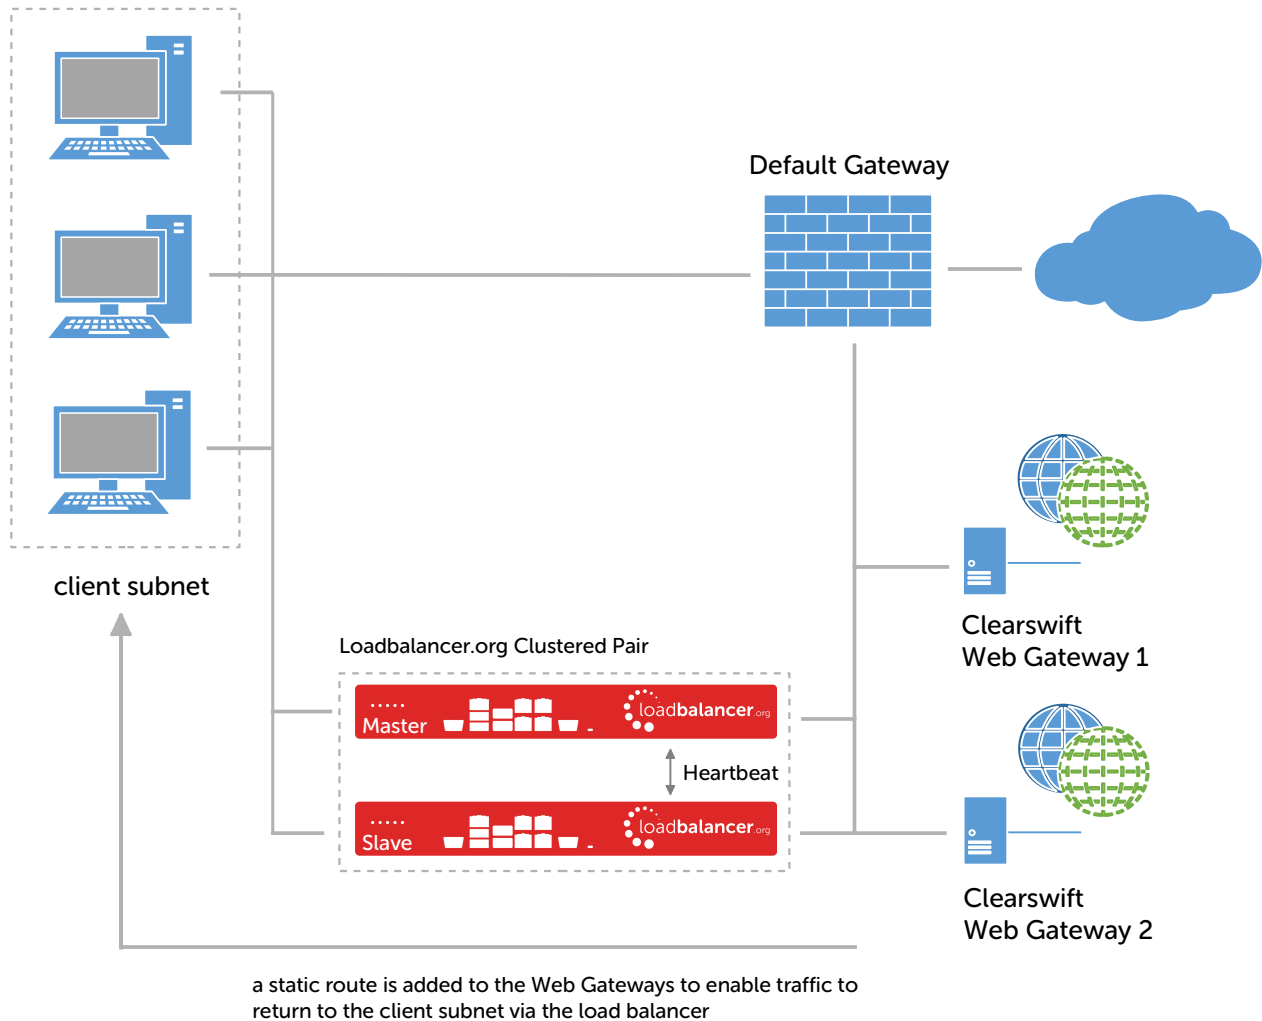
Please refer to the Clearswift Secure Web Gateway documentation for instructions on setting the default gateway. This should be done on all Web Gateways.

FINALIZE SETTINGS

Now refer to the section "*Configuration Settings Common to Options 1A, 1B & 1C*" on page [21](#) to finalize Web Gateway settings and configure client browser settings.

OPTION 1C – USING NAT MODE (PREFERRED NAT TOPOLOGY)

DEPLOYMENT ARCHITECTURE



Notes:

- Browser settings on client PC's must be changed to point at the Virtual Service (VIP) on the load balancer (see page [22](#))
- The load balancer is 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 master and slave in the event of a failover (see page [19](#))
- 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 slave unit is covered in section 1 of the Appendix on page [24](#)
- For more information on Clearswift Web Gateway deployment options please refer to [this URL](#)

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:

The screenshot shows the 'IP Address Assignment' configuration page. At the top, there are four network interface icons: eth0 (10 GB/s), eth1 (10 GB/s), eth2 (disabled), and eth3 (disabled). Below the icons, the configuration for eth0 and eth1 is shown. For eth0, the IP address is 192.168.4.200/24 and the MTU is 1500 bytes. For eth1, the IP address is 192.168.2.200/24 and the MTU is 1500 bytes.

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 master and slave 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 master or slave is active.

To configure a Floating IP:

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

The screenshot shows the 'Floating IP's' configuration page. There is a text input field labeled 'New Floating IP' with the value 192.168.4.205. A green button labeled 'Add Floating IP' is located at the bottom right of the form.

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**
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 *Persistence Timeout* to **3600** (i.e. 1 hour)
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	<input type="text" value="Proxy1"/>	?
Real Server IP Address	<input type="text" value="192.168.4.210"/>	?
Real Server Port	<input type="text" value="8080"/>	?
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"/>

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)

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 master or slave is active. Page [19](#) details how to create the Floating IP.

Note:

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

FINALIZE SETTINGS


Now refer to the section "*Configuration Settings Common to Options 1A, 1B & 1C*" below to finalize Web Gateway and client browser settings.

CONFIGURATION SETTINGS COMMON TO OPTIONS 1A, 1B & 1C

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

WEB GATEWAY OPERATING MODE

The Clearswift Web Gateway can easily be configured for Explicit Proxy Mode using the WebUI option: *System > Proxy Settings > Proxy Mode & Listening Port* and selecting "Standard" as shown below



PROXY PORT CONFIGURATION

The required proxy port can be set under the same WebUI option. As shown below, simply edit the setting that was configured during the initial setup wizard and change the port as required:



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:

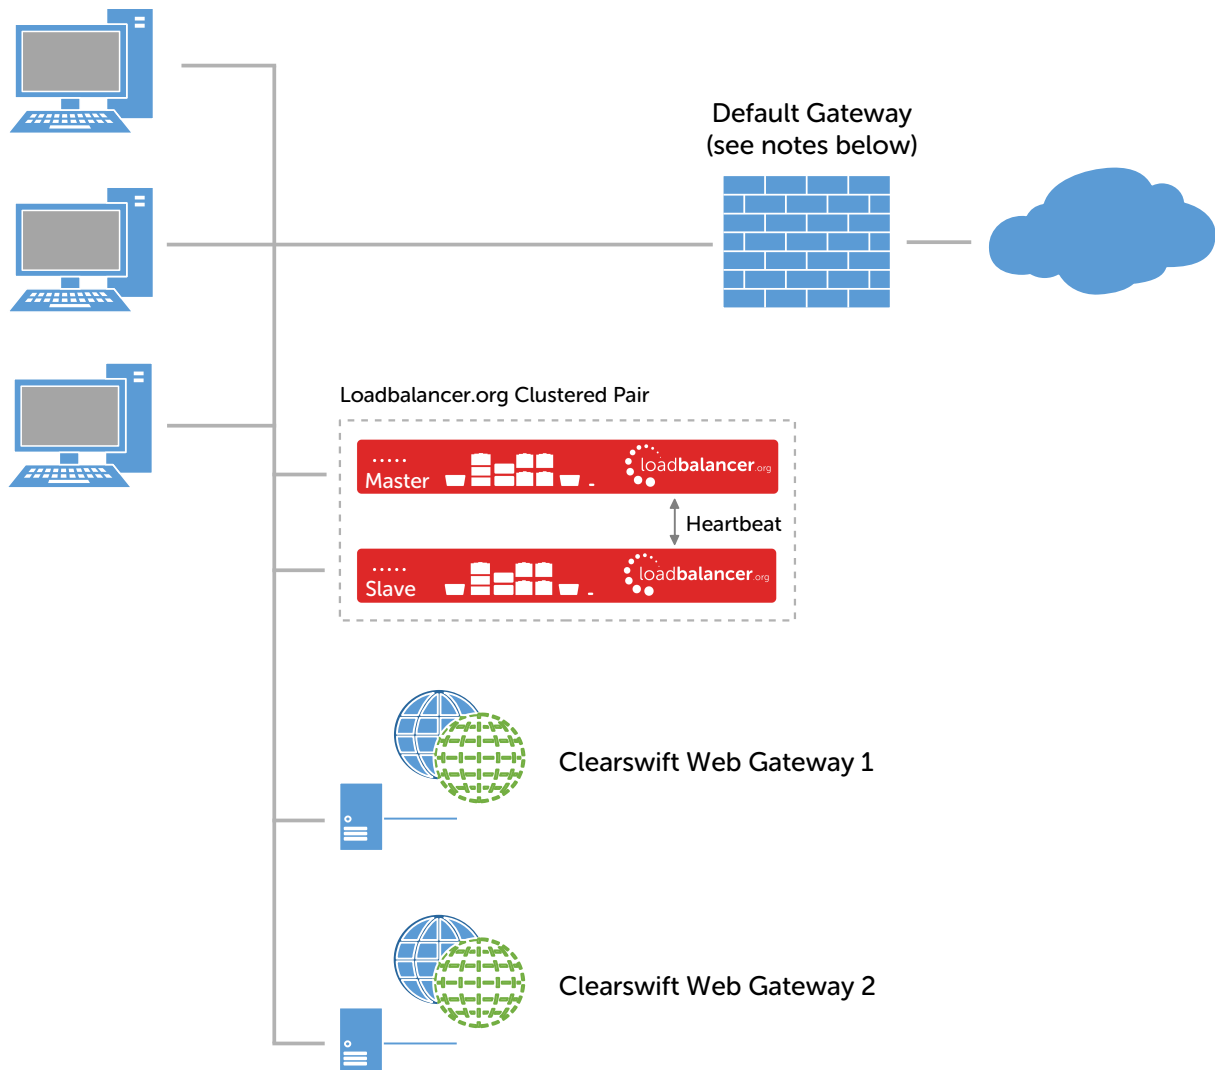
The screenshot shows the 'Local Area Network (LAN) Settings' dialog box. It has two main sections: 'Automatic configuration' and 'Proxy server'. In the 'Automatic configuration' section, the 'Automatically detect settings' checkbox is unchecked. In the 'Proxy server' section, the 'Use a proxy server for your LAN (These settings will not apply to dial-up or VPN connections)' checkbox is checked. The 'Address' field contains '192.168.2.202' and the 'Port' field contains '8080'. There is an 'Advanced' button next to the port field. The 'Bypass proxy server for local addresses' checkbox is also checked. At the bottom, there are 'OK' and 'Cancel' buttons.

The screenshot shows the 'Proxy Settings' dialog box. It has two main sections: 'Servers' and 'Exceptions'. In the 'Servers' section, there is a table with columns for 'Type', 'Proxy address to use', and 'Port'. The 'HTTP' row has '192.168.2.202' in the address field and '8080' in the port field. The 'Secure' row has '192.168.2.202' in the address field and '8080' in the port field. The 'FTP' row has '192.168.2.202' in the address field and '8080' in the port field. The 'Socks' row has empty fields. Below the table, the 'Use the same proxy server for all protocols' checkbox is checked. In the 'Exceptions' section, there is a text box with '192.168*' entered. Below the text box, it says 'Use semicolons (;) to separate entries.' At the bottom, there are 'OK' and 'Cancel' buttons.

Type	Proxy address to use	Port
HTTP:	192.168.2.202	8080
Secure:	192.168.2.202	8080
FTP:	192.168.2.202	8080
Socks:		

11. Option 2 - Transparent Routed Proxy Mode

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 page [26](#) 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 page [25](#))
- 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 slave unit is covered in section 1 of the Appendix on page [29](#)
- For more information on Clearswift Web Gateway deployment options please refer to [this URL](#)

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	<input type="text" value="Proxy"/>	?
Virtual Service	IP Address	<input type="text" value="1"/>
	Ports	<input type="text"/>
Protocol	<input type="text" value="Firewall Marks"/>	?
Forwarding Method	<input type="text" value="Direct Routing"/>	?

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 page [25](#) for more details.

6. Clear the *Virtual Service Ports* field, the ports are defined in the firewall rules on page [25](#)
7. Ensure that *Protocol* is set to **Firewall Marks**

Note:

The ports field will be disabled when this is done.

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 *Persistence Timeout* to **3600** (i.e. 1 hour)
12. Under the *Health Checks* section change *Check Type* to **Ping Server**
13. Click **Update**

ADD THE FLOATING IP

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

New Floating IP	<input type="text" value="192.168.2.202"/>	<input type="button" value="Add Floating IP"/>
-----------------	--	--

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

CONFIGURE FIREWALL RULES

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 section 2 in the Appendix if you intend to forward ALL traffic to the Web Gateways.

MAINTENANCE > FIREWALL SCRIPT

```
##### Manual Firewall Marks #####

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

# Example: Associate HTTP and HTTPS with Firewall Mark 1:
#VIP1="10.0.0.66"
#iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 80 -j MARK --set-mark 1
#iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 443 -j MARK --set-mark 1

# A Virtual Service may then be created in the web interface, using 1 as the
# service address.

##### Packet Filtering #####

# You should always use a network perimeter firewall to lock down all
# external access to the load balancer except the required Virtual Services
# and the required services from your admin machine / network (SSH & HTTPS)

# Allow unlimited traffic on the loopback interface:
#iptables -A INPUT -i lo -j ACCEPT
#iptables -A OUTPUT -o lo -j ACCEPT
```

- 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	<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"/>	?

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)

WEB GATEWAY CONFIGURATION

WEB GATEWAY OPERATING MODE

The Clearswift Web Gateway can easily be configured for transparent routed mode using the WebUI option: *System > Proxy Settings > Proxy Mode & Listening Port* and selecting "transparent" as shown below

Proxy Mode

Select the mode that the Web Proxy will run in:

Proxy Mode :

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.

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 -m tcp --dport 80 -j MARK --set-mark $FWMARK
iptables -t mangle -A PREROUTING -s $SUBNET -p tcp -m 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.

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 & Validation

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

System Overview

Reports > Layer 4 Status

Reports > Layer 4 Current Connections

Several reporting and dashboard options are also available on the Web Gateways, for this please refer to your specific vendors documentation.

LAYER 4 – CURRENT CONNECTIONS**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

IPVS connection entries

pro	expire	state	source	virtual	destination
TCP	13:07	ESTABLISHED	192.168.64.7:3565	192.168.111.88:8080	192.168.64.60:8080
TCP	13:07	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	13:03	ESTABLISHED	192.168.64.7:3564	192.168.111.88:8080	192.168.64.60:8080
TCP	13:03	ESTABLISHED	192.168.64.7:3568	192.168.111.88:8080	192.168.64.60:8080

Transparent Mode

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

LAYER 4 CURRENT CONNECTIONS

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 Clearswift Web Gateway user interface. For more details please refer to the appropriate Clearswift documentation.

13. 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

14. Further Documentation

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

15. Conclusion

Loadbalancer.org appliances provide a very cost effective solution for highly available load balanced Clearswift Web Gateway environments.

16. Appendix

1 – CLUSTERED PAIR CONFIGURATION – ADDING A SLAVE UNIT

If you initially configured just the master unit and now need to add a slave - our recommended procedure, please refer to the relevant section below for more details:

Note:

A number of settings are not replicated as part of the master/slave pairing process and therefore must be manually configured on the slave appliance. These are listed below:

- Hostname & DNS settings
- Network settings including IP addresses, bonding configuration and VLANs
- Routing configuration including default gateways and static routes
- Date & time settings
- Physical – Advanced Configuration settings including Internet Proxy IP address & port, Firewall table size, SMTP relay and Syslog server
- SNMP settings
- Graphing settings
- Firewall Script & Firewall Lockdown Script settings
- Software updates


Version 7:

Please refer to **Chapter 8 – Appliance Clustering for HA** in the [v7 Administration Manual](#).

Version 8:

To add a slave node – i.e. create a highly available clustered pair:

- Deploy a second appliance that will be the slave and configure initial network settings
- Using the WebUI, navigate to: *Cluster Configuration > High-Availability Configuration*



CREATE A CLUSTERED PAIR

loadbalancer.org

Local IP address
192.168.1.20

IP address of new peer
192.168.1.21

Password for loadbalancer user on peer
.....

Add new node

- Specify the IP address and the *loadbalancer* users password (the default is 'loadbalancer') for the slave (peer) appliance as shown above
- Click **Add new node**
- The pairing process now commences as shown below:

- Once complete, the following will be displayed:

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

Note:

Clicking the **Restart Heartbeat** button on the master appliance will also automatically restart heartbeat on the slave appliance.

Note:

Please refer to chapter 9 – Appliance Clustering for HA in the [Administration Manual](#) for more detailed information on configuring HA with 2 appliances.

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 on page [25](#), i.e.:

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

3 - COMPANY CONTACT INFORMATION

<i>Website</i>	URL: www.loadbalancer.org
<i>North America (US)</i>	<p>Loadbalancer.org, Inc. 4250 Lancaster Pike, Suite 120 Wilmington DE 19805 USA</p> <p>Tel: +1 888.867.9504 Fax: +1 302.213.0122 Email (sales): sales@loadbalancer.org Email (support): support@loadbalancer.org</p>
<i>North America (Canada)</i>	<p>Loadbalancer.org Ltd 300-422 Richards Street Vancouver, BC V6B 2Z4 Canada</p> <p>Tel: +1 866.998.0508 Fax: +1 302.213.0122 Email (sales): sales@loadbalancer.org Email (support): support@loadbalancer.org</p>
<i>Europe (UK)</i>	<p>Loadbalancer.org Ltd. Compass House North Harbour Business Park Portsmouth, PO6 4PS UK</p> <p>Tel: +44 (0)330 3801064 Fax: +44 (0)870 4327672 Email (sales): sales@loadbalancer.org Email (support): support@loadbalancer.org</p>
<i>Europe (Germany)</i>	<p>Loadbalancer.org GmbH Tengstraße 27 D-80798 München Germany</p> <p>Tel: +49 (0)89 2000 2179 Fax: +49 (0)30 920 383 6495 Email (sales): vertrieb@loadbalancer.org Email (support): support@loadbalancer.org</p>