

Load Balancing Web Proxies / Filters / Gateways

Version 1.9.0



Table of Contents

1. About this Guide	4
2. Loadbalancer.org Appliances Supported	4
3. Software Versions Supported	4
3.1. Loadbalancer.org Appliance	4
4. Web Proxies/Filters/Gateways	4
5. Benefits of Implementing a Load Balancer	4
6. Load Balancer Configuration Options	5
6.1. Deployment Modes	5
Layer 4 (Recommended)	5
Layer 7	5
6.2. Persistence / Server Affinity	5
Source IP Address (Recommended)	6
Destination Hash	6
7. Web Proxy Deployment Modes	6
7.1. 1 – Explicit Proxy Mode (Recommended)	6
7.2. 2 – Transparent Routed Proxy Mode	6
8. Summary of Deployment Options	6
9. Loadbalancer.org Appliance – the Basics	7
9.1. Virtual Appliance	7
9.2. Initial Network Configuration	7
9.3. Accessing the Appliance WebUI	8
Main Menu Options	9
9.4. Appliance Software Update	10
Determining the Current Software Version	10
Checking for Updates using Online Update	10
Using Offline Update	10
9.5. Ports Used by the Appliance	11
9.6. HA Clustered Pair Configuration	12
10. Option 1 – Explicit Proxy Mode (Recommended)	12
10.1. Option 1A – Using DR (Direct Return) Mode (Recommended)	12
Deployment Architecture	12
Load Balancer Configuration	13
Web Proxy Appliance Configuration	15
Finalize Settings	15
10.2. Option 1B – Using NAT Mode	15
Deployment Architecture	15
Load Balancer Configuration	17
Web Proxy Configuration	19
Finalize Settings	20
10.3. Option 1C – Using NAT Mode (Preferred NAT Topology)	20
Deployment Architecture	20
Load Balancer Configuration	21
Web Proxy Configuration	23
Finalize Settings	23
10.4. Configuration Settings Common to Options 1A, 1B & 1C	23
Web Proxy Operating Mode	24
Client Configuration	24
11. Option 2 - Transparent Routed Proxy Mode	25

11.1. Deployment Architecture	25
11.2. Load Balancer Configuration	26
Create the Virtual Service (VIP)	27
Add the Floating IP	27
Configure Firewall Rules	28
Define the Real Servers (RIPs)	28
11.3. Web Proxy Appliance Configuration	29
Web Proxy Operating Mode	29
Router/Default Gateway Configuration	29
11.4. Client Configuration	29
12. Testing & Verification	29
12.1. Layer 4 – Current Connections	30
Explicit Proxy Mode	30
Transparent Mode	30
13. Technical Support	31
14. Further Documentation	31
15. Appendix	32
15.1. Configuring HA - Adding a Secondary Appliance	32
Non-Replicated Settings	32
Configuring the HA Clustered Pair	33
15.2. 2 – Modified Transparent Mode Firewall Rules	34
16. Document Revision History	36

1. About this Guide

This guide details the steps required to configure a load balanced Web Proxy/Filter/Gateway environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Web Proxy/Filter/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 for load balancing Web Proxies/Filters. For full specifications of available models please refer to <https://www.loadbalancer.org/products>.

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.

4. Web Proxies/Filters/Gateways

Web Proxies/Filters/Gateways provide a number of functions to permit organizations to control the way their staff access the Internet. These products are often appliance based and provide functionality including:

- Web Security & Control
- URL Filtering
- Content Caching
- Anti SPAM/Anti Malware/Anti Virus
- User Authentication
- High Availability

5. Benefits of Implementing a Load Balancer

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



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

6.1. Deployment Modes

Layer 4 (Recommended)

DR Mode - Direct Server Return Mode (Recommended)

In this mode, traffic from the client to the Web Proxy 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 Proxy 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 Proxies to be the load balancer or by configuring a static route on the Web Proxies 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 Proxies) 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 Proxy load balancing deployments.

6.2. Persistence / Server Affinity

Persistence may or may not be required and depends on the specific web proxy 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 Proxy.

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 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 the Web Proxies.

7. Web Proxy Deployment Modes

There are two implementation methods that are typically used – Explicit Proxy Mode & Transparent Routed Proxy Mode. The specific terminology used by each vendor may vary, but will be similar.

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 Proxies. Please refer to [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/proxy 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. Please refer to [Option 2 - Transparent Routed Proxy Mode](#) for configuration details.

Note

Various limitations relating to HTTPS inspection and client authentication may affect your particular Web Proxy appliance when deployed in Transparent Mode. Please check with your particular vendor to determine if this is the case and help choose the most appropriate deployment mode to use.

8. Summary of Deployment Options

Option	Web Proxy Mode	Load Balancer Mode	Notes
Option 1A (Recommended)	Explicit Proxy Mode	DR Mode	The Web Proxies must be configured to accept traffic for the VIP. Please refer to Option 1A for configuration details.



Option	Web Proxy Mode	Load Balancer Mode	Notes
Option 1B	Explicit Proxy Mode	NAT Mode	The load balancer must be set as the default gateway for the Web Proxies. Please refer to Option 1B for configuration details.
Option 1C	Explicit Proxy Mode	NAT Mode	A static route must be configured on the Web Proxies to send client return traffic back via the load balancer. Please refer to Option 1C for configuration details.
Option 2	Transparent Routed Proxy Mode	DR Mode	Firewall rules must be added to the load balancer to transparently send traffic to the Web Proxies. Please refer to Option 2 for configuration details.

9. Loadbalancer.org Appliance – the Basics

9.1. Virtual Appliance

A fully featured, fully supported 30 day trial is available if you are conducting a PoC (Proof of Concept) deployment. The VA is currently available for VMware, Virtual Box, Hyper-V, KVM, XEN and Nutanix AHV and has been optimized for each Hypervisor. By default, the VA is allocated 2 vCPUs, 4GB of RAM and has a 20GB virtual disk. The Virtual Appliance can be downloaded [here](#).

Note

The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.

Note

Please refer to [Virtual Appliance Installation](#) and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.

Note

The VA has 4 network adapters. For VMware only the first adapter (**eth0**) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

9.2. Initial Network Configuration



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

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

9.3. Accessing the Appliance WebUI

The WebUI is accessed using a web browser. By default, users are authenticated using Apache authentication. Users can also be authenticated against LDAP, LDAPS, Active Directory or Radius - for more information, please refer to [External Authentication](#).

 **Note** There are certain differences when accessing the WebUI for the cloud appliances. For details, please refer to the relevant [Quick Start / Configuration Guide](#).

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

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

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

 **Note** If you need to change the port, IP address or protocol that the WebUI listens on, please refer to [Service Socket Addresses](#).

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

Username: loadbalancer

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

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

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

Primary | Secondary Active | Passive Link 15 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 ? 2024-03-15 16:27:21 UTC

Would you like to run the Setup Wizard?

Accept
Dismiss

VIRTUAL SERVICE ⌵ IP ⌵ PORTS ⌵ CONNS ⌵ PROTOCOL ⌵ METHOD ⌵ MODE ⌵

No Virtual Services configured.

Network Bandwidth

Bytes/s

20 k
15 k
10 k
5 k
0

Thu 18:00 Fri 00:00 Fri 06:00 Fri 12:00

■	RX	3k Min,	4k Avg,	32675k Total,
■	TX	6k Min,	7k Avg,	56693k Total,

System Load Average

System Load

1.0
0.8
0.6
0.4
0.2
0.0

Thu 18:00 Fri 00:00 Fri 06:00 Fri 12:00

■	1m average	0.00 Min,	0.12 Avg,	0.60 Max
■	5m average	0.00 Min,	0.06 Avg,	0.21 Max
■	15m average	0.00 Min,	0.02 Avg,	0.08 Max

Memory Usage

- You'll be asked if you want to run the Setup Wizard which can be used to configure layer 7 services. Click **Dismiss** if you're following a guide or want to configure the appliance manually or click **Accept** to start the wizard.

Main Menu Options

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

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

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

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

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

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

Live Chat - Start a live chat session with one of our Support Engineers

9.4. Appliance Software Update

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

Determining the Current Software Version

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

Copyright © Loadbalancer.org Inc. 2002 – 2024
ENTERPRISE VA Max - v8.11.1

English ▼

Checking for Updates using Online Update

Note

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

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

Information: Version v8.11.1 is the current release. No updates are available

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

Note

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

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

Information: Update completed successfully.

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

Using Offline Update

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



Note

Please contact support@loadbalancer.org to check if an update is available and obtain the latest offline update files.

To perform an offline update:

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

Software Update

Offline Update

The following steps will lead you through offline update.

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

Archive: No file chosen

Checksum: No file chosen

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

9.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22 *	SSH
TCP & UDP	53 *	DNS / 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



Protocol	Port	Purpose
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

 **Note**

The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to [Service Socket Addresses](#).


9.6. HA Clustered Pair Configuration

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

10. Option 1 – Explicit Proxy Mode (Recommended)

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

Deployment Architecture



Notes

- Browser settings on client PCs must be changed to point at the Virtual Service (VIP) on the load balancer (see [Client Configuration](#)).
- The load balancer is configured in one-arm Layer 4 DR mode.
- The Web Proxies must be configured to accept traffic for the VIP (see [Modify the Web Proxies to accept traffic for the VIP](#)).
- 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](#).

Load Balancer Configuration

Create the Virtual Service (VIP)



- Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Services*.
- Click **Add a New Virtual Service**.
- 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"/>	?	

- Enter an appropriate label (name) for the VIP, e.g. **Proxy**.
- Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.2.202**.
- Set the *Virtual Service Ports* field to the required port, e.g. **8080**.
- Ensure that *Protocol* is set to **TCP**.
- Ensure that *Forwarding Method* is set to **Direct Routing**.
- Click **Update**.
- Now click **Modify** next to the newly created VIP.
- Ensure that *Persistence* is enabled. Set *Persistence Timeout* to **3600** (i.e. 1 hour).
- Click **Update**.

Define the Real Servers (RIPs)

- Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers*.
- Click **Add a New Real Server** next to the newly created VIP.
- Find the following fields:

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"/>	?

Enter the following details:



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

Web Proxy Appliance Configuration

Modify the Web Proxies to accept traffic for the VIP

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 Proxies. 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 Proxy 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 Proxy 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, refer to the [Administration Manual](#) and search for 'ARP Problem'.

 **Note** Vendors such as *Blox* and *Smoothwall* have options in their Web User Interface that allow this to be easily configured, so command line entries are not required. Please consult your specific vendor or loadbalancer.org for more information.


Finalize Settings

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

10.2. 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 [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 Proxies is configured to be the load balancer. For an HA pair, a floating IP address must be configured to allow the gateway IP 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 Proxies](#)).
- 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](#).

Load Balancer Configuration

Configure Network Settings

Two interfaces are required. Typically **eth0** is used for the internal (Web Proxy) 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:

IP Address Assignment


eth0


eth1

eth0	192.168.4.200/24	MTU <input type="text" value="1500"/> bytes
eth1	192.168.2.200/24	MTU <input type="text" value="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 Proxies

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 Proxies. This will "float" between the Primary and Secondary units in the event of a failover or failback. This ensures that the Web Proxies 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 IPs*.

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"/>	▼	?

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 that *Persistence* is enabled. 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"/>	?

4. Enter an appropriate label (name) for the first Web Proxy, 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 Proxy(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 Proxies) 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	<input type="text" value="eth1 (Default)"/>	?
Multi-threaded	<input type="text" value="yes"/>	?

2. 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.
3. Click **Update**.

Web Proxy Configuration

Configure the Default Gateway

As mentioned, Option 1B requires the default gateway on the Web Proxies 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 Proxies details how to create the Floating IP.

Note


Please refer to the relevant Web Proxy documentation for instructions on setting the default gateway. This should be done on all Web Proxies.

Finalize Settings

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

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

Deployment Architecture



Notes



- Browser settings on client PCs 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, a static route is configured on the Web Proxies 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 Proxies](#)).
- This method is more efficient & faster than Option 1B since the Web Proxies 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 Secondary Appliance](#).

Load Balancer Configuration

Configure Network Settings

Two interfaces are required. Typically **eth0** is used for the internal (Web Proxy) 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:

IP Address Assignment



eth0



eth1

eth0	<input style="width: 95%; height: 40px; border: 1px solid #ccc;" type="text" value="192.168.4.200/24"/>	MTU <input style="width: 40px; border: 1px solid #ccc;" type="text" value="1500"/> bytes
eth1	<input style="width: 95%; height: 40px; border: 1px solid #ccc;" type="text" value="192.168.2.200/24"/>	MTU <input style="width: 40px; border: 1px solid #ccc;" type="text" value="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 Proxies

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 Proxies. This will "float" between the Primary and Secondary units in the event of a failover or failback. This ensures that the Web Proxies 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 IPs*.

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"/>	▼	?

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**.
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"/>	?

4. Enter an appropriate label (name) for the first Web Proxy, 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 Proxy(s).

Web Proxy Configuration

Configure a Static Route

As mentioned, Option 1C requires a Static Route to be defined on the Web Proxies 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 Proxies](#) details how to create the Floating IP.

Note

Please refer to the relevant Web Proxy documentation for instructions on configuring a Static Route. This should be done on all Web Proxies.

Finalize Settings

Now refer to [Configuration Settings Common to Options 1A, 1B & 1C](#) below to finalize web proxy and client browser settings.

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



The steps in the following 2 sub sections must be followed for options 1A, 1B & 1C.

Web Proxy Operating Mode

Typically, there is a setting in the WebUI of the Web Proxy to allow the selection of either Explicit Proxy Mode or Transparent Routed Proxy Mode.

For Options 1A, 1B & 1C this should be set to Explicit Proxy Mode. The exact terminology does vary between vendors so please check your specific appliance.


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. 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 Proxies (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



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 [Configure Firewall Rules](#) for more details.

6. Clear the *Virtual Service Ports* field, the ports are defined in [Configure Firewall Rules](#).
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 that *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"/>
-----------------	--



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

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 section 2 in the Appendix if you intend to forward ALL traffic to the web proxies.

4. 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 Proxy, 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 Proxy(s).

11.3. Web Proxy Appliance Configuration

Web Proxy Operating Mode

Typically, there is a setting in the WebUI of the Web Proxy to allow the selection of either Explicit Proxy Mode or Transparent Routed Proxy Mode.

For Option 2 this should be set to Transparent Routed Proxy Mode. The exact terminology does vary between vendors so please check your specific appliance.

Note

When using Transparent Mode, it's not necessary to modify the Web Proxy 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 Proxies. 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 Proxy, 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 Proxy.

Note

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

11.4. 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:

System Overview

[Reports > Layer 4 Status](#)

[Reports > Layer 4 Current Connections](#)

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

12.1. 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 Proxy (192.168.64.60).

Layer 4 Current Connections

Check Status

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

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

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.

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	Fallback Page	Fallback page configuration
Maintenance	Firewall Script	Firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings



⚠ Important

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


Configuring the HA Clustered Pair

📌 Note

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


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

Create a Clustered Pair




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



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](#).

15.2.2 – Modified Transparent Mode Firewall Rules

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

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.



16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.7.0	10 September 2019	Styling and layout	General styling updates	RJC
1.7.1	17 January 2020	Added note explaining how to disable "Secure Mode" to unlock the firewall script page	Required update	RJC
1.7.2	22 July 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.8.0	1 January 2022	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.8.1	5 January 2023	Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section	Housekeeping across all documentation	AH
1.8.2	2 February 2023	Updated screenshots	Branding update	AH
1.8.3	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.9.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH



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://twitter.com/loadbalancer.org)

About Loadbalancer.org

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

