



Load Balancing Barracuda Web Filter

Deployment Guide **v1.2.5**

Table of Contents

1. About this Guide.....	4
2. Loadbalancer.org Appliances Supported.....	4
3. Loadbalancer.org Software Versions Supported.....	4
4. Barracuda Web Filter 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 Filter Deployment Mode.....	6
Explicit 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. Explicit Proxy Mode.....	10
Option 1A – Using DR (Direct Return) Mode (Recommended).....	10
Deployment Architecture.....	10
Load Balancer Configuration.....	11
Web Filter Configuration.....	12
Finalize Settings.....	12
Option 1B – Using NAT Mode.....	13
Deployment Architecture.....	13
Load Balancer Configuration.....	14
Web Filter Configuration.....	16
Finalize Settings.....	16
Option 1C – Using NAT Mode (Preferred NAT Topology).....	17
Deployment Architecture.....	17
Load Balancer Configuration.....	18
Web Filter Configuration.....	20
Finalize Settings.....	20
Configuration Settings Common to Options 1A, 1B & 1C.....	20
Web Filter Operating Mode.....	20
Proxy Port Configuration.....	20
Client Configuration.....	21
11. Testing & Validation.....	22
Layer 4 – Current Connections.....	22
12. Technical Support.....	23

13. Further Documentation.....	23
14. Conclusion.....	23
15. Appendix.....	24
1 – Clustered Pair Configuration – Adding a Slave Unit.....	24
2 - Company Contact Information.....	26

1. About this Guide

This guide details the steps required to configure a load balanced Barracuda Web Filter environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Barracuda Web Filter 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 Barracuda Web Filters. 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. Barracuda Web Filter Appliances Supported

- All versions

5. Benefits of Implementing a Load Balancer

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

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

DEPLOYMENT MODES

LAYER 4 (RECOMMENDED)

DR Mode - Direct Server Return Mode (Recommended)

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

PERSISTENCE / SERVER AFFINITY

Persistence may or may not be required and depends on the specific Web Filter 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 Filter.

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 Filter 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 Filters.

7. Web Filter Deployment Mode

When using a load balancer Barracuda suggest using *Forward Proxy Mode*. This is also commonly referred to as *Explicit Proxy Mode*. For more information please refer to 'Method 3' in the 'High Availability Deployment Options' section at [this Barracuda URL](#).

EXPLICIT PROXY MODE

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 Filters. Please refer to the section starting on page [10](#) for configuration details.

8. Summary of Deployment Options

Option	Web Filter Mode	Load Balancer Mode	Notes
Option 1A (Recommended)	Forward Proxy Mode	DR Mode	The Web Filters must be configured to accept traffic for the VIP. Please refer to page 10 for configuration details.
Option 1B	Forward Proxy Mode	NAT Mode	The load balancer must be set as the default gateway for the Web Filters. Please refer to page 13 for configuration details.

Option 1C	Forward Proxy Mode	NAT Mode	<p>A static route must be configured on the Web Filters to send client return traffic back via the load balancer.</p> <p>Please refer to page 17 for configuration details.</p>
-----------	--------------------	----------	---

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 Loadbalancer.org Enterprise VA MAX interface. On the left is a navigation menu with items like System Overview, Local Configuration, Cluster Configuration, Maintenance, View Configuration, Reports, Logs, and Support. The top right shows 'Enterprise VA MAX' branding. Below the navigation is a status bar with 'Master | Slave', 'Active | Passive', and 'Link' indicators, along with a refresh button and '5 Seconds' timer. 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 'No Virtual Services configured.' Three graphs are shown: 'Network Bandwidth' (Bytes/s), 'System Load Average', and 'Memory Usage' (Bytes). Each graph has a legend with min, avg, and max values.

Direction	Min	Avg	Total
RX	2k	4k	1853k
TX	11k	45k	18736k

Average	Min	Avg	Max
1m average	0.36	0.38	0.39
5m average	0.09	0.13	0.17
15m average	0.03	0.05	0.07

Category	Min	Avg	Max
Used	117.78M	122.85M	127.92M
Page	79.52M	79.92M	80.32M
Buffer	10.86M	11.14M	11.43M
Free	1812.93M	1819.67M	1826.41M

(shows v8.2.x)

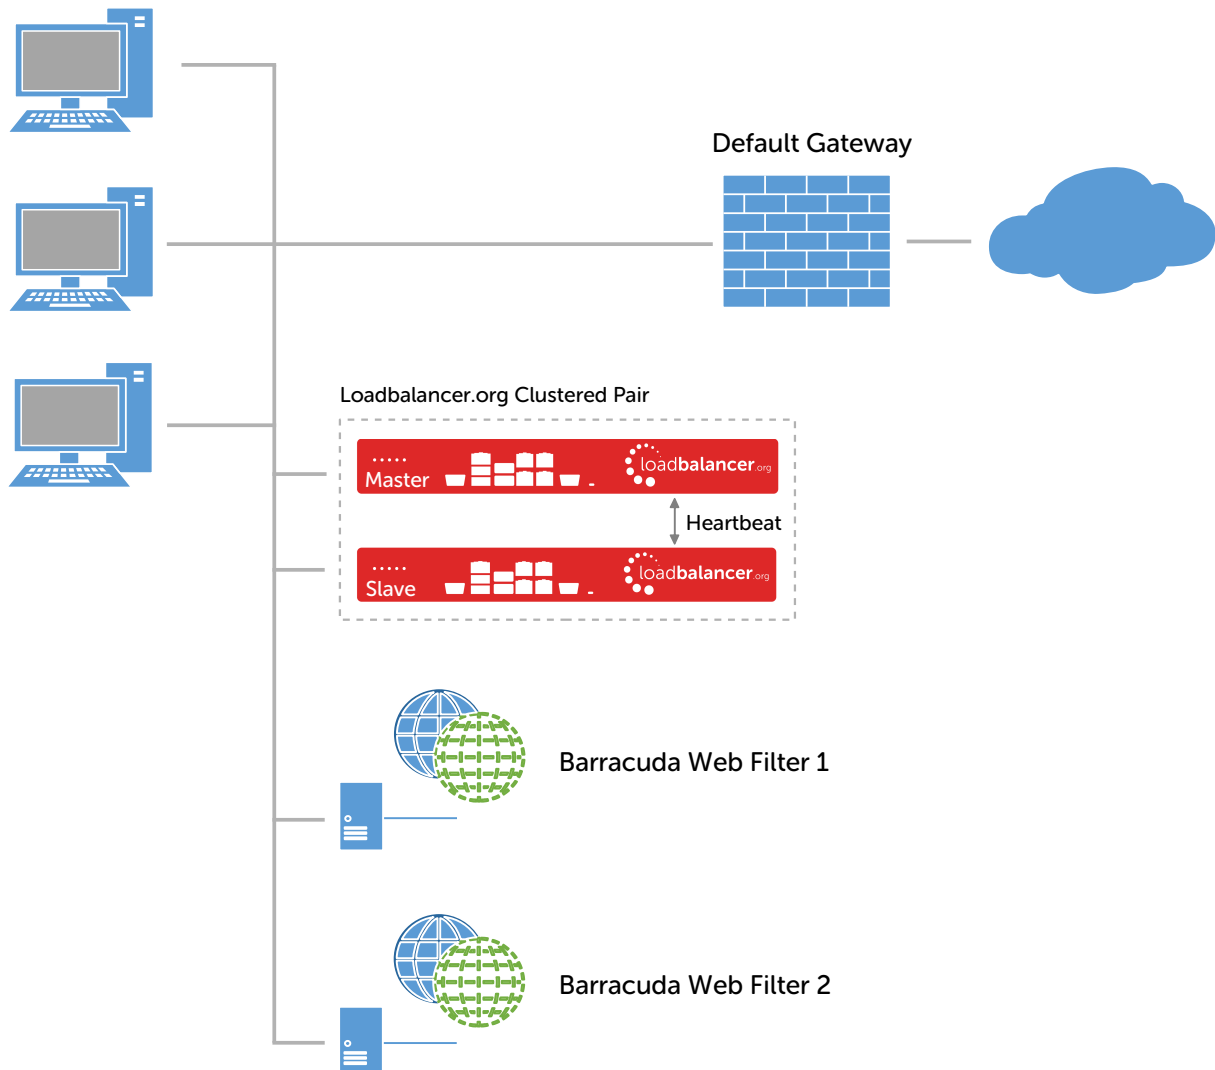
HA CLUSTERED PAIR CONFIGURATION

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 [24](#).

10. Explicit Proxy Mode

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 [21](#))
- The load balancer is configured in one-arm Layer 4 DR mode
- The Web Filters 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 [24](#)
- For more information on Barracuda Web Filter 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 Filter, 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 Filter(s)

WEB FILTER CONFIGURATION

Modify the Web Filters 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 Filters. 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 Filter 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 Filter 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 Barracuda Appliance

Note:

For Barracuda Web Filters you must contact Barracuda support directly to enable DR mode (aka DSR mode). Please refer to [this Barracuda URL](#) for more information. Since the appliance does not permit root access by default, you'll also need to consult Barracuda support on how to add the required iptables rule to an appropriate startup script.

Suggested steps:

1. Login as root
2. Edit the file `/etc/rc.local` and add the following line to the file, setting the VIP address as required:

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

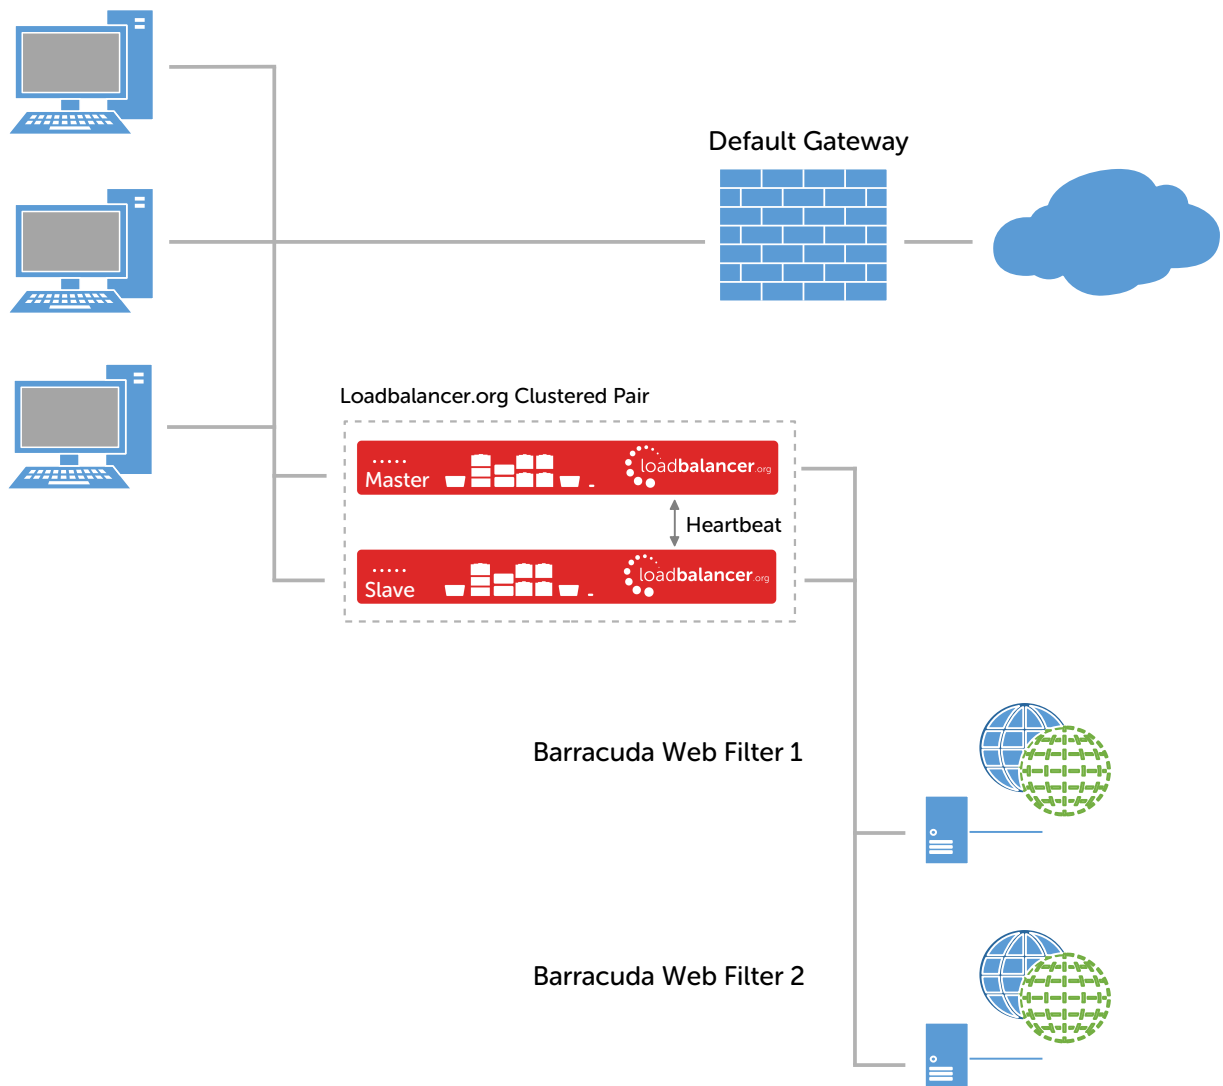
3. Reboot the Web Filter to apply the new setting

FINALIZE SETTINGS

Now refer to the section "*Configuration Settings Common to Options 1A, 1B & 1C*" on page [20](#) to finalize Web Filter 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 [21](#))
- 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 Filters 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 [14](#))
- 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 Barracuda Web Filter 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 Filter) 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 and eth1 are active (10 GB/s), while eth2 and eth3 are disabled (indicated by a red 'X'). 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 Filters

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 Filters. This will 'float' between the master and slave units in the event of a failover or failback. This ensures that the Web Filters 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' containing the value 192.168.4.205. To the right of the input field is a green button labeled 'Add Floating IP'.

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

3. Click **Add Floating IP**

Create the Virtual Service (VIP)

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

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

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

Configure the Default Gateway

As mentioned, Option 1B requires the default gateway on the Web Filters 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 [14](#) details how to create the Floating IP.

Note:

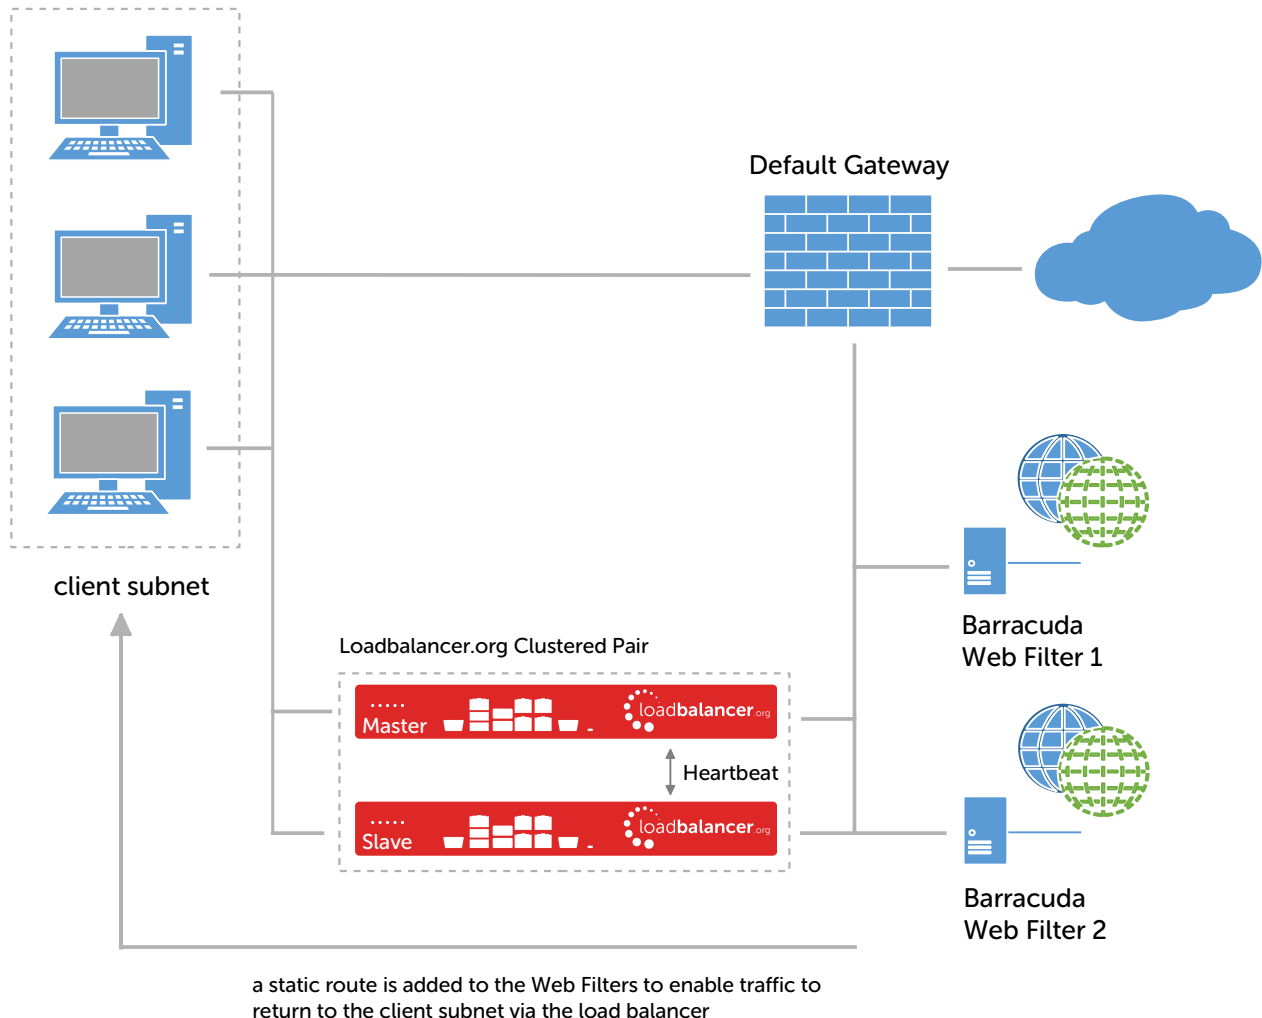
This should be done on all Web Filters. Please refer to the Barracuda Web Filter documentation for instructions on setting the default gateway.

FINALIZE SETTINGS

Now refer to the section "*Configuration Settings Common to Options 1A, 1B & 1C*" on page [20](#) to finalize Web Filter 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 [21](#))
- 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 Filters 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 [18](#))
- This method is more efficient & faster than Option 1B since the Web Filters 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 Barracuda Web Filter 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 Filter) 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 and eth1 are active (10 GB/s), while eth2 and eth3 are disabled (indicated by a red 'X'). 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 gateway for the Static Route on the Web Filters

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 Filters. This will 'float' between the master and slave units in the event of a failover or failback. This ensures that the Web Filters 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' containing the value 192.168.4.205. To the right of the input field is a green button labeled 'Add Floating IP'.

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

3. Click **Add Floating IP**

Create the Virtual Service (VIP)

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

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

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

WEB FILTER CONFIGURATION

Configure a Static Route

As mentioned, Option 1C requires a Static Route to be defined on the Web Filters 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 [18](#) details how to create the Floating IP.

Note:

This should be done on all Web Filters. Please refer to the Barracuda Web Filter documentation for instructions on configuring a Static Route.

FINALIZE SETTINGS

Now refer to the section "*Configuration Settings Common to Options 1A, 1B & 1C*" below to finalize Web Filter 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 FILTER OPERATING MODE

By default the Barracuda Web Filter supports Forward Proxy Mode on port 8080 with no configuration changes.

PROXY PORT CONFIGURATION

If required the proxy port can be changed using the WebUI option: *Advanced > Proxy* and defining the required port in the *Proxy Port* field.

The screenshot shows the Barracuda Web Filter 610Vx administration interface. The user is logged in as 'admin' and is viewing the 'Proxy Settings' page under the 'ADVANCED' tab. The interface includes a navigation menu with options like Backup, Energize Updates, Firmware Update, External Servers, Appearance, Remote Filtering, Syslog, Linked Management, Secure Administration, Delegated Admin, Troubleshooting, and Proxy. The Proxy Settings page has a 'Save Changes' and 'Help' button. The settings include:

- Send VIA Header: Yes No
- Send Forwarded-For Header: Yes No
- Send CUDA_CLIP: Yes No
- Peer Proxy IP: [Text Input]
- Peer Proxy Port: [Text Input]
- Only Peer Proxy These Domains: [Table with columns: Domain, Disable, Bulk Edit]
- Transparent Proxy Port: 8080

Help text on the right side of the settings:

- By enabling this feature, the Barracuda Web Filter proxy will expose its identity to Web servers.
- By enabling this feature, the HTTP requests will expose the proxy that the Barracuda Web Filter is using and the client IP address for which it is forwarding.
- Enabling this feature will send the proprietary CUDA_CLIP header in the request
- IP address and port of any pre-existing proxy in front of the Barracuda.
- Domain names for which you want to send traffic to the specified peer proxy, using this format: .example.com
- Port to use to directly proxy through the Barracuda Web Filter. Default: 8080

CLIENT CONFIGURATION

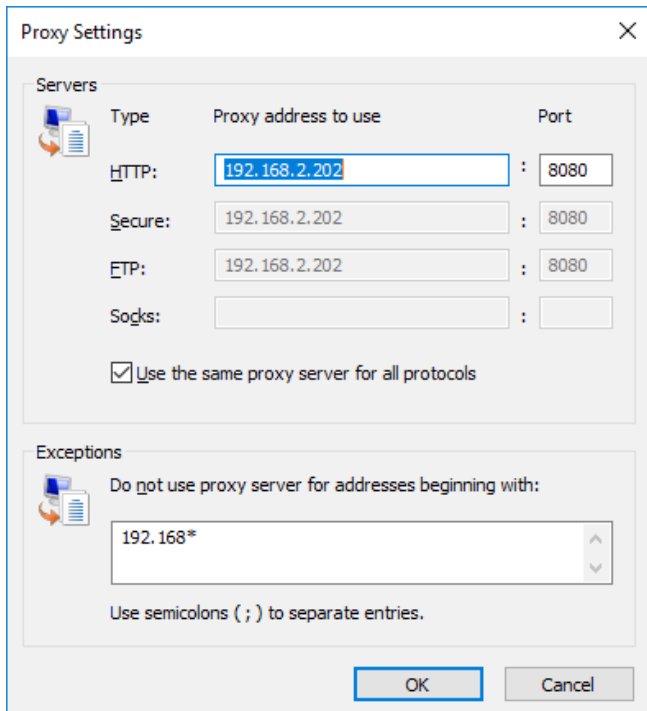
Client browser settings must be set so that browsers connect via the VIP. In a Microsoft based LAN environment, this is typically achieved using AD group policy.

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 Windows 'Local Area Network (LAN) Settings' dialog box. The 'Automatic configuration' section has 'Automatically detect settings' and 'Use automatic configuration script' both unchecked. The 'Address' field is empty. The 'Proxy server' section has 'Use a proxy server for your LAN (These settings will not apply to dial-up or VPN connections)' 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. The 'OK' and 'Cancel' buttons are at the bottom.



11. 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

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

LAYER 4 – CURRENT CONNECTIONS

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 Filter (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

12. 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

13. 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>

14. Conclusion

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

15. 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 - COMPANY CONTACT INFORMATION

Website	URL: www.loadbalancer.org
North America (US)	<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>
North America (Canada)	<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>
Europe (UK)	<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>
Europe (Germany)	<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>