



Appliance Quick Start Guide

v8.0

rev. 1.0.6

Copyright © 2002 – 2015 Loadbalancer.org, Inc



Table of Contents

About this Guide.....	5
About the Appliance.....	5
Appliance Configuration Overview.....	5
Appliance Security.....	5
Loadbalancer.org Terminology.....	6
What is a Virtual IP Address?.....	7
What is a Floating IP Address?.....	7
What are Your Objectives?.....	7
What is the Difference Between a One-Arm and a Two-Arm Configuration?.....	8
What Load Balancing Methods are Supported?.....	8
Direct Routing (DR).....	10
Network Address Translation (NAT).....	11
Source Network Address Translation (SNAT).....	12
Physical Appliance.....	13
Initial Network Interface Configuration.....	14
Using the Network Setup Wizard (Recommended).....	14
Using Linux Commands.....	15
IPv6 Support.....	16
Accessing the Web User Interface (WUI).....	17
Appliance Configuration Using the Setup Wizard.....	17
Running the Wizard.....	17
Appliance Configuration Using the Web User Interface (WUI).....	19
Adding Virtual Services (VIPs).....	21
Defining Real Servers (RIPs).....	22
Configuring the Physical Real Servers.....	23
Configuring the Real Servers for Layer 4 NAT Mode.....	23
Configuring the Real Servers for Layer 4 DR Mode.....	23
Detecting the ARP Problem.....	23
Solving the ARP Problem for Linux.....	23
Method 1 (using iptables).....	23
Method 2 (using arp_ignore sysctl values).....	24
Solving the ARP Problem for Solaris & MAC OS X / BSD.....	25
Solving the ARP Problem for Windows Servers.....	26
Windows Server 2008.....	26
Windows Server 2012.....	30
Verifying netsh Settings for Windows 2008 & 2012.....	34
Configuring the Application/Service to Respond to both the RIP and VIP.....	35
Configuring the Real Servers for Layer 7 SNAT Mode.....	35
High-Availability Configuration of two Loadbalancer.org Appliances.....	36
Clustered Pair Configuration.....	36
Testing Load Balanced Services.....	38
Diagnosing VIP Connection Problems.....	38
Taking Real Servers Offline (Halting).....	39
Using Reports & Log Files.....	40
Testing High-Availability for a Loadbalancer.org HA-Pair.....	40
Does Your Application Cluster Correctly Handle its Own State?.....	42
Replication Solutions for Shared Data.....	42
Solutions for Session Data.....	42
Persistence (aka Affinity).....	42
What do You do if Your Application is not Stateless?.....	43
Loadbalancer.org Persistence Options.....	43
Loadbalancer.org Technical Support.....	43
Appendix.....	45
Company Contact Information.....	45
Front & Rear Panel Layouts.....	47
Physical Appliance (back page reference).....	48

About this Guide

This quickstart guide provides enough information to deploy the appliance and configure a basic load balancing solution. For more detailed information please refer to the administration manual available here:

<http://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf>

About the Appliance

The Loadbalancer.org appliance runs the GNU/Linux operating system with a custom kernel configured for load balancing.

The core software is based on customized versions of Centos 6.x / RHEL 6.x, Linux 3.10.x, LVS, HA-Linux, HAProxy, Pound, STunnel & Ldirectord. Full root access is provided which enables complete control of all settings.

The appliance is available in the following formats: hardware, virtual (VMware, HyperV, KVM, XEN) and cloud based (Amazon, Azure).

Appliances can be deployed as single units or as a clustered pair.



NOTE : Loadbalancer.org always recommend that clustered pairs should be used where possible for high availability and resilience, this avoids introducing a single point of failure to your network. For more information on configuring an HA pair please refer to page 36 .

Appliance Configuration Overview

Initial network configuration can be carried out on the console by using the Network Setup Wizard, using standard Linux network setup commands, or by connecting to the default IP address:port in a browser (**192.168.2.21:9080**) and making changes using the WUI.

Once the network is configured, the appliance can be configured manually or by using the Setup Wizard (for layer 7 services). The WUI is accessible using HTTP on port **9080** and HTTPS on port **9443**. It's also possible to configure the load balancer at the console using the text based Links browser, although using the WUI is the recommended method.

For a clustered pair, we recommend that the master is fully configured first, then the slave should be added. For more information on configuring a clustered pair, please refer to chapter 9 in the administration manual. Once a pair is configured, from v8.0 load balanced services can be configured / modified on the master **OR** the slave and the peer appliance will be kept automatically in sync.

Appliance Security

The appliance includes a security lockdown command (**lbsecure**) that enables passwords to set, network access to be locked down and SSH key regeneration in one simple step. This command can be run on a single appliance or an HA pair. For more details please refer to chapter 5 in the administration manual.

Loadbalancer.org Terminology

<u>Acronym</u>	<u>Terminology</u>
Load Balancer	An IP based traffic manager for server clusters
VIP	Virtual IP address – the address of the load balanced cluster of RIPv, the address presented to connecting clients
Floating IP	The Floating IP Address is automatically created whenever a VIP is configured, the FIP address is the same as the VIP address, it enables services to be moved between the master and slave appliance
RIP	The Real IP address of a back-end server in the cluster
GW	The Default Gateway for a back-end server in the cluster
WUI	Web User Interface
Layer 4	Part of the seven layer OSI model, descriptive term for a network device that can route packets based on TCP/IP header information
Layer 7	Part of the seven layer OSI model, descriptive term for a network device that can read and write the entire TCP/IP header and payload information at the application layer
DR	Direct Routing (aka DSR / Direct Server Return) is a standard load balancing technique that distributes packets by altering only the destination MAC address of the packet
NAT	Network Address Translation – Standard load balancing technique that changes the destination of packets to and from the VIP (external subnet to internal cluster subnet)
SNAT (HAProxy)	Source Network Address Translation – the load balancer acts as a proxy for all incoming & outgoing traffic
SSL Termination (Pound & STunnel)	The SSL certificate is installed on the load balancer in order to decrypt HTTPS traffic on behalf of the cluster
MASQUERADE	Descriptive term for standard firewall technique where internal servers are represented as an external public IP address. Sometimes referred to as a combination of SNAT & DNAT rules
One-Arm	The load balancer has one physical network card connected to one subnet
Two-Arm	The load balancer has two interfaces connected to two subnets - this can be achieved using two physical network cards or by assigning two addresses to one physical network card
Eth0	Usually the internal interface also known as Gb0
Eth1	Usually the external interface also known as Gb1

What is a Virtual IP Address?

Most load balancer vendors use the term Virtual IP address (VIP) to describe the address that the cluster is accessed from. It's important to understand that the Virtual IP address (VIP) refers to both the physical IP address and also to the logical load balancer configuration. Likewise the real IP (RIP) address refers to both the Real Servers physical IP address and its representation in the logical load balancer configuration.



NOTE : It's not possible to configure a VIP on the same IP address as any of the network interfaces. This ensures services can 'float' (move) between master and slave appliances.

What is a Floating IP Address?

A floating IP address is automatically created whenever a VIP is configured. The FIP address is the same as the VIP address. Since the FIP must be able to move between the master and slave appliance, it's not possible to configure a VIP/FIP on the same IP address as an interface as mentioned in the note above. This restriction also applies to single appliance (i.e. master only) deployments. FIPs can also be manually defined to provide a 'floating default gateway' for layer 4 NAT mode configurations. This allows the default gateway for the NAT mode Real Servers to be brought up on the slave should the master fail.

What are Your Objectives?

It's important to have a clear focus on your objectives and the required outcome for the successful implementation of your load balancing solution. If the objective is clear and measurable, you know when you have achieved the goal.

Load balancers have a number of flexible features and benefits for your technical infrastructure and applications. The first question to ask is:

Are you looking for increased performance, reliability, ease of maintenance or all three?

Performance	A load balancer can increase performance by allowing you to utilize several commodity servers to handle the workload of one application
Reliability	Running an application on one server gives you a single point of failure. Utilizing a load balancer moves the point of failure to the load balancer. At Loadbalancer.org we advise that you only deploy load balancers as clustered pairs to remove this single point of failure (for more details on configuring a clustered pair please refer to page 36)
Maintenance	Using the appliance, you can easily bring servers on and off line to perform maintenance tasks, without disrupting your users



NOTE : In order to achieve all three objectives of performance, reliability & maintenance in a web based application, your application must handle persistence correctly (see page 42 for more details).

What is the Difference Between a One-Arm and a Two-Arm Configuration?

The number of 'arms' is normally a descriptive term for how many physical connections (Ethernet interfaces) are used to connect a device to a network. It's very common for a load balancer that uses a routing method (NAT) to have a two-arm configuration. Proxy based load balancers (SNAT) commonly use a one-arm configuration.

One-Arm	The load balancer has one physical network card connected to one subnet
Two-Arm	The load balancer has two network interfaces connected to two subnets – this can be achieved by using two physical network cards or by assigning two addresses to one physical network card

What Load Balancing Methods are Supported?

The Loadbalancer.org appliance is one of the most flexible load balancers on the market. The design allows different load balancing modules to utilize the core high availability framework of the appliance. Multiple load balancing methods can be used at the same time or in combination with each other.

Layer 4	DR (Direct Routing)	Ultra-fast local server based load balancing <i>Requires solving the 'ARP problem' on the Real Servers</i>	One-Arm
Layer 4	NAT (Network Address Translation)	Fast Layer 4 load balancing, the appliance becomes the default gateway for the Real Servers	Two-Arm
Layer 4	TUN	Similar to DR but works across IP encapsulated tunnels	One-Arm
Layer 7	SSL Termination (Pound & STunnel)	Usually required in order to process cookie persistence in HTTPS streams on the load balancer <i>Processor intensive</i>	One or Two-Arm
Layer 7	SNAT (Source Network Address Translation: HAProxy)	Layer 7 allows great flexibility including full SNAT and WAN load balancing, cookie insertion and URL switching <i>Not as fast as Layer 4</i>	One or Two-Arm

Key:



Recommended for high performance fully transparent and scalable solutions



Recommended if HTTP cookie persistence is required, also used for several Microsoft applications such as Exchange, Sharepoint & Remote Desktop Services and for overall deployment simplicity since real servers can be on any accessible subnet and no Real-Server changes are required



Only required for Direct Routing implementation across routed networks (rarely used)

Loadbalancer.org Recommendation:

Where feasible, one-arm layer 4 Direct Routing (DR) mode is our recommended method because it's a very high performance solution with little change to your existing infrastructure.



NOTE : Sometimes it's not possible to use DR mode. The two most common reasons being: if the application cannot bind to the RIP & VIP at the same time; or if the host operating system cannot be modified to handle the ARP problem (see page 23 for more details).

A second option is layer 4 Network Address Translation (NAT) mode. This is a fairly high performance solution but it 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). Network engineers with experience of hardware load balancers will have often used this method.

The third option is layer 7 Source Network Address Translation (SNAT) mode using HAProxy. If your application requires that the load balancer handles cookie insertion, RDP cookies, Session Broker integration or SSL termination then this option is appropriate. This can be deployed in one-arm or two-arm mode and does not require any changes to the application servers. HAProxy is a high-performance solution that operates as a full proxy, but due to this it cannot perform as fast as the layer 4 solutions.



NOTE : If your application doesn't maintain its own state information then you may need to use cookie insertion to maintain server persistence (affinity).

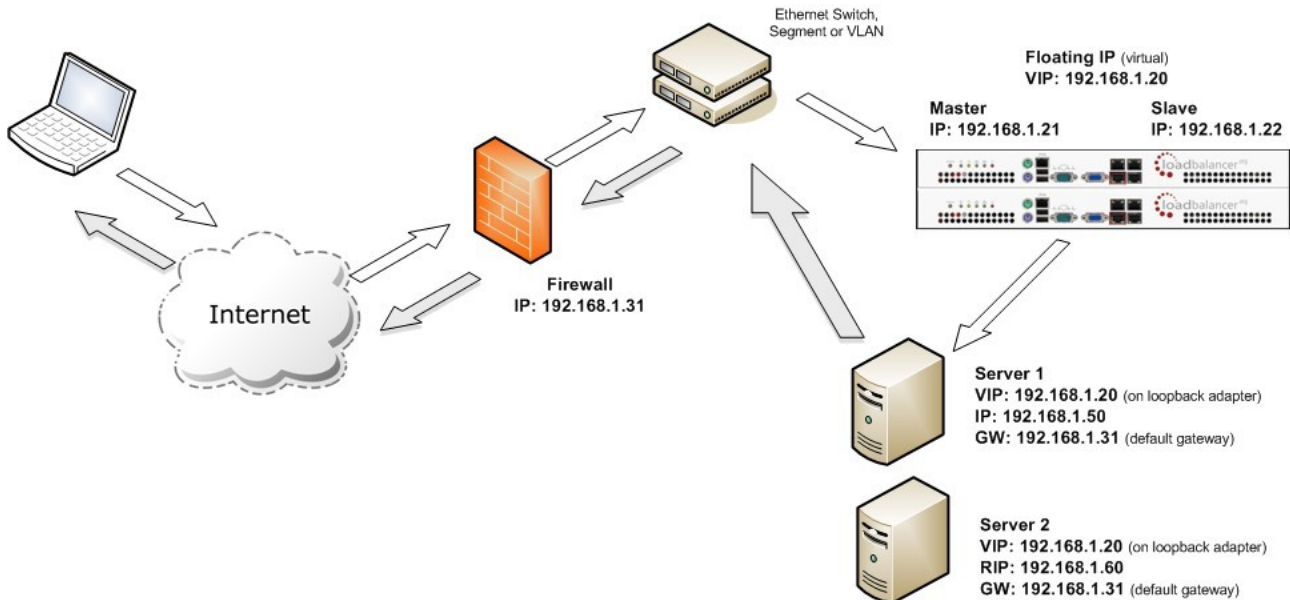
The following sections describe these configurations in more details.



IMPORTANT NOTE : If you are using Microsoft Windows Real Servers (i.e. back-end servers) make sure that Windows NLB (Network Load Balancing) is completely disabled to ensure that this does not interfere with the operation of the load balancer.

Direct Routing (DR)


One-arm Direct Routing (DR) mode is a very high performance solution that requires little change to your existing infrastructure. *N.B. Kemp, Brocade, Barracuda & A10 Networks call this Direct Server Return (DSR) and F5 call it nPath.*



- Direct Routing mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast
- When the packet reaches the Real Server it expects it to own the Virtual Services IP address (VIP). This means that you need to ensure that the Real Server (and the load balanced application) respond to both the Real Servers own IP address and the VIP
- The Real Server should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Servers in this way is referred to as *Solving the ARP Problem*. Please refer to page 23 for more details on this
- On average, DR mode is 8 times quicker than NAT for HTTP, 50 times quicker for Terminal Services and much, much faster for streaming media or FTP
- The load balancer must have an Interface in the same subnet as the Real Servers to ensure layer 2 connectivity required for DR mode to work
- The VIP can be brought up on the same subnet as the Real Servers, or on a different subnet provided that the load balancer has an interface in that subnet
- Port translation is not possible in DR mode i.e. having a different RIP port than the VIP port
- DR mode is transparent, i.e. the Real Server will see the source IP address of the client

Network Address Translation (NAT)

Sometimes it's not possible to use DR mode. The two most common reasons being: if the application cannot bind to the RIP & VIP at the same time; or if the host operating system cannot be modified to handle the ARP problem. The second choice is Network Address Translation (NAT) mode. This is also a high performance solution but it 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).



- The load balancer translates all requests from the external Virtual Service to the internal Real Servers
- Normally eth0 is used for the *internal* network and eth1 is used for the *external* network although this is not mandatory. If the Real Servers require Internet access, Autonat should be enabled using the WUI option: *Cluster Configuration > Layer 4 – Advanced Configuration*, the external interface should be selected
- NAT mode can be deployed in the following ways:

2-arm (using 2 Interfaces), 2 subnets (as shown above) - One interface on the load balancer is connected to subnet1 and the second interface and Real Servers are connected to subnet2. The VIP is brought up in subnet1. The default gateway on the Real Servers is set to be an IP address in subnet2 on the load balancer. Clients can be located in subnet1 or any remote subnet provided they can route to the VIP


2-arm (using 1 Interface), 2 subnets - same as above except that a single interface on the load balancer is allocated 2 IP addresses, one in each subnet

1-arm (using 1 Interface), 1 subnet - Here, the VIP is brought up in the same subnet as the Real Servers. For clients located in remote networks the default gateway on the Real Servers must be set to be an IP address on the load balancer. For clients located on the same subnet, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer - for more details on 'One-Arm NAT Mode' please refer to chapter 6 in the [administration manual](#)

- If you want Real Servers to be accessible on their own IP address for non-load balanced services, e.g. SMTP or RDP, you will need to setup individual SNAT and DNAT firewall script rules for each Real Server or add additional VIPs for this - please refer to chapter 6 in the [administration manual](#)
- NAT mode is transparent, i.e. the Real Server will see the source IP address of the client
- Port translation is possible in NAT mode, i.e. VIP:80 → RIP8080 is possible

Source Network Address Translation (SNAT)

If your application requires that the load balancer handles cookie insertion then you need to use the SNAT configuration. This mode is also used with numerous Microsoft applications such as Exchange, Sharepoint, Lync etc.



This mode has the advantage of a one arm configuration and does not require any changes to the application servers. However, since the load balancer is acting as a full proxy it doesn't have the same raw throughput as the layer 4 methods.

The network diagram for the Layer 7 HAProxy SNAT mode is very similar to the Direct Routing example except that no re-configuration of the Real Servers is required. The load balancer proxies the application traffic to the servers so that the source of all traffic becomes the load balancer.

- SNAT is a full proxy and therefore load balanced Real Servers do not need to be changed in any way
- Because SNAT is a full proxy any server in the cluster can be on any accessible subnet including across the Internet or WAN
- SNAT is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancers IP address. If required, this can be solved by either enabling TProxy on the load balancer, or for HTTP, using X-forwarded-For headers. Please refer to chapter 6 in the [administration manual](#) for more details.
- SNAT mode can be deployed using either a 1-arm or 2-arm configuration




NOTE : For detailed configuration examples using various load balancing methods, please refer to [chapter 11](#) in the [administration manual](#).

Physical Appliance

- Remove all packaging
- Rack mount the appliance if required
- The power supply is an auto sensing unit (100v to 240v)
- Connect the power lead from the power socket to the mains or UPS
- Connect a network cable from the switch to one of the Ethernet ports – typically *eth0* but this is not mandatory
- If using a two-armed configuration connect another cable to a second Ethernet port – typically *eth1* but this is not mandatory
- For a clustered hardware pair (see page 36), the units must be able to communicate either via network (ucast), via serial cable or both. By default, ucast only is used. If serial is preferred or you want to use both methods, connect a serial cable (1 supplied with each appliance) between the two appliances.

N.B. If a serial cable is used, Heartbeat must be configured for this using the WUI option: Cluster Configuration > Heartbeat Configuration and enabling 'Serial'

- Attach a monitor to the VGA port and keyboard to the USB or PS/2 port
- Check mains power is on and press the power switch to start the appliance (the fans should start & front panel LED's should light)
- Allow a minute for booting
- Refer to page 14 for details on configuring network settings
- Refer to page 17 for details on accessing the Web User Interface



N.B. The above image shows the Enterprise MAX, for connecting other models please refer to page 47

Initial Network Interface Configuration

By default the load balancer is pre-configured with the following IP address & subnet mask:

192.168.2.21 / 24 (192.168.2.21 / 255.255.255.0)

This default address can be changed at the console in two ways:

- Using the built-in Network Setup Wizard
- Using traditional Linux commands



NOTE : For the VA, four NICs are included but only eth0 is connected by default at power on. If the other NICs are required, these should be connected using the network configuration screen within the Hypervisor.

Using the Network Setup Wizard (Recommended)

To run the wizard, login to the console of the appliance as the 'setup' user. This is explained in the initial console start-up message as shown below:

```
Welcome to the Loadbalancer.org appliance.

To perform initial network configuration, log in to the console as
Username: setup
Password: setup

To access the web interface and wizard, point your browser at
http://192.168.2.21:9080/
or
https://192.168.2.21:9443/

lbmaster login: _
```

- login to the console:
Username: setup
Password: setup
- Once logged in, enter the IP address/mask, default gateway & DNS servers at the prompts as shown below:

```
Loadbalancer.org basic network set up

Static IP address (eg. 192.168.0.26)      : 192.168.67.23/18
Default gateway (eg. 192.168.0.1)         : 192.168.64.1
DNS Servers
  Primary (eg. 192.168.0.250)             : 192.168.64.1
  Secondary (Leave blank to omit)          :
```


After the required settings have been entered, a summary will be presented along with details of how to access the WUI as shown below:

```
Summary of settings
Static IP address:      192.168.67.23/18
Default gateway:       192.168.64.1
DNS servers:           192.168.64.1

You may now connect the eth0 network interface to your switch, and
continue configuration through the web interface on:

    http://192.168.67.23:9080/lbadmin/

Press any key...
```

As mentioned in the text the IP address is now configured for interface eth0.

IP addresses for the other interfaces can now be configured using the WUI option: *Local Configuration > Network Interface Configuration* (to access the WUI please refer to page 17) or by using Linux commands as explained in the following section.

Using Linux Commands

To set the IP address, login to the console or an SSH session as root:

Username: root
Password: loadbalancer

set the IP address using the following command:

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

e.g.

```
ip addr add 192.168.1.100/24 dev eth0
```

set the default gateway using the following command:

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

e.g.



```
route add default gw 192.168.1.254 eth0
```

N.B. Setting the IP address in this way is temporary, the IP address MUST be set via the WUI to make this permanent otherwise settings will be lost after a reboot

IPv6 Support

The appliance fully supports IPv6. This allows Virtual Services to be configured using IPv4 addresses or IPv6 addresses. It's also possible to mix IPv4 and IPv6 addresses on a single appliance as illustrated below using the WUI. Please refer to the next section to access the WUI

IP Address Assignment

	 eth0 10 GB/s	 eth1 10 GB/s
eth0	<div>192.168.1.20/24 192.168.2.20/24</div>	<div>MTU 1500 bytes</div>
eth1	<div>10.15.1.1/16 2001:cdba:0000:0000:0000:0000:3257:9655/64</div>	<div>MTU 1500 bytes</div>

Configure Interfaces

Once the required addresses are defined, use the **Configure Interfaces** button to apply the new settings.

Accessing the Web User Interface (WUI)

- Using a web browser, access the WUI using the following URL:

http://192.168.2.21:9080/lbadmin/

(replace 192.168.2.21 with your IP address if it's been changed)

N.B. If you prefer you can use the HTTPS administration address:

https://192.168.2.21:9443/lbadmin/

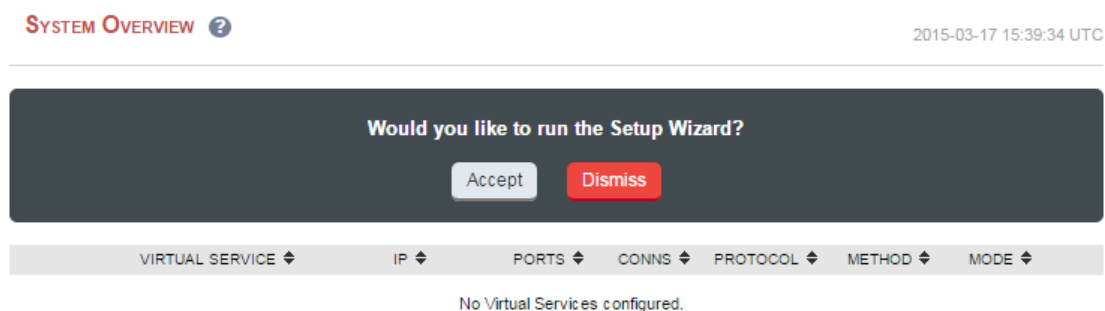
(replace 192.168.2.21 with your IP address if it's been changed)

- Login to the WUI:

Username: loadbalancer

Password: loadbalancer

- Once logged in, you'll be asked if you want to run the web based setup wizard. If you click **[Accept]** the Layer 7 Virtual Service configuration wizard will start. If you prefer to configure the appliance manually, simple click **[Dismiss]**



Appliance Configuration Using the Setup Wizard

The wizard can be used to setup one or more Layer 7 Virtual Services and associated Real Servers. Layer 4 services must currently be configured manually.

Running the Wizard

First, set the IP address using one of the methods described starting on page 14. Then connect to the WUI and start the wizard by clicking the **[Accept]** link shown above, or by using the WUI option: *Cluster Configuration > Setup Wizard* and clicking **General Layer 7 Virtual Service**, then continue as detailed below:

- Define the required Virtual Service settings as shown in the example below:

SETUP WIZARD - GENERAL LAYER 7 VIRTUAL SERVICE

Load balancer configuration	
	Master
Hostname	lbmaster
Static IP Addresses	192.168.111.222/18
Floating IP Addresses	

Create a new Layer 7 Virtual Service

Label	VIP1
Virtual Service	IP Address
	192.168.111.225
	Ports
	80
Layer 7 Protocol	TCP Mode

Create Virtual Service

Select the Layer 7 protocol to be handled by this Virtual Service.
Advanced options may be set by editing this Virtual Service once it has been created.

- Click **Create Virtual Service**
- Now continue and add the associated Real Servers as shown below:

Attach Real Servers

Label	IP Address	Port	Weight	
RIP1	192.168.111.226	80	100	✕
RIP2	192.168.111.227	80	100	✕

Add Real Server

Attach Real Servers

- Use the **Add Real Server** button to define additional Real Servers, once all are defined click **Attach Real Servers**
- Finally reload HAProxy using the **Reload HAProxy** button in the blue box at the top of the screen or by using the WUI option: *Maintenance > Restart Services* and clicking **Reload HAProxy**



NOTE : Running the wizard again will permit additional Layer 7 VIPs and associated RIPv6s to be defined.



NOTE : To restore manufacturer's settings use the WUI option: *Maintenance > Backup & Restore > Restore Manufacturer's Defaults*. *N.B. this will reset the IP address to 192.168.2.21/24*

Appliance Configuration Using the Web User Interface (WUI)

If you have already used the web based wizard, then you will already be using the WUI. From here all administration tasks can be carried out. If not, access the WUI as follows:

With a web browser access the WUI: ***http://192.168.2.21:9080/lbadmin/***

(replace 192.168.2.21 with the correct IP address)

log in to the WUI: ***Username:*** loadbalancer
 Password: loadbalancer

*N.B. If you prefer you can use the HTTPS administration address: ***https://192.168.2.21:9443/lbadmin/****

Once logged in, the WUI is displayed:

- System Overview
- Local Configuration
- Cluster Configuration
- Maintenance
- View Configuration
- Reports
- Logs
- Support

SYSTEM OVERVIEW ⓘ

2015-06-18 14:21:20 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



System Load Average



Memory Usage



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

Adding Virtual Services (VIPs)

Layer 7 VIPs can be created using the wizard or manually. Layer 4 VIPs must be created manually.



NOTE : It's not possible to configure a VIP on the same IP address as any of the network interfaces. This ensures services can 'float' (move) between master and slave appliances.

Manually adding a layer 4 VIP

- Using the WUI, go to *Cluster Configuration > Layer 4 – Virtual Services* and click **[Add a New Virtual Service]**
- Enter the required details for the new VIP, example details are shown below:

Label	<input type="text" value="HTTP-Cluster"/>		?
Virtual Service	IP Address	<input type="text" value="192.168.1.100"/>	?
	Ports	<input type="text" value="80"/>	?
Protocol	<input type="text" value="TCP"/>		?
Forwarding Method	<input type="text" value="Direct Routing"/>		?

- Enter an appropriate label for the VIP, e.g. **HTTP-Cluster**
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.1.100**
- Set the *Virtual Service Ports* field to the required port, e.g **80**
- Leave *Protocol* set to **TCP**
- Set *Forwarding Method* as required, either **Direct Routing**, **NAT** or **Tunneling**
- Click **Update**

Manually adding a layer 7 VIP

- Using the WUI, go to *Cluster Configuration > Layer 7 – Virtual Services* and click **[Add a New Virtual Service]**
- Enter the required details for the new VIP, example details are shown below:

Label	<input type="text" value="HTTP-Cluster"/>		?
Virtual Service	IP Address	<input type="text" value="192.168.1.100"/>	?
	Ports	<input type="text" value="80"/>	?
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>		?
Manual Configuration	<input type="checkbox"/>		?






- Enter an appropriate label for the VIP, e.g. **HTTP-Cluster**
- Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.1.100**
- Set the *Virtual Service Ports* field to the required port, e.g **80**
- Set *Layer 7 Protocol* set to **HTTP Mode** for HTTP servers, and **TCP Mode** for non HTTP servers
- Click **Update**

Defining Real Servers (RIPs)

Layer 7 RIPs can be defined using the wizard or manually. Layer 4 RIPs must be defined manually.

Manually adding a layer 4 RIP

- Using the WUI, go to *Cluster Configuration > Layer 4 – Real Servers* and click **[Add a New Real Server]** next to the relevant VIP
- Enter the required details for the new RIP, example details are shown below:





Label	<input type="text" value="IIS1"/>	
Real Server IP Address	<input type="text" value="192.168.1.110"/>	
Weight	<input type="text" value="100"/>	
Minimum Connections	<input type="text" value="0"/>	
Maximum Connections	<input type="text" value="0"/>	

- Enter an appropriate label for the RIP, e.g. **IIS1**
- Set the *Real Server IP address* field to the required IP address, e.g. **192.168.1.110**
- Leave Weight set to **100**
- Leave *Minimum Connections* and *Maximum Connections* set to 0 (unrestricted)
- Click **Update**

N.B. For DR mode RIPs, there is no port field because port redirection is not possible

Manually adding a layer 7 RIP

- Using the WUI, go to *Cluster Configuration > Layer 7 – Real Servers* and click **[Add a New Real Server]** next to the relevant VIP
- Enter the required details for the new RIP, example details are shown below:

Label	<input type="text" value="IIS1"/>	
Real Server IP Address	<input type="text" value="192.168.1.110"/>	
Real Server Port	<input type="text" value="80"/>	
Weight	<input type="text" value="100"/>	

- Enter an appropriate label for the RIP, e.g. **IIS1**
- Set the *Real Service IP address* field to the required IP address, e.g. **192.168.1.110**
- Set the *Real Server Port* field to the required port, e.g. **80**
- Set *Layer 7 Protocol* to **HTTP Mode** for HTTP servers, and **TCP Mode** for non HTTP servers
- Click **Update**

Configuring the Physical Real Servers

Depending on the deployment method (DR, NAT or SNAT) used, the actual physical backend servers may need additional configuration to allow the load balancer to operate correctly. The following sections define what is needed for each mode.

Configuring the Real Servers for Layer 4 NAT Mode

If you are using a two-arm NAT load balancing method, the Real Server configuration is a simple case of configuring the load balancer as the default gateway. Normally, a floating IP address is added using *Cluster Configuration > Floating IPs*. This is important when a master / slave configuration is used to allow failover & failback of the default gateway address.



NOTE : Failure to correctly configure the Real Servers default gateway is the most common mistake when using NAT mode.

Configuring the Real Servers for Layer 4 DR Mode

When using DR mode, each Real Server (RIP) must be configured to accept packets destined for both the VIP address **and** the RIP address. This is because in DR mode the destination address of load balanced packets is the VIP address, whilst for other traffic such as health-checks, administration traffic etc. it's the Real Server's own IP address (the RIP). Each Real Server must also be configured so that it does not respond to ARP requests for the VIP address – only the load balancer should do this.

Configuring the Real Servers in this way is referred to as 'Solving the ARP problem'. The steps required depend on the OS used as detailed in the following sections.

Detecting the ARP Problem

Attempt to connect to the VIP, then use *Reports > Layer 4 Current Connections* to check whether the connection state is **SYN_RECV** as shown below. If it is, this is normally a good indication that the real servers have not been correctly configured for DR mode.

```
IPVS connection entries
pro expire state      source          virtual         destination
TCP 00:51 SYN_RECV    192.168.2.7:64763 192.168.2.109:80 192.168.2.99:80
```

Solving the ARP Problem for Linux

Method 1 (using iptables)

You can use iptables (netfilter) on each Real Server to re-direct incoming packets destined for the Virtual Service IP address. To make this permanent, simply add the command to an appropriate start-up script such as /etc/rc.local. If the Real Server is serving multiple VIPs, add additional iptables rules for each VIP.

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

```
e.g.. iptables -t nat -A PREROUTING -p tcp -d 10.0.0.21 -j REDIRECT
```

(Change the IP address to be the same as your Virtual Service)

This means redirect any incoming packets destined for 10.0.0.21 (the Virtual Service) locally, i.e. to the primary address of the incoming interface on the Real Server.



NOTE : Method 1 may not always be appropriate if you're using IP-based virtual hosting on your web server. This is because the iptables rule above redirects incoming packets to the primary address of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 2 below instead.

Also, Method 1 does not work with IPv6 Virtual Services, use method 2 below instead.

Method 2 (using arp_ignore sysctl values)

This is the preferred method as it supports both IPv4 and IPv6. Each Real Server needs the loopback adapter to be configured with the Virtual Services IP address. This address must not respond to ARP requests and the web server also needs to be configured to respond to this address. To set this up follow steps 1-4 below.

Step 1: re-configure ARP on the Real Servers (this step can be skipped for IPv6 Virtual Services)

To do this add the following lines to /etc/sysctl.conf:

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

Step 2: re-configure DAD on the Real Servers (this step can be skipped for IPv4 Virtual Services)

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

Step 3: apply these settings

Either reboot the Real Server or run the following command to apply these settings:

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

Step 4: add the Virtual Services IP address to the loopback adapter

Run the following command for each VIP. To make this permanent, simply add the command to an appropriate startup script such as /etc/rc.local.

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

for IPv6 addresses use:

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


*N.B. Steps 1, 2 & 3 can be replaced by writing directly to the required files using the following commands:
(temporary until the next reboot)*

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

Solving the ARP Problem for Solaris & MAC OS X / BSD

Solaris:

With Solaris the loopback interface does not respond to ARP requests so you just add your VIPs to it.

```
ifconfig lo0:1 plumb
ifconfig lo0:1 VIP netmask 255.255.255.255 up
```

You will need to add this to the startup scripts for your server.

MAC OS X or BSD:

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

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

You will need to add this to the startup scripts for your server.



NOTE : Don't forget that the service on the Real Servers needs to listen on both the RIP address and VIP address as mentioned previously.



NOTE : Failure to correctly configure the Real Servers to handle the ARP problem is the most common mistake in DR mode configurations.

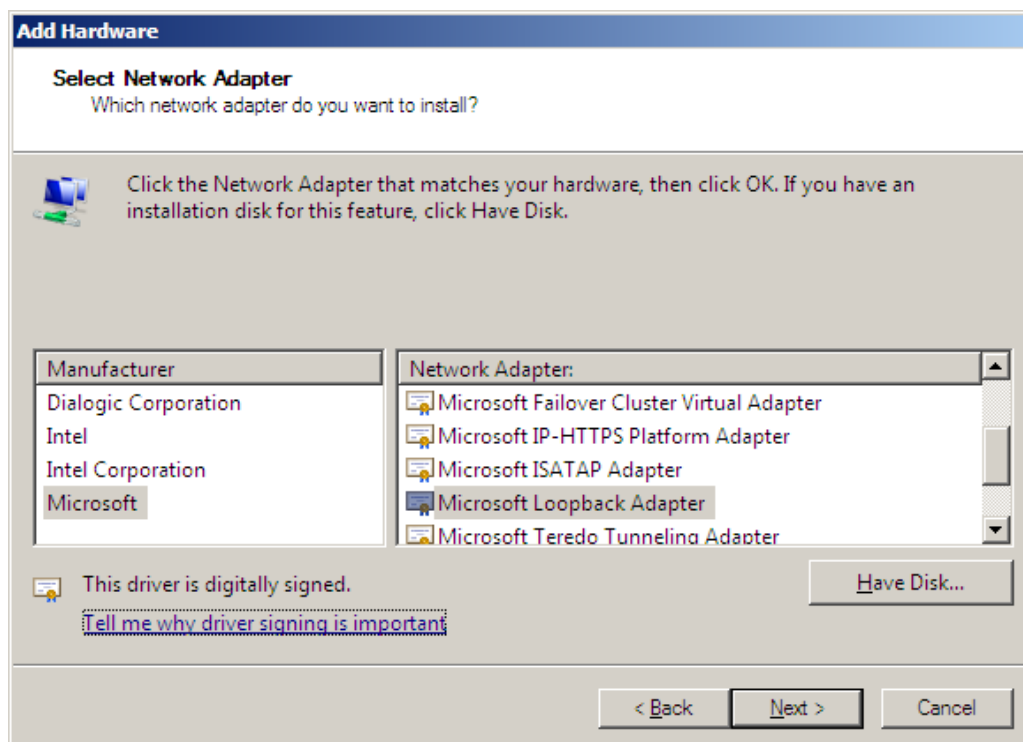
Solving the ARP Problem for Windows Servers

Windows Server 2008

Windows Server 2008 supports Direct Routing (DR) by using the MS Loopback Adapter to handle the traffic and a series of netsh commands to modify the servers strong / weak host behavior. The IP address on the Loopback Adapter must be set to be the same as the Virtual Services IP address (VIP). If the Real Server is included in multiple VIPs, you can add additional IP addresses to the Loopback Adapter that correspond to each VIP.

Step 1: Install the Microsoft Loopback Adapter

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




6. Click **Next** to start the installation, when complete click **Finish**

Step 2: Configure the Loopback Adapter


1. Open Control Panel and click **View Network status and tasks** under **Network and internet**
2. Click **Change adapter settings**
3. Right-click the new Loopback Adapter and select **Properties**

- Un-check all items except **Internet Protocol Version 4 (TCP/IPv4)** and **Internet Protocol Version 6 (TCP/IPv6)** as shown below


N.B. leaving both checked ensures that both IPv4 and IPv6 are supported. If preferred, only the protocol to be used can be checked



- If configuring IPv4 addresses select **Internet Protocol Version 4 (TCP/IPv4)**, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) with a subnet mask of 255.255.255.255 , e.g. 192.168.2.20 / 255.255.255.255 as shown below



- If configuring IPv6 addresses select **Internet Protocol Version 6 (TCP/IPv6)**, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) and set the *Subnet Prefix Length* to be the same as your network setting , e.g. 2001:470:1f09:e72::15 / 64 as shown below



7. Click **OK**, then click **Close** to save and apply the new settings
8. Now repeat the above process on the other Windows 2008 Real Servers

N.B. For Windows 2008, it's not necessary to modify the interface metric on the advanced tab and should be left set to Automatic

Step 3: Configure the strong / weak host behavior

Windows 2008 supports strong host sends and receives for both IPv4 and IPv6 by default. To ensure that Windows 2008 is running in the correct mode to be able to respond to the VIP, the following commands must be run on each Real Server:

For IPv4 addresses :

```
netsh interface ipv4 set interface "net" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostsend=enabled
```

For these commands to work, the LAN connection NIC must be named "net" and the loopback NIC must be named "loopback" as shown below. If you prefer to leave your current NIC names, then the commands above must be modified accordingly. For example, if your network adapters are named "LAN" and "LOOPBACK", the commands required would be:


```
netsh interface ipv4 set interface "LAN" weakhostreceive=enabled
netsh interface ipv4 set interface "LOOPBACK" weakhostreceive=enabled
netsh interface ipv4 set interface "LOOPBACK" weakhostsend=enabled
```


For IPv6 addresses :

```
netsh interface ipv6 set interface "net" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostsend=enabled
netsh interface ipv6 set interface "loopback" dadtransmits=0
```


For these commands to work, the LAN connection NIC must be named "net" and the loopback NIC must be named "loopback" as shown below. If you prefer to leave your current NIC names, then the commands above must be modified accordingly. For example, if your network adapters are named "LAN" and "LOOPBACK", the commands required would be:

```
netsh interface ipv6 set interface "LAN" weakhostreceive=enabled
netsh interface ipv6 set interface "LOOPBACK" weakhostreceive=enabled
netsh interface ipv6 set interface "LOOPBACK" weakhostsend=enabled
netsh interface ipv6 set interface "LOOPBACK" dadtransmits=0
```



N.B. The names for the NICs are case sensitive, so make sure that the name used for the interface and the name used in the commands match exactly.

1. Start Powershell or use a command window to run the appropriate netsh commands as shown in the example below



N.B. This shows an IPv6 example, use the IPv4 commands if you're using IPv4 addresses


2. Now repeat these 4 commands on the other Windows 2008 Real Servers

Windows Server 2012

Windows Server 2012 supports Direct Routing (DR) by using the MS Loopback Adapter to handle the traffic and a series of netsh commands to modify the servers strong / weak host behavior. The IP address on the Loopback Adapter must be set to be the same as the Virtual Services IP address (VIP). If the Real Server is included in multiple VIPs, you can add additional IP addresses to the Loopback Adapter that correspond to each VIP.

Step 1: Install the Microsoft Loopback Adapter

1. Click **Start**, then run **hdwwiz** to start the Hardware Installation Wizard
2. When the Wizard has started, click **Next**
3. Select **Install the hardware that I manually select from a list (Advanced)**, click **Next**
4. Select **Network adapters**, click **Next**
5. Select **Microsoft & Microsoft KM-Test Loopback Adapter**, click **Next**




6. Click **Next** to start the installation, when complete click **Finish**

Step 2: Configure the Loopback Adapter


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

- Un-check all items except **Internet Protocol Version 4 (TCP/IPv4)** and **Internet Protocol Version 6 (TCP/IPv6)** as shown below


N.B. leaving both checked ensures that both IPv4 and IPv6 are supported. If preferred, only the protocol to be used can be checked



- If configuring IPv4 addresses select **Internet Protocol Version 4 (TCP/IPv4)**, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) with a subnet mask of 255.255.255.255 , e.g. 192.168.2.20 / 255.255.255.255 as shown below



- If configuring IPv6 addresses select **Internet Protocol Version 6 (TCP/IPv6)**, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) and set the *Subnet Prefix Length* to be the same as your network setting , e.g. 2001:470:1f09:e72::15 / 64 as shown below



7. Click **OK** on TCP/IP Properties, then click **Close** on Ethernet Properties to save and apply the new settings
8. Now repeat the above process on the other Windows 2012 Real Servers

N.B. For Windows 2012, it's not necessary to modify the interface metric on the advanced tab and should be left set to Automatic

Step 3: Configure the strong / weak host behavior

Windows 2012 supports strong host sends and receives for both IPv4 and IPv6 by default. To ensure that Windows 2008 is running in the correct mode to be able to respond to the VIP, the following commands must be run on each Real Server:

For IPv4 addresses :

```
netsh interface ipv4 set interface "net" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostsend=enabled
```

For these commands to work, the LAN connection NIC must be named "net" and the loopback NIC must be named "loopback" as shown below. If you prefer to leave your current NIC names, then the commands above must be modified accordingly. For example, if your network adapters are named "LAN" and "LOOPBACK", the commands required would be:

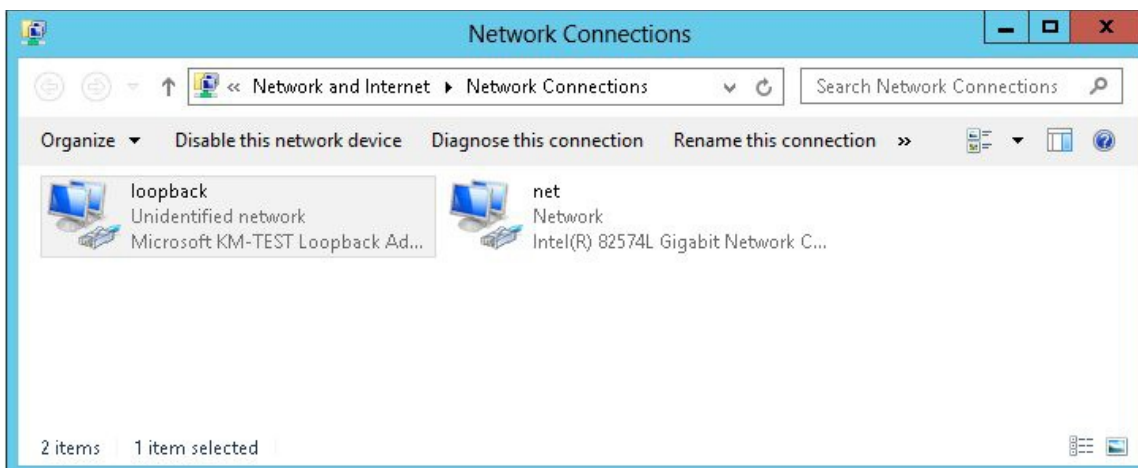
```
netsh interface ipv4 set interface "LAN" weakhostreceive=enabled
netsh interface ipv4 set interface "LOOPBACK" weakhostreceive=enabled
netsh interface ipv4 set interface "LOOPBACK" weakhostsend=enabled
```


For IPv6 addresses :

```
netsh interface ipv6 set interface "net" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostsend=enabled
netsh interface ipv6 set interface "loopback" dadtransmits=0
```


For these commands to work, the LAN connection NIC must be named "net" and the loopback NIC must be named "loopback" as shown below. If you prefer to leave your current NIC names, then the commands above must be modified accordingly. For example, if your network adapters are named "LAN" and "LOOPBACK", the commands required would be:

```
netsh interface ipv6 set interface "LAN" weakhostreceive=enabled
netsh interface ipv6 set interface "LOOPBACK" weakhostreceive=enabled
netsh interface ipv6 set interface "LOOPBACK" weakhostsend=enabled
netsh interface ipv6 set interface "LOOPBACK" dadtransmits=0
```



N.B. The names for the NICs are case sensitive, so make sure that the name used for the interface and the name used in the commands match exactly.

1. Start Powershell or use a command window to run the appropriate netsh commands as shown in the example below



N.B. This shows an IPv6 example, use the IPv4 commands if you're using IPv4 addresses

2. Now repeat these 4 commands on the other Windows 2012 Real Servers

Verifying netsh Settings for Windows 2008 & 2012

To verify that settings have been configured correctly, run the following command on each Real Server to clearly list the settings that have been applied to the interface:

```
netsh interface ipv4 show interface <interface name>
```

i.e.

for the 'loopback' adapter run: netsh interface ipv4 show interface loopback

for the 'net' adapter run: netsh interface ipv4 show interface net

(N.B. For IPv6, simply replace 'ipv4' with 'ipv6' in the above commands)

e.g.

```
C:\Users\Administrator>netsh interface ipv4 show interface loopback
```

```
Interface loopback Parameters
```

```
-----
IfLuid                      : ethernet_9
IfIndex                     : 15
State                       : connected
Metric                     : 30
Link MTU                    : 1500 bytes
Reachable Time              : 28500 ms
Base Reachable Time         : 30000 ms
Retransmission Interval     : 1000 ms
DAD Transmits               : 3
Site Prefix Length          : 64
Site Id                     : 1
Forwarding                  : disabled
Advertising                 : disabled
Neighbor Discovery          : enabled
Neighbor Unreachability Detection : enabled
Router Discovery            : dhcp
Managed Address Configuration : enabled
Other Stateful Configuration : enabled
Weak Host Sends              : enabled
Weak Host Receives          : enabled
Use Automatic Metric        : enabled
Ignore Default Routes       : disabled
Advertised Router Lifetime  : 1800 seconds
Advertise Default Route     : disabled
Current Hop Limit           : 0
Force ARPND wake up patterns : disabled
Directed MAC wake up patterns : disabled
```

```
C:\Users\Administrator>
```

This shows that the settings have been applied correctly.



NOTE : For Windows server 2008 / 2012, if you want to leave the built-in firewall enabled, you'll either need to enable the relevant default firewall exceptions or create your own to enable access to the web server. By default these exceptions will allow traffic on both the network and loopback adapters.



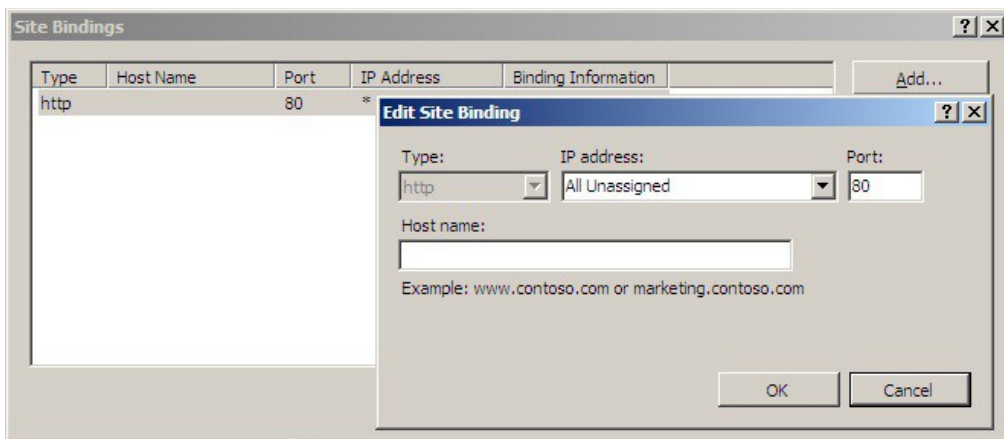
NOTE : Failure to correctly configure the Real Servers to handle the ARP problem is the most common problem in DR configurations.

Configuring the Application/Service to Respond to both the RIP and VIP

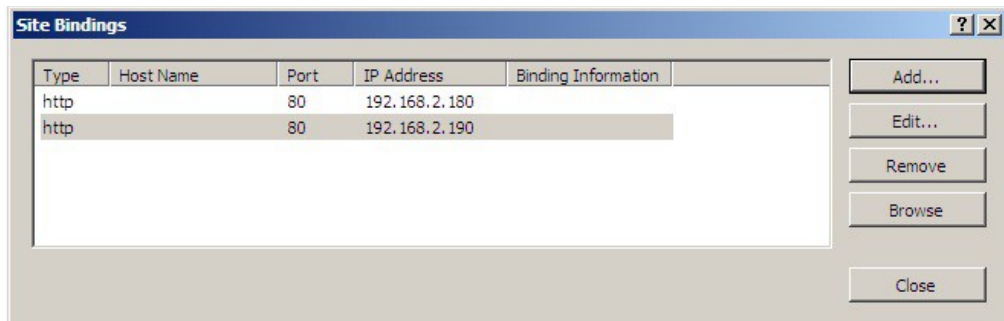
For DR & TUN modes, it's also important to make sure that the application running on the Real Servers (e.g. IIS) responds to both the VIP and RIP addresses.

IIS Example - Windows 2008 / 2012

By default, IIS listens on all configured IP addresses, this is shown in the example below (shows Windows 2008 example). As can be seen the IP address field is set to "All Unassigned".



If the default configuration is left, no further IIS configuration is required. If you do change the IP address in the bindings from "All Unassigned" to a specific IP address, then you need to make sure that you also add a binding for the Virtual Service IP address (VIP) as shown in the example below:



NOTE : These examples illustrates how IIS must be configured to ensure that its listening on both the RIP and VIP address. It's important to remember that this applies equally to all applications when running in DR mode.

Configuring the Real Servers for Layer 7 SNAT Mode

When using Layer7 (HAProxy) Virtual Services, no changes are required to the Real Servers.

High-Availability Configuration of two Loadbalancer.org Appliances


Loadbalancer.org's recommended configuration is to use a clustered pair of load balancers to provide a highly available and resilient load balancing solution. In this configuration, the pair uses a heartbeat to determine the state of the other appliance. Should the active device (normally the master) suffer a failure, the passive device (normally the slave) will take over the resources hosted on the shared floating IP addresses.

Clustered Pair Configuration

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

- In the WUI, open: *Cluster Configuration > High-Availability Configuration*

CREATE A CLUSTERED PAIR

 loadbalancer.org

Local IP address
192.168.110.40


IP address of new peer
192.168.110.41


Password for *loadbalancer* user on peer
••••••••

Add new node

- 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:

CREATE A CLUSTERED PAIR

 192.168.110.40 loadbalancer.org

 192.168.110.41 loadbalancer.org

Creating pool..

Local IP address
192.168.110.40

IP address of new peer
192.168.110.41

Password for *loadbalancer* user on peer
••••••••

configuring

- Once complete, the following will be displayed:



- To finalize the configuration, restart heartbeat as prompted in the blue message box

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



NOTE : For more details on configuring an HA pair please refer to chapter 9 in the administration manual.

Testing Load Balanced Services

For example, to test a web server based configuration, add a page to each web servers root directory e.g. test.html and put the server name on this page for easy identification during the tests.

Use two or more clients to do the testing. Open up a web browser on each test clients and enter the URL for the VIP e.g. **http://192.168.110.150**

Each client should see a different server name because of the load balancing algorithm in use i.e. they are being load balanced across the cluster.

Why test using two clients? If you use a single client it will most likely keep on hitting the same server for multiple requests. This is to do with the way that the load balancing algorithms are optimized.

Diagnosing VIP Connection Problems

1. **Make sure that the device is active** - this can be checked in the WUI. For a single appliance, the status bar should report **Master & Active** as shown below:

Master | Slave Active | Passive Link

2. **Check that the VIP/floating IP is up** - Using *View Configuration > Network Configuration* verify that the VIP is active on the load balancer, if not check *Logs > Heartbeat* for errors.

```
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 00:0c:29:cf:18:03 brd ff:ff:ff:ff:ff:ff
    inet 192.168.110.85/18 brd 192.168.127.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet 192.168.110.90/18 brd 192.168.127.255 scope global secondary eth0
        valid_lft forever preferred_lft forever
```

The above example shows that the interface address (192.168.110.85) and the VIP address (192.168.110.90) are both up.

3. **Check that the Real Servers are up** - Using *System Overview* make sure that none of your VIPs are colored red. If they are, the entire cluster is down (i.e. all Real Servers). Green indicates a healthy cluster, yellow indicates that your cluster may need attention (one or more of the Real Servers may be down), and blue indicates all Real Server have been deliberately taken offline (by using either Halt or Drain).

SYSTEM OVERVIEW ?							2015-03-18 11:37:15 UTC
	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
↑	HTTP-Cluster	192.168.110.150	80	0	TCP	Layer 4	DR
⚠	RDP-Cluster	192.168.110.150	3389	0	TCP	Layer 4	DR
↓	HTTP-Cluster-2	192.168.110.152	80	0	HTTP	Layer 7	Proxy
⚙	RDP-Cluster-2	192.168.110.152	3389	0	TCP	Layer 7	Proxy

4. Check the connection state -

For Layer 4 DR mode VIPs check *Reports > Layer 4 Current Connections* to view the current traffic in detail. Any packets with state **SYN_RECV** imply that the 'ARP Problem' has not been correctly solved on the Real Servers. See page 23 for more details on solving the ARP problem.

For layer 4 NAT mode VIPs check *Reports > Layer 4 Current Connections* to view the current traffic in detail. Any packets with state **SYN_RECV** often imply that the default gateway on the Real Servers has not been set to be an IP address on the load balancer.

For Layer 7 VIPs check *Reports > Layer 7 Status*. The default credentials required are:

username: loadbalancer

password: loadbalancer

This will open a second tab in the browser and display a statistics/status report as shown in the example below:

Statistics Report for pid 3261

> General process information

pid = 3261 (process #1, nbproc = 1)

uptime = 0d 0h00m42s

system limits: memmax = unlimited; ulimit-n = 81000

maxsock = 80024; maxconn = 40000; maxpipes = 0

current conns = 1; current pipes = 0/0; conn rate = 2/sec

Running tasks: 1/5; idle = 100 %

active UP

active UP, going down

active DOWN, going up

active or backup DOWN

active or backup DOWN for maintenance (MAINT)

backup UP

backup UP, going down

backup DOWN, going up

not checked

Display option:

Hide 'DOWN' servers

Refresh now

CSV export

External resources:

Primary site

Updates (v1.5)

Online manual

Note: UP with load-balancing disabled is reported as "NOLB".

L7

	Queue			Session rate			Sessions					Bytes		Denied		Errors			Warnings		Server									
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntime	Thrt	
Frontend				0	15	-	0	4	40 000	56		21 696	3 385 782	0	0	0					OPEN									
backup	0	0	-	0	0		0	0	-	0	0	0	0	0	0	0	0	0	0	0			1	-	Y					-
RIP1	0	0	-	0	16		0	2	-	56	56	21 696	3 385 782	0	0	0	0	0	0	0	42s UP	L4OK in 0ms	1	Y	-	0	0	0s	-	
Backend	0	0		0	16		0	2	4 000	56	56	21 696	3 385 782	0	0	0	0	0	0	0	42s UP		1	1	1			0	0s	-

stats

	Queue			Session rate			Sessions					Bytes		Denied		Errors			Warnings		Server									
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntime	Thrt	
Frontend				2	4	-	1	1	2 000	8		1 464	33 111	0	0	4					OPEN									
Backend	0	0		0	0		0	0	200	0	0	1 464	33 111	0	0	0	0	0	0	0	42s UP		0	0	0			0		









Taking Real Servers Offline (Halting)

Using the *System Overview* check that when you Halt one of the Real Servers the connections are redirected to the other server in the cluster.

Remove the network cable from one of the web servers or stop the web service/process, wait a few seconds (for the load balancer to detect the change) and then refresh the browsers on both clients. They should now both switch to the same server (since one has been removed from the load balancing list).

Replace the network cable, wait a few seconds and then refresh the browsers again. After a few refreshes they should again show different web servers.

The *System Overview* will also show the updated status as these tests are performed:

VIRTUAL SERVICE ⚙		IP ⚙	PORTS ⚙	CONNS ⚙	PROTOCOL ⚙	METHOD ⚙	MODE ⚙	
<div><div></div><div></div><div></div><div></div><div></div></div>	 HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	 RIP1	192.168.110.240	80	100	0	Drain	Halt	
	 RIP2	192.168.110.241	80	0	0	Online (halt)		
	 RIP3	192.168.110.242	80	100	0	Drain	Halt	

In this example:

'rip1' is green, this indicates that it's operating normally.

'rip2' is blue, this indicates that it has been either Halted or Drained. in this example Halt has been used as indicated by *Online (Halt)* being displayed. If it had been drained it would show as *Online (Drain)*.

'rip3' is red, this indicates that it has failed a health check.



NOTE : From v7.6.4 the System Overview supports sorting of VIPs. This can be done by clicking on the column headings or by drag & drop. For more details please refer to chapter 8 in the administration manual.

Using Reports & Log Files

The appliance includes several logs and reports that are very useful when diagnosing issues. Both are available as main menu options in the WUI. Details of both can be found in chapter 13 of the administration manual.

Testing High-Availability for a Loadbalancer.org HA-Pair

To test fail-over of a clustered pair, once fully configured power off the master and check that the slave unit takes over all the floating IP(s). If fail-over to the slave unit does not occur correctly, check *Logs > Heartbeat* on both nodes for any errors.



NOTE : It's very important to verify that master / slave failover occurs correctly before going live. This proves the resilience of the cluster and makes you aware of the failover / failback process. Please refer to chapter 9 in the administration manual for more details.



NOTE : When testing load balancer fail-over, don't just pull the network cable out. This will not cause a fail-over but will cause a split brain (i.e. both units active) to occur. It is possible to configure fail-over on network failure but this is not enabled by default. To enable this, a ping node must be configured under *Cluster Configuration > Heartbeat Configuration*.

The status of the appliance is shown at the top of the screen. For a working pair, the normal view is shown below:

This shows that the master unit is active and that the heartbeat link is up between master & slave.

N.B. If no VIPs are defined, the status on master & slave appears as follows:

Other states:

Master Slave	Active Passive	Link	this is a master unit, it's active, no slave unit has been defined
Master Slave	Active Passive	Link	this is a master unit, it's active, a slave has been defined but the link to the slave is down. Action: check & verify the heartbeat configuration
Master Slave	Active Passive	Link	this is a slave unit, it's active (a failover from the master has occurred) and the heartbeat link to the master has been established
Master Slave	Active Passive	Link	this is a master unit, a slave unit has been defined, but the link is down (e.g. serial cable unplugged) so the state cannot be determined. In this case the floating IP's may be active on both units. Action: check & verify the heartbeat configuration, check the serial cable (if applicable), check heartbeat logs & if required restart heartbeat on both units

N.B. Restarting heartbeat will cause a temporary outage of all load balanced services

Does Your Application Cluster Correctly Handle its Own State?



NOTE : Load balancers work most effectively if the application servers are completely stateless. This means that if a web server fails and is automatically taken out of the cluster; then all the current user sessions will be transferred to other servers in the cluster without the users needing to re login to the application again. ***If your application doesn't have a persistent data store then you can't have seamless fail over for your back-end servers.***

Do your web servers store persistent information on local drives?

- Images (jpeg, png, gif etc.)
- Files (html, php, asp etc.)

If so, these files either need to be on shared storage such as an NFS/CIFS mount, or they need to be replicated to all of the nodes in the cluster.

Replication Solutions for Shared Data

On UNIX you can use the RSYNC command to replicate files, on Windows Server you can use RSYNC as well but you may prefer ROBOCOPY that's included by default in newer versions of Windows Server or in the resource kit for older versions. Usually you will upload your content to one master server and then replicate it to the other servers in the cluster.

Solutions for Session Data

Standard ASP and PHP session data is stored locally by default, leaving your session data in a local store will prevent you from implementing seamless application server fail-over in your cluster. If an application server fails, all of the local session data will be lost and your user will need to re-log in and possibly lose shopping baskets etc.

This problem is easily resolvable by implementing a shared persistent data store for the cluster. This is usually either done with a shared back-end database or a shared memory solution.

Persistence (aka Affinity)

Persistence is a feature that is required by many web applications. Once a user has interacted with a particular server all subsequent requests are sent to the same server thus persisting to that particular server. It is normally required when the session state is stored locally to the web server as opposed to a database.

What do You do if Your Application is not Stateless?

Some applications require state to be maintained such as:

- Terminal Services / Remote Desktop Services
- SSH
- FTP (upload)
- SMTP (incoming)

You may also find that you are unable to modify your HTTP/HTTPS based application to handle shared session data.

For these cases, you can use persistence based on source IP address. You lose the ability to have transparent fail-over, but you do still get increased capacity and manageability. This persistence problem occurs with all load balancers and all vendors use standard methods and technologies to mitigate the issue.

Loadbalancer.org Persistence Options

- Source IP (subnet)
- Cookie (Active or Passive)
- SSL session ID
- Microsoft Connection Broker / Session Broker Integration

The standard Layer 4 persistence method is source IP persistence, you can handle millions of persistent connections at Layer 4. Just modify your Virtual Service to be persistent if you require source IP persistence.

Cookies are a Layer 7 based persistence method that can offer more even traffic distribution and also handle any clients where the source IP address may change during the session (e.g. mega proxies).

SSL session ID based persistence is useful in certain circumstances, although due to the way some browsers operate – notably Internet Explorer, the session ID can be renegotiated frequently which effectively breaks the persistence.

Loadbalancer.org Technical Support

If you have any questions regarding the appliance don't hesitate to contact the support team via their email address: support@loadbalancer.org or your local reseller.

For more details please refer to the administration manual:

<http://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf>


Appendix

Company Contact Information


Website	URL : www.loadbalancer.org
North America (US)	<p>Loadbalancer.org, Inc. 270 Presidential Drive Wilmington, DE 19807 USA</p> <p>Tel : +1 888.867.9504 (24x7) 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 : +1 855.681.6017 (24x7) Fax : +1 302.213.0122 Email (sales) : sales@loadbalancer.org Email (support) : support@loadbalancer.org</p>
Europe (UK)	<p>Loadbalancer.org Ltd. Portsmouth Technopole Kingston Crescent Portsmouth PO2 8FA England, UK</p> <p>Tel : +44 (0)330 3801064 (24x7) Fax : +44 (0)870 4327672 Email (sales) : sales@loadbalancer.org Email (support) : support@loadbalancer.org</p>
Europe (Germany)	<p>Loadbalancer.org GmbH Alt Pempelfort 2 40211 Düsseldorf Germany</p> <p>Tel : +49 (0)30 920 383 6494 Fax : +49 (0)30 920 383 6495 Email (sales) : vertrieb@loadbalancer.org Email (support) : support@loadbalancer.org</p>

Front & Rear Panel Layouts

Enterprise R20 & Enterprise Max



Enterprise 10G & Enterprise R320




Physical Appliance (back page reference)

- Remove all packaging
- Rack mount the appliance if required
- The power supply is an auto sensing unit (100v to 240v)
- Connect the power lead from the power socket to the mains or UPS
- Connect a network cable from the switch to one of the Ethernet ports – typically *eth0* but this is not mandatory
- If using a two-armed configuration connect another cable to a second Ethernet port – typically *eth1* but this is not mandatory
- For a clustered hardware pair (see page 36), the units must be able to communicate either via network (ucast), via serial cable or both. By default, ucast only is used. If serial is preferred or you want to use both methods, connect a serial cable (1 supplied with each appliance) between the two appliances.

N.B. If a serial cable is used, Heartbeat must be configured for this using the WUI option: Cluster Configuration > Heartbeat Configuration and enabling 'Serial'

- Attach a monitor to the VGA port and keyboard to the USB or PS/2 port
- Check mains power is on and press the power switch to start the appliance (the fans should start & front panel LED's should light)
- Allow a minute for booting
- Refer to page 14 for details on configuring network settings
- Refer to page 17 for details on accessing the Web User Interface



N.B. The above image shows the Enterprise MAX, for connecting other models please refer to page 47