



Load Balancing Microsoft Exchange 2016

Version 1.4.2

Table of Contents

1. About this Guide	4
2. Loadbalancer.org Appliances Supported	4
3. Loadbalancer.org Software Versions Supported	4
4. Microsoft Exchange Software Versions Supported	4
5. Exchange Server 2016	4
6. Exchange 2016 Server Roles	4
7. Load Balancing Exchange 2016	5
Load Balancing & HA Requirements	5
Database Availability Group (DAG)	5
Persistence (aka Server Affinity)	5
Port Requirements	6
SSL Termination	6
HTTPS Namespaces & IP addresses	6
Health-Checks	7
Load Balancer Deployment Concept	7
Virtual Service (VIP) Requirements	8
Load Balancer Deployment Modes	8
Layer 7 SNAT Mode	8
Layer 4 DR Mode	9
Our Recommendation	10
Is SSL Offloading Required?	10
8. Configuring Exchange 2016 for Load Balancing	11
1) External Access Domain	11
2) Virtual Directories	12
3) Outlook Anywhere	12
4) Autodiscover	12
Internal	12
External	13
5) Certificates	14
6) Send & Receive Connectors	14
Adding Connectors	15
7) DNS Configuration	15
8) Additional Exchange Server Configuration Steps (depends on Load balancing method)	15
SNAT Mode	15
DR Mode	15
9) IIS Restart (Important)	15
9. Loadbalancer.org Appliance – the Basics	15
Virtual Appliance	15
Initial Network Configuration	16
Accessing the WebUI	16
Main Menu Options	17
HA Clustered Pair Configuration	18
10. Appliance Configuration – Using Layer 7 SNAT Mode (without SSL Offload)	18
Load Balancer Deployment Overview	18
Load Balancer Configuration	19
Configure VIP1 – Mailbox Server Role HTTPS Services	19
Configure VIP2 – Mailbox Server Role IMAP4/POP3 Services	20
Configure VIP3 – Mailbox Server Role SMTP Services	22
Configuring Firewall Rules to Lockdown SMTP	23

Additional Settings if using Kerberos Authentication	23
Finalizing the Configuration	23
Exchange Server Configuration Steps	24
11. Appliance Configuration – Using Layer 7 SNAT Mode (with SSL Offload)	24
Load Balancer Deployment Overview	24
Load Balancer Configuration	25
Configure VIP1 – Mailbox Server Role HTTP/HTTPS Services	25
Configure VIP2 – Mailbox Server Role IMAP4/POP3 Services	28
Configure VIP3 – Mailbox Server Role SMTP Services	29
Configuring Firewall Rules to Lockdown SMTP	30
Additional Settings if using Kerberos Authentication	30
Finalizing the Configuration	31
Exchange Server Configuration Steps	31
Configure IIS logging to Capture XFF Header IP Addresses	31
12. Appliance Configuration – Using Layer 4 DR Mode	31
Load Balancer Deployment Overview	31
Load Balancer Configuration	32
Configure VIP1 – Mailbox Server Role HTTPS Services	32
Configure VIP2 – Mailbox Server Role IMAP4/POP3 Services	33
Configure VIP3 – Mailbox Server Role SMTP Services	35
Exchange Server Configuration Steps	36
13. Testing & Verification	36
Useful Exchange 2016 & Other Microsoft Tools	36
Testing Server Health-checks using Set-ServerComponentState	36
Testing Mailflow	37
Testing SMTP Mail flow using Telnet	38
Microsoft Exchange Testing Tool	39
Useful Appliance based Tools & Features	39
Using System Overview	39
Layer 4 Status Report	40
Layer 7 Statistics Report	40
Appliance Logs	40
14. Technical Support	40
15. Further Documentation	41
16. Conclusion	41
17. Appendix	42
Configuring Firewall Rules to Lockdown SMTP	42
To add firewall rules	42
Enabling Layer 7 Transparency using TPROXY	43
Using a Layer 4 Virtual Service for SMTP	44
Layer 4 DR Mode – Solving the ARP Problem:	44
Configuring an HTTP to HTTPS redirect for OWA	44
Configuring HA - Adding a Secondary Appliance	45
Solving the ARP Problem	47
Windows Server 2012 & Later	47
Update the Network Adapter Priority Order	52
18. Document Revision History	54

1. About this Guide

This guide details the steps required to configure a load balanced Microsoft Exchange 2016 environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Microsoft Exchange 2016 configuration changes that are required to enable load balancing.

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

2. Loadbalancer.org Appliances Supported

All our products can be used with Exchange 2016. For full specifications of available models please refer to: <https://www.loadbalancer.org/products>.

Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

3. Loadbalancer.org Software Versions Supported

- V8.3.8 and later

Note

The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If using an older software version, note that the screenshots presented here may not match the WebUI exactly.

4. Microsoft Exchange Software Versions Supported

- Microsoft Exchange 2016 – all versions

5. Exchange Server 2016

Exchange 2016 is Microsoft's latest enterprise level messaging and collaboration server. Exchange 2016 has been designed for simplicity of scale, hardware utilization, and failure isolation. This has greatly simplified both the deployment process and the implementation of a load balancer.

Note

Exchange 2016 has since been superseded by Exchange 2019. The deployment guide for Exchange 2019 is available [here](#).

6. Exchange 2016 Server Roles

In Exchange 2016 the functionality of the Exchange 2013 CAS and Mailbox server roles have been consolidated into a single role: the **Mailbox Server Role**. In addition, the **Edge Transport Role** is also included.

Role	Purpose
Mailbox Server	This role consolidates the Mailbox and Client Access roles from Exchange Server 2013. Compared to Exchange Server 2010 this role consolidates all of the functions of the Client Access, Mailbox, Hub Transport, and Unified Messaging server roles. The Mailbox server role in Exchange Server 2016 is the only mandatory server role, and the consolidation reinforces the recommended practice since Exchange Server 2010 to deploy Exchange as a multi-role server instead of deploying individual roles to separate servers.
Edge Transport Server	This role is much the same as Edge Transport in previous versions of Exchange. It's designed to sit in perimeter networks and provide secure inbound and outbound mail flow for the organization. Edge Transport servers are not mandatory.

Outlook Client Protocols

- MAPI over HTTPS – *Outlook 2013 SP1 minimum*
- RPC over HTTPS – *aka Outlook Anywhere*

Mail Flow

In Exchange Server 2016, mail flow occurs through the transport pipeline. The transport pipeline is a collection of services, connections, components, and queues that work together to route all messages to the categorizer in the Transport service on an Exchange 2016 Mailbox server. For more information please refer to the the following Microsoft link: <https://technet.microsoft.com/en-us/library/aa996349%28v=exchg.160%29.aspx>

7. Load Balancing Exchange 2016

Note

It's highly recommended that you have a working Exchange 2016 environment first before implementing the load balancer.

Load Balancing & HA Requirements

In Exchange Server 2016, there is a single building block that provides the client access services and the high availability architecture necessary for any enterprise messaging environment. High availability is provided by implementing multiple Mailbox Servers, configuring a Database Availability Group (DAG) and deploying a load balancer.

Database Availability Group (DAG)

A DAG is a group of up to 16 Mailbox Servers with 100 active and passive databases. It provides automatic database-level recovery from failures that affect individual servers or databases.

Note

DAG's utilize Microsoft Clustering Services which cannot be enabled on the same server as Microsoft Network Load Balancing (NLB). Therefore, using Microsoft NLB is not an option in this case. Using a Loadbalancer.org hardware or virtual appliance provides an ideal solution.

Persistence (aka Server Affinity)

As with Exchange 2013, Exchange 2016 does not require session affinity at the load balancing layer.

Port Requirements

The following table shows the port list that must be load balanced. Some services such as IMAP4 or POP3 may not be required in your environment.

TCP Port	Role	Uses
25	MBOX	Inbound SMTP
110	MBOX	POP3 clients
143	MBOX	IMAP4 clients
443	MBOX	HTTPS (Outlook Web App, AutoDiscovery, Web Services, ActiveSync, MAPI over HTTP, RPC over HTTP – a.k.a. Outlook Anywhere, Offline Address Book, Exchange Administration Center) <i>Note: Outlook Web App has been renamed as Outlook on the Web in Exchange 2016</i>
993	MBOX	Secure IMAP4 clients
995	MBOX	Secure POP3 clients

SSL Termination

We generally recommend that SSL is terminated on the Exchange servers for scalability and effective load sharing. However, if you're load balancing Exchange using layer 7 SNAT mode, by default, the client IP address will be lost and replaced by the load balancer's own IP and therefore audit logs will contain the load balancer's IP address and not the clients. If this is an issue for your environment, X-Forwarded-For headers can be inserted by the load balancer which enable IIS on each Exchange server to be configured to log the client address from the XFF header as described in [this Microsoft article](#). In this case, SSL must be terminated on the load balancer to allow the header to be inserted. Once inserted, traffic can be re-encrypted from the load balancer to the Exchange servers. For more details on configuring layer 7 SNAT mode with SSL offload, please refer to [Appliance Configuration – Using Layer 7 SNAT Mode \(with SSL Offload\)](#).

HTTPS Namespaces & IP addresses

The following examples show 2 different approaches to HTTPS namespace configuration and the related load balancing considerations for each.

Example 1 – simple namespace configuration

Namespace	Purpose
mail.lbtestdom.com	Outlook Web App, ActiveSync, MAPI over HTTP, RPC over HTTP, Offline Address Book, Exchange Web Services
autodiscover.lbtestdom.com	Auto Discover

Notes

1. In this case a single VIP is used for all HTTPS namespaces/services.

- Both DNS entries should then point at the same VIP.
- This method is simple to setup, but only permits a single Exchange URL to be health checked. However, a successful full HTTPS service check on the OWA virtual directory is a good indication that the other Virtual Directories & applications are also functioning correctly.

Example 2 – expanded namespace configuration

Namespace	Purpose
owa.lbtestdom.com	Outlook Web Access
outlook.lbtestdom.com	Outlook Anywhere
ews.lbtestdom.com	Exchange Web Services
autodiscover.lbtestdom.com	Autodiscover
activesync.lbtestdom.com	ActiveSync
oab.lbtestdom.com	Offline Address Book

Notes

- In this case multiple VIPs are used – one for each HTTPS namespace/service.
- Each related DNS entry should then point at the corresponding VIP.
- This method is more complex to setup, but does enable more granular health checks to be configured.
- This guide uses the config of example 1 above, i.e. a single IP address for all services.

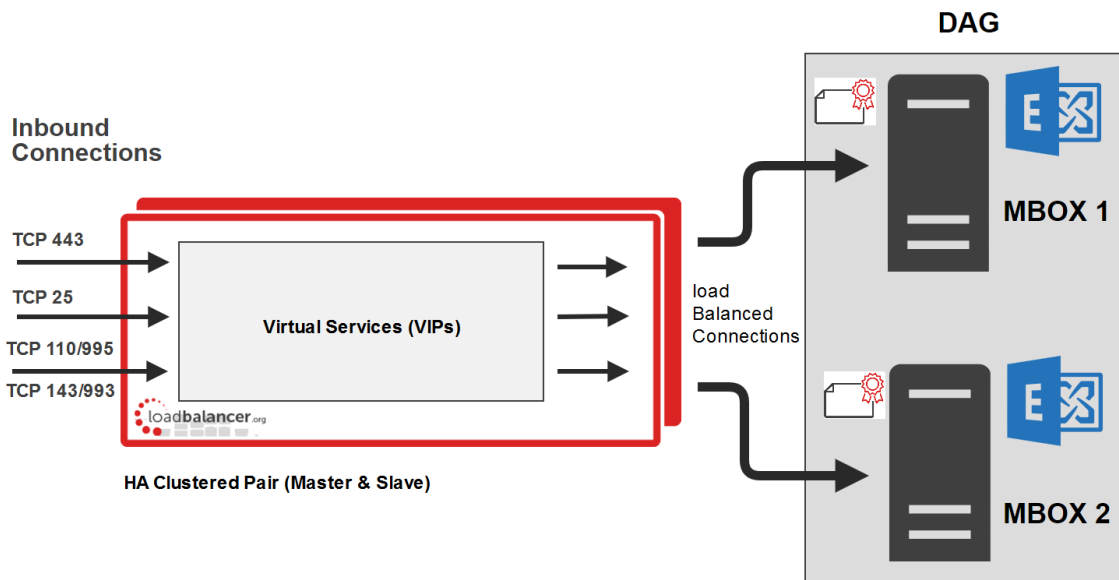
Health-Checks

In this guide, the health check for HTTPS services accesses `owa/healthcheck.htm` on each server and checks for a '200 OK' response. A different virtual directory (e.g. ECP, EWS etc.) can be chosen if preferred or more appropriate. Note that `healthcheck.htm` is generated in-memory based on the component state of the protocol in question and does not physically exist on disk.

Load Balancer Deployment Concept

Exchange 2016 can be deployed in various ways, in this example two servers are used. Each server hosts the Mailbox role in a DAG configuration. This provides high availability and uses a minimum number of Exchange Servers.

Clients then connect to the Virtual Services (VIPs) on the load balancer rather than connecting directly to one of the Exchange servers. These connections are then load balanced across the Exchange servers to distribute the load according to the load balancing algorithm selected.



VIP = Virtual IP Addresses

Note

The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to [Configuring HA - Adding a Secondary Appliance](#) for more details on configuring a clustered pair.

Virtual Service (VIP) Requirements

To provide load balancing and HA for Exchange 2016, the following VIPs are required:

- HTTPS (for all HTTPS based services)
- SMTP

Optionally, additional VIPs may be required as follows:

- HTTP (for redirecting to HTTPS, please refer to [Configuring an HTTP to HTTPS redirect for OWA](#) for more details)
- IMAP4
- POP3

Note

IMAP4 and POP3 are not typically used. Therefore these VIPs are not generally required.

Load Balancer Deployment Modes

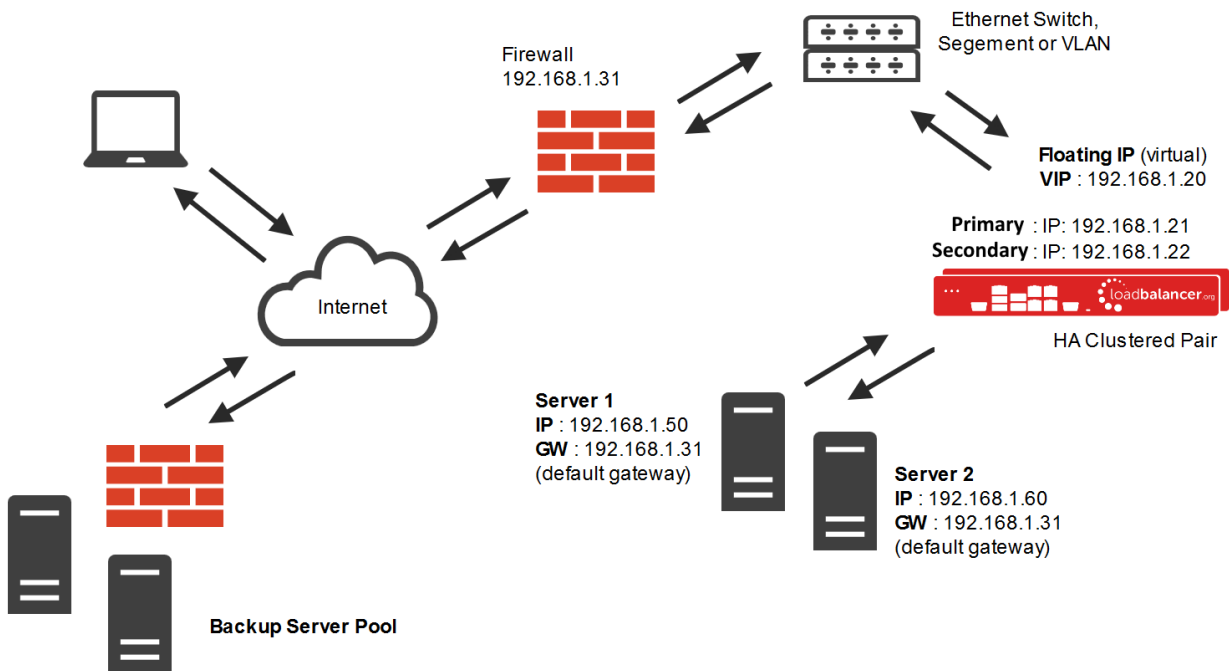
The load balancer can be deployed in 4 fundamental ways: *Layer 4 DR mode*, *Layer 4 NAT mode*, *Layer 4 SNAT mode* and *Layer 7 SNAT mode*.

For Exchange 2016, either layer 7 SNAT mode or layer 4 DR is normally used. These modes are described below and are used for the configurations presented in this guide.

Layer 7 SNAT Mode

Layer 7 SNAT mode uses a proxy (HAProxy) at the application layer. Inbound requests are terminated on the load balancer and HAProxy generates a new corresponding request to the chosen Real Server. As a result, Layer 7 is

typically not as fast as the Layer 4 methods. Layer 7 is typically chosen when either enhanced options such as SSL termination, cookie based persistence, URL rewriting, header insertion/deletion etc. are required, or when the network topology prohibits the use of the layer 4 methods.

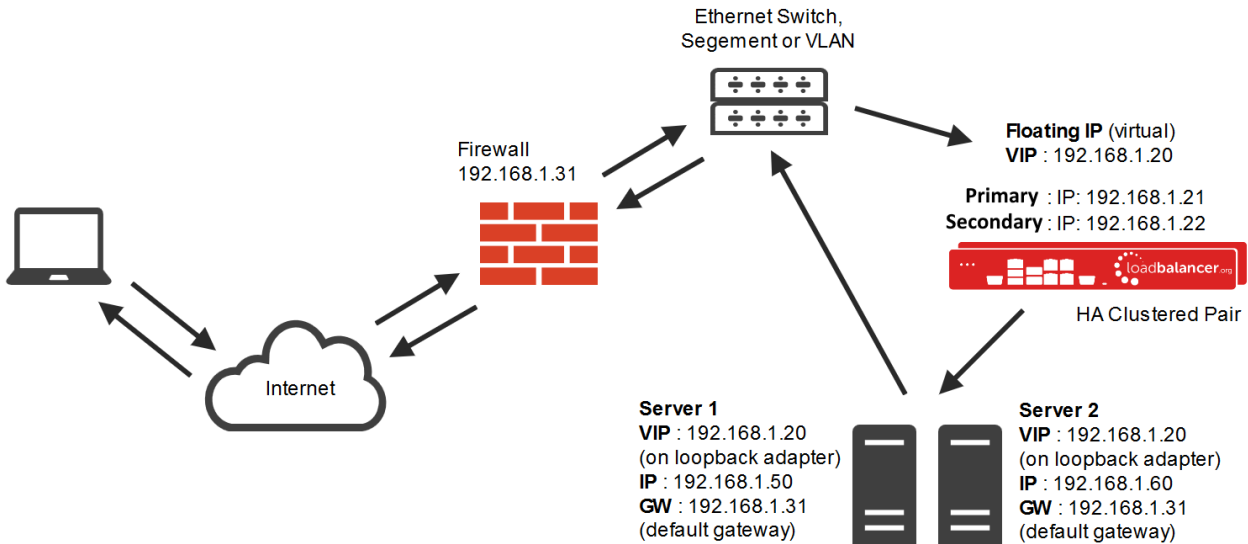


- Because layer 7 SNAT mode is a full proxy, any server in the cluster can be on any accessible subnet including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's own IP address by default, or any other local appliance IP address if preferred (e.g. the VIP address). This can be configured per layer 7 VIP. If required, the load balancer can be configured to provide the actual client IP address to the Real Servers in 2 ways. Either by inserting a header that contains the client's source IP address, or by modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. For more information on these methods please refer to [Transparency at Layer 7](#).
- Layer 7 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, eth0 is normally used for the internal network and eth1 is used for the external network although this is not mandatory.
- Requires no additional configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

Layer 4 DR Mode

One-arm direct routing (DR) mode is a very high performance solution that requires little change to your existing infrastructure.

Note | Kemp, Brocade, Barracuda & A10 Networks call this *Direct Server Return* and F5 call it *N-Path*.



- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that you need to ensure that the Real Server (and the load balanced application) respond to both the Real Server's own IP address and the VIP.
- The Real Servers 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*. For more information please refer to [DR Mode Considerations](#).
- 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 with DR mode, e.g. VIP:80 → RIP:8080 is not supported.
- DR mode is transparent, i.e. the Real Server will see the source IP address of the client.

Our Recommendation

For simplicity we recommend using layer 7 SNAT mode. This mode requires no changes to the Exchange Servers and enables the Exchange Servers to be located on any route-able network.

Is SSL Offloading Required?

We generally recommend that SSL is terminated on the Exchange servers for scalability and effective load sharing. However, when using layer 7 SNAT mode, by default the client IP address is lost and is replaced by the load balancer's own IP address. Therefore, Exchange audit logs contain the load balancer's IP address and not the clients.

If this is an issue for your environment, X-Forwarded-For headers can be inserted by the load balancer which then enables IIS on each Exchange server to be configured to log the client address – for more information, please refer to [this Microsoft article](#). To allow the header to be inserted, SSL must be terminated on the load balancer. Once inserted, traffic is re-encrypted from the load balancer to the Exchange Servers.

- To configure the appliance using Layer 7 SNAT mode *without* SSL termination, refer to [Appliance Configuration – Using Layer 7 SNAT Mode \(without SSL Offload\)](#).
- For configuring appliance using Layer 7 SNAT mode *with* SSL termination, refer to [Appliance Configuration – Using Layer 7 SNAT Mode \(with SSL Offload\)](#).

System Administrators typically want to lock down a receive connector to accept SMTP connections only from a controlled set of devices such as external smart mail hosts, printers, networked photocopiers etc. However, when using layer 7 SNAT mode - which is not transparent, this is not possible. Instead, we recommend using the load balancer's built in firewall to configure SMTP lockdown as described in [Configuring Firewall Rules to Lockdown SMTP](#).

Other Options:

Note

1 – Configure a layer 4 VIP for SMTP rather than a layer 7 based VIP. Layer 4 is transparent by default so the source IP address is maintained. This is covered in [Using a Layer 4 Virtual Service for SMTP](#). This requires the ARP problem to be solved – this requires loopback adapters to be installed on each Exchange Server and also modification to each servers strong / weak host model.

2 – Enable full layer 7 transparency using TPROXY. This is covered in [Enabling Layer 7 Transparency using TPROXY](#). This requires the load balancer to be deployed in a 2-arm configuration where the load balancer becomes the default gateway for the Exchange Servers.

8. Configuring Exchange 2016 for Load Balancing

1) External Access Domain

This can be configured using the EAC. Select *servers > virtual directories* and then click the spanner icon. This will open the form shown below. All Mailbox Servers should be configured with a valid external name, e.g. `mail.lbtestdom.com`

configure external access domain

Select the Client Access servers to use with the external URL.

+ -

NAME
EXCH2016-MBOX1
EXCH2016-MBOX2

Enter the domain name you will use with your external Client Access servers (example:mail.contoso.com).

2) Virtual Directories

The Internal and External URLs for the various virtual directories need to be configured to suit your environment. The External URLs are automatically set to be the same as the external access domain when this is configured, but can be changed if needed. The Internal URLs must be set individually by clicking the Edit (pen) icon for each virtual directory. All settings can be configured using the EAC option: *servers > virtual directories* as shown below:

servers databases database availability groups **virtual directories** certificates

Select server: All servers
Select type: All



NAME	SERVER	TYPE	V...	LAST MODIFI...	
PowerShell (Default W...	EXCH2016-MBOX1	Power...	V...	02/12/2015 ...	
PowerShell (Default W...	EXCH2016-MBOX2	Power...	V...	02/12/2015 ...	
owa (Default Web Site)	EXCH2016-MBOX1	OWA	V...	02/12/2015...	owa (Default Web Site) Website: Default Web Site Authentication: Basic, FBA Outlook Web App version: Exchange2013 External URL: https://mail.lbtestdom.com/owa
owa (Default Web Site)	EXCH2016-MBOX2	OWA	V...	02/12/2015 ...	
OAB (Default Web Site)	EXCH2016-MBOX1	OAB	V...	02/12/2015 ...	

3) Outlook Anywhere

This is configured using the EAC. Select *servers > servers* and then click the edit (pen) icon next to each sever, click the Outlook Anywhere option as shown below to change the setting. The external and internal names for each server should be configured as required, e.g. **mail.lbtestdom.com**.

EXCH2016-MBOX1

general

databases and database availability groups

POP3

IMAP4

unified messaging

DNS lookups

transport limits

transport logs

▶ Outlook Anywhere

Outlook Anywhere allows your users to connect to their Exchange mailboxes via Outlook. [Learn more](#)

Specify the external host name (for example, contoso.com) that users will use to connect to your organization.

mail.lbtestdom.com

*Specify the internal host name (for example, contoso.com) that users will use to connect to your organization.

mail.lbtestdom.com

*Specify the authentication method for external clients to use when connecting to your organization:

Negotiate

Allow SSL offloading

Save

Cancel

4) Autodiscover

Internal


The Service Connection Point (SCP) object contains the authoritative list of Autodiscover service URLs for the

forest. The `Set-ClientAccessService` cmdlet can be used to update the SCP object as shown in the following example:

```
Set-ClientAccessService -Identity "EXCH2016-MBOX1" -
AutoDiscoverServiceInternalUri
"https://autodiscover.lbtestdom.com/autodiscover/autodiscover.xml"
```

Once configured, the *Test Email AutoConfiguration* option available when <CTRL> right-clicking the Outlook icon in the taskbar can be used to view these settings as shown below:

Note | The minimum Outlook client for Exchange 2016 is Outlook 2010.



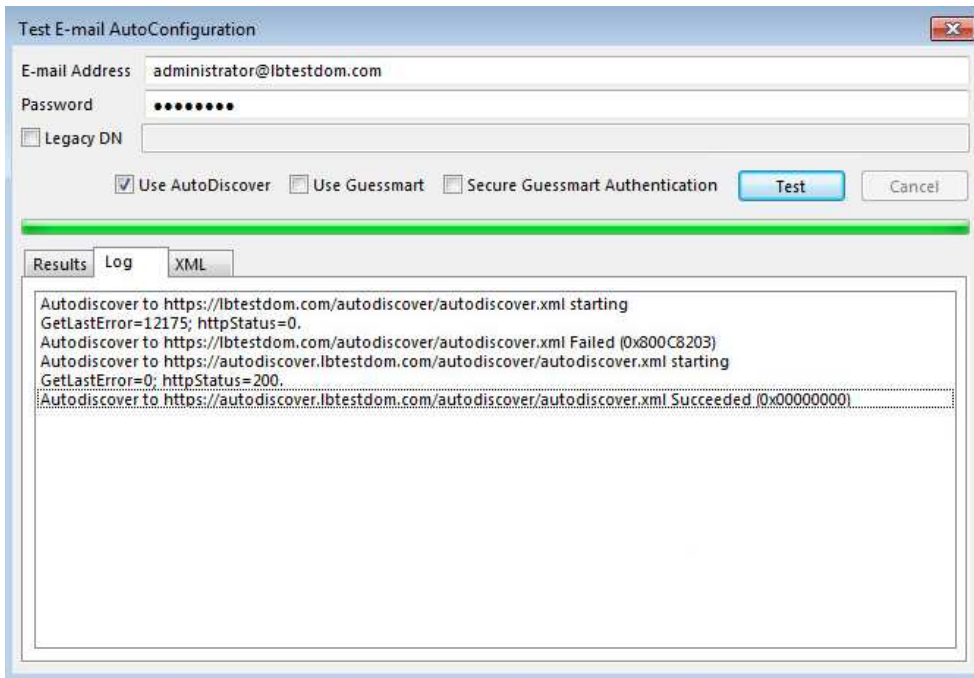
External

When Outlook is started on a client that is not domain-connected, it first tries to locate the Autodiscover service by looking up the SCP object in Active Directory. Because the client is unable to contact Active Directory, it tries to locate the Autodiscover service by using DNS. In this scenario, the client will determine the domain of the user's e-mail address, and then check DNS by using two predefined URLs. For the SMTP domain `lbtestdom.com`, Outlook will try the following two URLs to try to connect to the Autodiscover service:

`https://lbtestdom.com/autodiscover/autodiscover.xml`

`https://autodiscover.lbtestdom.com/autodiscover/autodiscover.xml`

Again, this can be seen using the *Test Email AutoConfiguration* option as shown below:



5) Certificates

The recommended approach is to use SAN certificates and specify all required namespaces. It's also possible to use wildcard certs if preferred. Certificate requests can be generated using either the graphical based Exchange Admin Center or the command based Exchange Management Shell.

The EAC can also be used to import/export certificates using the `server > certificates > More` option.

Important

The same certificate and private key must be deployed on all Exchange Servers.

6) Send & Receive Connectors

By default no send connectors are created when Exchange 2016 is installed. A send connector must be created manually that either sends outbound email messages to a smart host or directly to their recipient using DNS.

Five receive connectors are automatically created by default. The table below lists these connectors:

Receive Connector	Role	Purpose
Default <server name>	MBOX	Accepts connections from Mailbox servers running the Transport service and from Edge servers
Client Proxy <server name>	MBOX	Accepts connections from front-end servers. Typically, messages are sent to a front-end server over SMTP
Default FrontEnd <server name>	MBOX	Accepts connections from SMTP senders over port 25. This is the common messaging entry point into your organization
Outbound Proxy Frontend <server name>	MBOX	Accepts messages from a Send Connector on a back-end server, with front-end proxy enabled
Client Frontend <server name>	MBOX	Accepts secure connections, with Transport Layer Security (TLS) applied

For more information on mail connectors please refer to the following Technet article:

Adding Connectors

Connectors can be created using the Exchange Administration Center (EAC) or the Exchange Management Shell. Receive connectors must use a unique combination of IP address bindings, port number assignments, and remote IP address ranges from which mail is accepted. Multiple send connectors can be created, this is typically done to enable multiple outbound email routes to be specified that have different costs.

The exact connector configuration depends on your specific environment and requirements.

7) DNS Configuration

Configure appropriate internal and external DNS entries for the various Internal and External URL's that have been defined in steps 1) to 4). The DNS entries should point at the HTTPS VIP on the load balancer – assuming a simple namespace design as shown below:

DNS record	Purpose
mail.lbtestdom.com	Points at the VIP used for all HTTPS based services
autodiscover.lbtestdom.com	Points at the VIP used for all HTTPS based services

Note

If multiple VIPs are defined for the various Virtual Directories, DNS should be configured accordingly.

8) Additional Exchange Server Configuration Steps (depends on Load balancing method)

The steps required depend on the load balancing mode used as described below.

SNAT Mode

When using SNAT mode, no additional configuration changes to the Exchange Servers are required.

DR Mode

The 'ARP problem' must be solved on each Exchange Server for DR mode to work. For detailed steps on solving the ARP problem for Windows 2012 and later, please refer to [Solving the ARP Problem](#).

9) IIS Restart (Important)

Once all Exchange configuration is complete restart IIS on each server (or reboot the server) to ensure all changes are applied. This can be done using the following command in a command or Powershell Window:

```
iisreset /restart
```

9. Loadbalancer.org Appliance – the Basics

Virtual Appliance

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

The Virtual Appliance can be downloaded [here](#).

Note

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

Note

Please refer to [The Virtual Appliance - Hypervisor Deployment](#) and the ReadMe.txt text file included in the VA download for more detailed information on deploying the VA using various Hypervisors.

Note

For the VA, 4 NICs are included but only eth0 is connected by default at power up. If the other NICs are required, these should be connected using the network configuration screen within the Hypervisor.

Initial Network Configuration

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

Important

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

Accessing the WebUI

The WebUI is accessed using a web browser. By default, user authentication is based on local Apache .htaccess files. User administration tasks such as adding users and changing passwords can be performed using the WebUI menu option: *Maintenance > Passwords*.

Note

A number of compatibility issues have been found with various versions of Internet Explorer and Edge. The WebUI has been tested and verified using both Chrome & Firefox.

Note

If required, users can also be authenticated against LDAP, LDAPS, Active Directory or Radius. For more information please refer to [External Authentication](#).

1. Using a browser, access the WebUI using the following URL:

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

2. Log in to the WebUI:

Username: loadbalancer

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

Note

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

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

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

WARNING: YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS.

Buy with confidence. All purchases come with a 90 day money back guarantee. Already bought? Enter your license key [here](#)

Buy Now

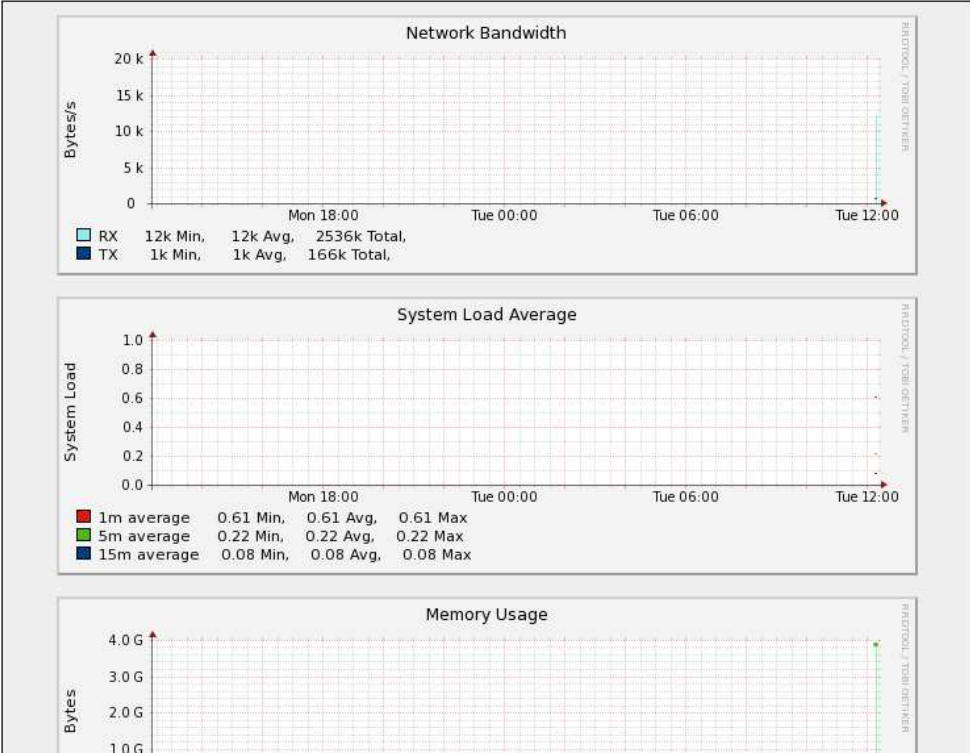
System Overview ?

2022-01-11 12:18:37 UTC

Would you like to run the Setup Wizard?

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

No Virtual Services configured.



Note

The WebUI for the VA is shown, the hardware and cloud appliances are very similar. The yellow licensing related message is platform & model dependent.

- You'll be asked if you want to run the Setup Wizard. If you click **Accept** the Layer 7 Virtual Service configuration wizard will start. If you want to configure the appliance manually, simple click **Dismiss**.

Main Menu Options

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

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

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

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

View Configuration - Display the saved appliance configuration settings

Reports - View various appliance reports & graphs

Logs - View various appliance logs

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

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

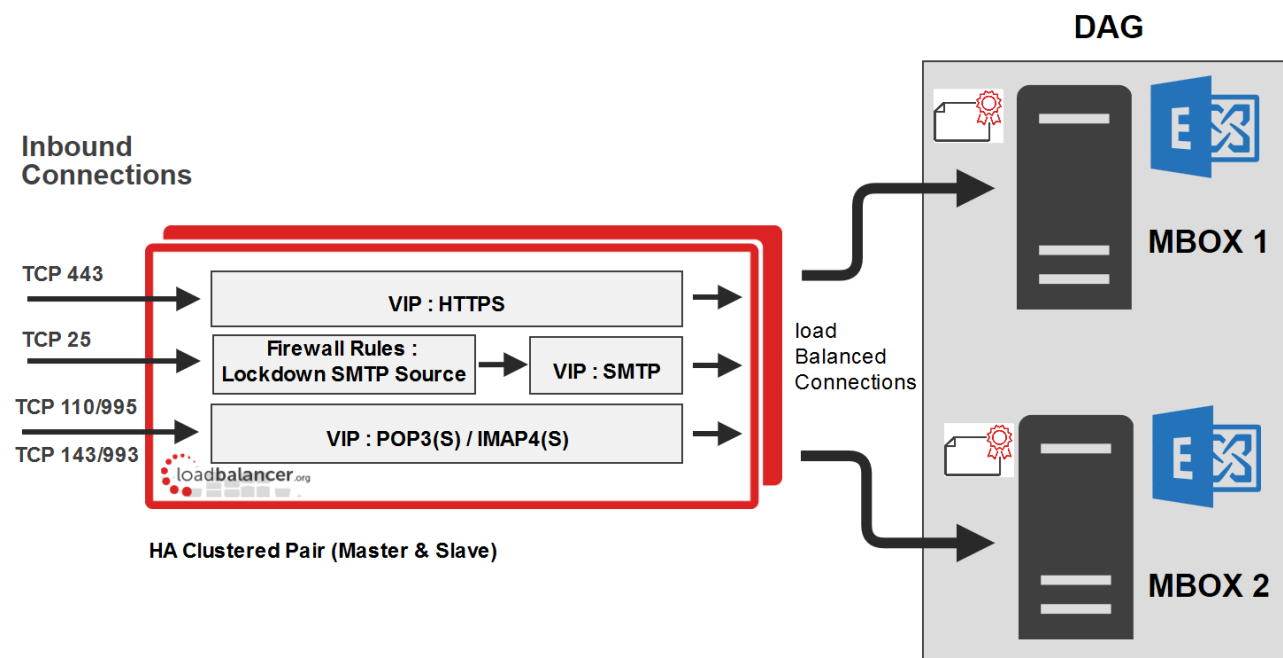
HA Clustered Pair Configuration

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

10. Appliance Configuration – Using Layer 7 SNAT Mode (without SSL Offload)

Load Balancer Deployment Overview

The diagram below illustrates how the load balancer is configured and deployed.



Notes

- Layer 7 is not transparent by default. This means that the client source IP address is lost and is replaced by the IP address of the load balancer. All Exchange audit logs will show the IP address of the load balancer, not the clients. If this is an issue, please refer to the configuration option in [Appliance Configuration – Using Layer 7 SNAT Mode \(with SSL Offload\)](#) where X-Forwarded-For headers are used to record the client IP address in the Exchange server's IIS logs.
- System Administrators typically want to lock down a receive connector to accept SMTP connections only from a controlled set of devices such as external smart mail hosts, printers, networked photocopiers etc. However, when using layer 7 SNAT mode - which is not transparent, this is not possible. Instead, we recommend using the load balancer's built in firewall to configure SMTP lockdown as described in [Configuring Firewall Rules to Lockdown SMTP](#).

Other Options:

- 1 - Configure a layer 4 VIP for SMTP rather than a layer 7 based VIP. Layer 4 is transparent by default so the

source IP address is maintained. This is covered in [Using a Layer 4 Virtual Service for SMTP](#). This requires the ARP problem to be solved – this requires loopback adapters to be installed on each Exchange Server and also modification to each servers strong / weak host model.






2 - Enable full layer 7 transparency using TPROXY. This is covered in [Enabling Layer 7 Transparency using TPROXY](#). This requires the load balancer to be deployed in a 2-arm configuration where the load balancer becomes the default gateway for the Exchange Servers.

Load Balancer Configuration

Configure VIP1 – Mailbox Server Role HTTPS Services

a) Setting up the Virtual Service

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Service* and click **Add a New Virtual Service**.
2. Enter the following details:

Label	<input type="text" value="MBOX-HTTPS"/>	
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	
Ports	<input type="text" value="443"/>	
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	
Manual Configuration	<input type="checkbox"/>	
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate label for the VIP, e.g. **MBOX-HTTPS**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **443**.
6. Set *Layer 7 Protocol* set to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Set *Balance mode* to **Weighted Round Robin**.

Note

Microsoft recommends that 'Round Robin' rather than 'Least Connection' should be used to help prevent over loading servers when they are brought online. This could occur if Least Connection was selected, since the load balancer would try to balance the number of connections across all real servers and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

10. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.

- In the *Health Checks* section set *Health Checks* to **Negotiate HTTPS (GET)**.
- Set *Request to send* to `owa/healthcheck.htm`.

Note

As mentioned earlier, any other Exchange virtual directory (e.g. ECP, EWS etc.) can be used if preferred or more appropriate. All have an associated `healthcheck.htm` that can be used in the same way. Note that `healthcheck.htm` is generated in-memory based on the component state of the protocol in question and does not physically exist on disk.

- Leave *Response expected* blank, this will configure the load balancer to look for a '200 OK' response.
- Scroll down to the *Other* section and click **[Advanced]**.
- Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes).
- Click **Update**.

b) Setting up the Real Servers

- Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
- Enter the following details:

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Real Server Port	<input type="text" value="443"/>	?
Weight	<input type="text" value="100"/>	?

- Enter an appropriate label for the RIP, e.g. **MBOX1**.
- Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
- Change the *Real Server Port* field to **443**.
- Click **Update**.
- Repeat the above steps to add your other Mailbox Server(s).

c) Configure HTTP to HTTPS OWA Redirect

If required, the load balancer can be configured to automatically redirect users who attempt to connect to `http://<URL-to-access-OWA>` to `https://<URL-to-access-OWA>`. For details on configuring this, please refer to [Configuring an HTTP to HTTPS redirect for OWA](#).

Configure VIP2 – Mailbox Server Role IMAP4/POP3 Services

a) Setting up the Virtual Service

Note

These steps show IMAP4 settings, for POP3 change the port numbers from 143 & 993 to 110 & 995.

- Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Service* and click **Add a New Virtual Service**.
- Enter the following details:

Label	<input type="text" value="MBOX-IMAP4"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	?
Ports	<input type="text" value="143,993"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
Manual Configuration	<input type="checkbox"/>	?

- Enter an appropriate label for the VIP, e.g. **MBOX-IMAP4**.
- Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
- Set the *Virtual Service Ports* field to **143,993**.
- Set *Layer 7 Protocol* to **TCP Mode**.
- Click **Update**.
- Now click **Modify** next to the newly created VIP.
- Set *Balance mode* to **Weighted Round Robin**.

Note

Microsoft recommends that 'Round Robin' rather than 'Least Connection' should be used to help prevent over loading servers when they are brought online. This could occur if Least Connection was selected, since the load balancer would try to balance the number of connections across all real servers and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

- Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
- Scroll down to the *Other* section and click **[Advanced]**.
- Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes).
- Click **Update**.

b) Setting up the Real Servers

- Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
- Enter the following details:

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Real Server Port	<input type="text"/>	?
Weight	<input type="text" value="100"/>	?

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

Configure VIP3 – Mailbox Server Role SMTP Services

a) Setting up the Virtual Service

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Service* and click **Add a New Virtual Service**.
2. Enter the following details:





Label	<input type="text" value="MBOX-SMTP"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	?
Ports	<input type="text" value="25"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
Manual Configuration	<input type="checkbox"/>	?

3. Enter an appropriate label for the VIP, e.g. **MBOX-SMTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **25**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
10. Scroll down to the *Other* section and click **[Advanced]**.

11. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes).
12. Click **Update**.

b) Setting up the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration* > *Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	<input type="text" value="MBOX1"/>	
Real Server IP Address	<input type="text" value="192.168.30.20"/>	
Real Server Port	<input type="text" value="25"/>	
Weight	<input type="text" value="100"/>	

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **25**.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

Configuring Firewall Rules to Lockdown SMTP

Because layer 7 is not transparent by default, it's not possible to filter inbound SMTP connections by IP address at the receive connector. Our recommended way to address this is to use the load balancer's built-in firewall to control which hosts can connect to the SMTP VIP on port 25. Please refer to [Configuring Firewall Rules to Lockdown SMTP](#) for details of how to configure this.

Additional Settings if using Kerberos Authentication

If you're using Kerberos to authenticate your Exchange users and these users are members of a large number of AD security groups and/or have a large SID history, Kerberos tickets may become so large that they no longer fit in the standard 16K HAProxy response buffer. For Windows 2012 and later, the default **MaxTokenSize** is set to 48K. In addition, there is a new KDC policy setting that can be enabled to log an event in the system event log if a Kerberos ticket is larger than a certain size (the default setting is 12k). If you determine that tickets in your environment are larger than 16K, the default response buffer size on the load balancer must be increased.

To increase the Request buffer size:

1. Go to *Cluster Configuration* > *Layer 7 – Advanced Configuration*.
2. Set the *Request buffer length* to the required value, e.g. **51200** (i.e. 50K).

Finalizing the Configuration

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the blue box at the top

of the screen or by using the *Restart Services* menu option:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.

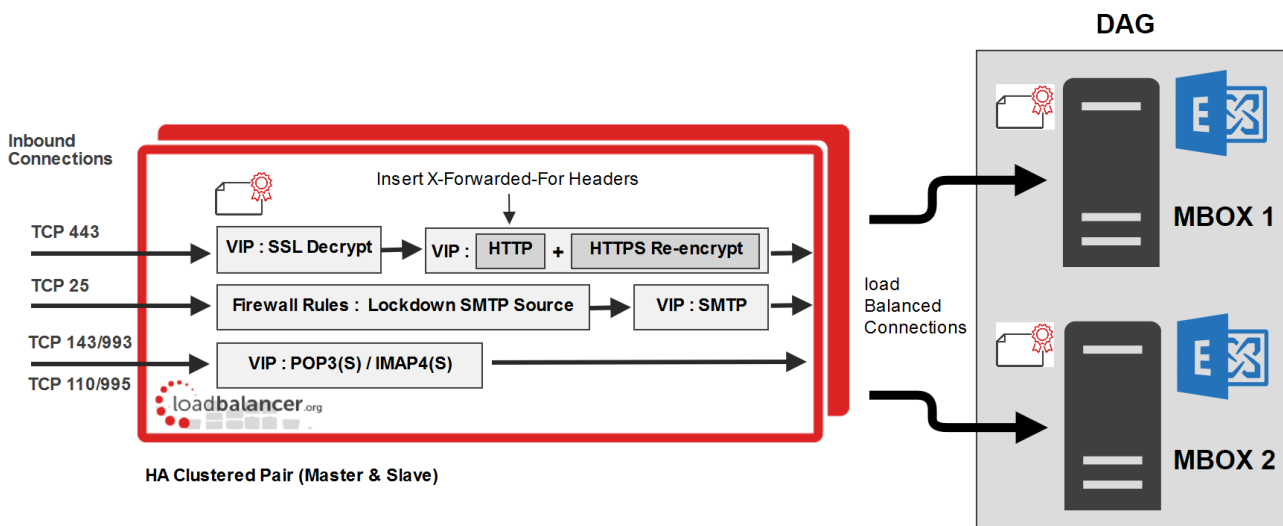
Exchange Server Configuration Steps

No additional configuration is required when SSL is terminated on the Exchange Servers.

11. Appliance Configuration – Using Layer 7 SNAT Mode (with SSL Offload)

Load Balancer Deployment Overview

The diagram below illustrates how the load balancer is configured and deployed. The key difference to the previous configuration is that SSL is terminated on the load balancer.



Notes

- Layer 7 is not transparent by default. This means that the client source IP address is lost and is replaced by the IP address of the load balancer. To allow the client IP address to be passed to the Exchange Servers, SSL is terminated on the load balancer which enables X-forwarded-For headers to be inserted. The Exchange servers can then be configured so that this address is included in the IIS logs as described in [this Microsoft article](#).
- System Administrators typically want to lock down a receive connector to accept SMTP connections only from a controlled set of devices such as external smart mail hosts, printers, networked photocopiers etc. However, when using layer 7 SNAT mode - which is not transparent, this is not possible. Instead, we recommend using the load balancer's built in firewall to configure SMTP lockdown as described in [Configuring Firewall Rules to Lockdown SMTP](#).

Other Options:

1 - Configure a layer 4 VIP for SMTP rather than a layer 7 based VIP. Layer 4 is transparent by default so the source IP address is maintained. This is covered in [Using a Layer 4 Virtual Service for SMTP](#). This requires the ARP problem to be solved – this requires loopback adapters to be installed on each Exchange Server and also modification to each servers strong / weak host model.

2 - Enable full layer 7 transparency using TPROXY. This is covered in [Enabling Layer 7 Transparency using TPROXY](#). This requires the load balancer to be deployed in a 2-arm configuration where the load balancer becomes the default gateway for the Exchange Servers.

Load Balancer Configuration

Configure VIP1 – Mailbox Server Role HTTP/HTTPS Services

a) Setting up the Virtual Service

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Service* and click **Add a New Virtual Service**.
2. Enter the following details:

Label	<input type="text" value="MBOX-HTTP"/>	
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	
Ports	<input type="text" value="80"/>	
Protocol		
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	
Manual Configuration	<input type="checkbox"/>	
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate label for the VIP, e.g. **MBOX-HTTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **80**.
6. Set *Layer 7 Protocol* set to **HTTP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Set *Balance mode* to **Weighted Round Robin**.

Note

Microsoft recommends that 'Round Robin' rather than 'Least Connection' should be used to help prevent over loading servers when they are brought online. This could occur if Least Connection was selected, since the load balancer would try to balance the number of connections across all real servers and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

10. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
11. In the *Health Checks* section set *Health Checks* to **Negotiate HTTPS (GET)**.
12. Set *Request to send* to **owa/healthcheck.htm**

Note

As mentioned earlier, any other Exchange virtual directory (e.g. ECP, EWS etc.) can be used if preferred or more appropriate. All have an associated healthcheck.htm that can be used in the same way. Note that healthcheck.htm is generated in-memory based on the component state of the protocol in question and does not physically exist on disk.

13. Leave *Response expected* blank, this will configure the load balancer to look for a '200 OK' response.
14. Scroll down to the *Other* section and click **[Advanced]**.
15. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes).
16. Ensure that *Set X-forwarded-For Header* is enabled (checked).
17. Click **Update**.

b) Setting up the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration* > *Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Real Server Port	<input type="text" value="443"/>	?
Re-Encrypt to Backend	<input checked="" type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **443**.
6. Enable (check) the *Re-Encrypt to Backend* checkbox.
7. Click **Update**.
8. Repeat the above steps to add your other Mailbox Server(s).

c) Export Your SSL Certificate

When you export your certificate from Exchange, make sure that you include the private key.

d) Upload Your SSL Certificate to The Load Balancer

To upload a Certificate:

1. Using the WebUI, navigate to: *Cluster Configuration* > *SSL Certificates*.
2. Click **Add a new SSL Certificate** & select *Upload prepared PEM/PFX file*.

I would like to:

Upload prepared PEM/PFX file

Create a new SSL Certificate Signing Request (CSR)

Create a new Self-Signed SSL Certificate.

Label

File to upload No file chosen

3. Enter a suitable *Label* (name) for the certificate, e.g. **ExchangeCert**.
4. Browse to and select the certificate file to upload (PEM or PFX format).
5. Enter the password, if applicable.
6. Click **Upload Certificate**. If successful, a message similar to the following will be displayed:

Information: cert1 SSL Certificate uploaded successfully.

e) Configure SSL Termination

To configure an SSL VIP:

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Termination* and click **Add a new Virtual Service**.

Label

Associated Virtual Service

Virtual Service Port

SSL Operation Mode

SSL Certificate

Source IP Address

Enable Proxy Protocol

Bind Proxy Protocol to L7 VIP

2. Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **MBOX-HTTP**.

Note

Once the VIP is selected, the *Label* field will be auto-populated with **SSL-MBOX-HTTP**. This can be changed if preferred.

3. Ensure that the *Virtual Service Port* is set to **443**.
4. Leave *SSL Operation Mode* set to **High Security**.

5. Select the required *SSL Certificate*.
6. Click **Update**.

f) Configure HTTP to HTTPS OWA Redirect

If required, the load balancer can be configured to automatically redirect users who attempt to connect to <http://<URL-to-access-OWA>> to <https://<URL-to-access-OWA>>. For details on configuring this, please refer to [Configuring an HTTP to HTTPS redirect for OWA](#).

Configure VIP2 – Mailbox Server Role IMAP4/POP3 Services

a) Setting up the Virtual Service

Note | These steps show IMAP4 settings, for POP3 change the port numbers from 143 & 993 to 110 & 995.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Service* and click **Add a New Virtual Service**.
2. Enter the following details:

Label	<input type="text" value="MBOX-IMAP4"/>	
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	
Ports	<input type="text" value="143,993"/>	
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	
Manual Configuration	<input type="checkbox"/>	

3. Enter an appropriate label for the VIP, e.g. **MBOX-IMAP4**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **143,993**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Set *Balance mode* to **Weighted Round Robin**.

Note | Microsoft recommends that 'Round Robin' rather than 'Least Connection' should be used to help prevent over loading servers when they are brought online. This could occur if Least Connection was selected, since the load balancer would try to balance the number of connections across all real servers and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

10. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
11. Scroll down to the *Other* section and click **[Advanced]**.
12. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes).
13. Click **Update**.

b) Setting up the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Real Server Port	<input type="text"/>	?
Weight	<input type="text" value="100"/>	?

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

Configure VIP3 – Mailbox Server Role SMTP Services

a) Setting up the Virtual Service





1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a New Virtual Service**.
2. Enter the following details:

Label	<input type="text" value="MBOX-SMTP"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	?
Ports	<input type="text" value="25"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
Manual Configuration	<input type="checkbox"/>	?

3. Enter an appropriate label for the VIP, e.g. **MBOX-SMTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **25**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
10. Scroll down to the *Other* section and click **[Advanced]**.
11. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes).
12. Click **Update**.

b) Setting up the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	<input type="text" value="MBOX1"/>	
Real Server IP Address	<input type="text" value="192.168.30.20"/>	
Real Server Port	<input type="text" value="25"/>	
Weight	<input type="text" value="100"/>	

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **25**.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

Configuring Firewall Rules to Lockdown SMTP

Because layer 7 is not transparent by default, it's not possible to filter inbound SMTP connections by IP address at the receive connector. Our recommended way to address this is to use the load balancer's built-in firewall to control which hosts can connect to the SMTP VIP on port 25. Please refer to [Configuring Firewall Rules to Lockdown SMTP](#) for details of how to configure this.

Additional Settings if using Kerberos Authentication

If you're using Kerberos to authenticate your Exchange users and these users are members of a large number of AD security groups and/or have a large SID history, Kerberos tickets may become so large that they no longer fit in the standard 16K HAProxy response buffer. For Windows 2012 and later, the default **MaxTokenSize** is set to 48K. In

addition, there is a new KDC policy setting that can be enabled to log an event in the system event log if a Kerberos ticket is larger than a certain size (the default setting is 12k). If you determine that tickets in your environment are larger than 16K, the default response buffer size on the load balancer must be increased.

To increase the Request buffer size:

1. Go to *Cluster Configuration > Layer 7 – Advanced Configuration*.
2. Set the *Request buffer length* to the required value, e.g. **51200** (i.e. 50K).

Finalizing the Configuration

To apply the new settings, HAProxy and STunnel must both be reloaded. This can be done using the buttons in the blue box at the top of the screen or by using the *Restart Services* menu option:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.
3. Click **Reload STunnel**.

Exchange Server Configuration Steps

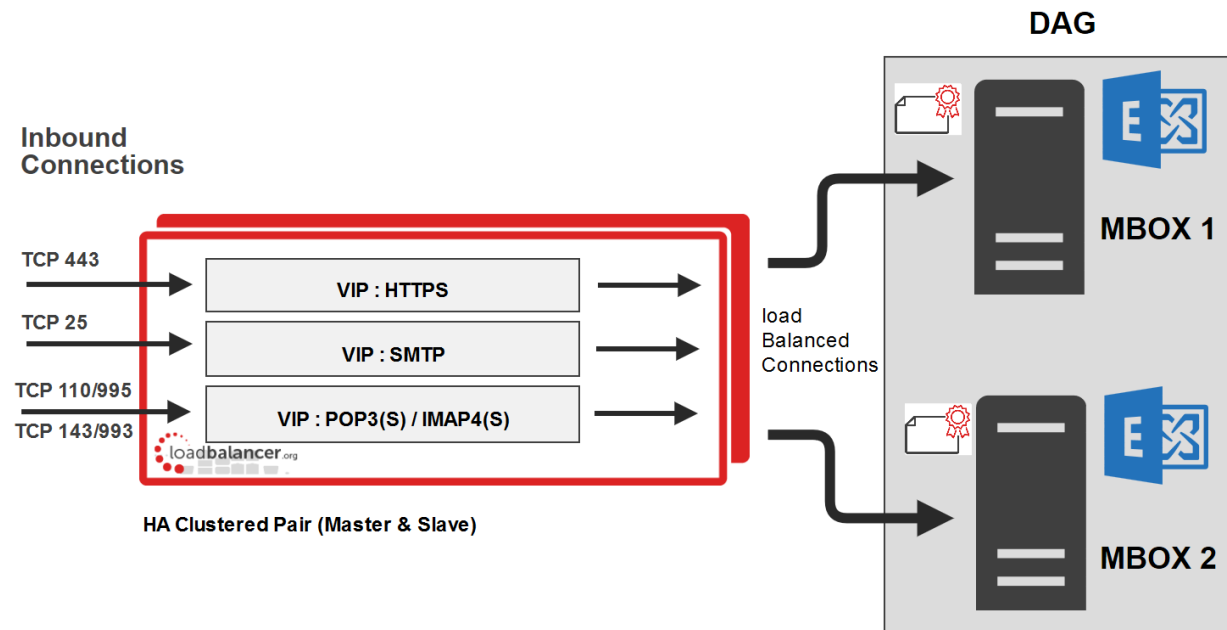
Configure IIS logging to Capture XFF Header IP Addresses

Please refer to [this Microsoft article](#) for configuration steps.

12. Appliance Configuration – Using Layer 4 DR Mode

Load Balancer Deployment Overview

The diagram below illustrates how the load balancer is configured and deployed.



Notes

- Layer 4 DR mode is transparent by default. This means that the client source IP address is maintained through to the Exchange Servers & the audit logs.

- When using DR mode, System Administrators are able to lock down the receive connector to accept SMTP connections only from a controlled set of devices such as external smart mail hosts, printers, networked photocopiers etc. As mentioned earlier, this is because DR mode is transparent, so source IP addresses are preserved through the load balancer to the Exchange Servers.

Load Balancer Configuration

Configure VIP1 – Mailbox Server Role HTTPS Services

a) Setting up the Virtual Service

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

Label	<input type="text" value="MBOX-HTTPS"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	?
Ports	<input type="text" value="443"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate label for the VIP, e.g. **MBOX-HTTPS**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **443**.
6. Leave *Protocol* set to **TCP**.
7. Leave *Forwarding Method* set to **Direct Routing**.
8. Click **Update**.
9. Now click **Modify** next to the newly created VIP.
10. Set *Balance mode* to **Weighted Round Robin**.

Note

Microsoft recommends that 'Round Robin' rather than 'Least Connection' should be used to help prevent over loading servers when they are brought online. This could occur if Least Connection was selected, since the load balancer would try to balance the number of connections across all real servers and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

11. Un-check the *Persistence* option.
12. Set *Check Type* to **Negotiate**.

13. Set *Protocol* to **HTTPS**.
14. Set *Request to send* to **owa/healthcheck.htm**.

Note

As mentioned earlier, any other Exchange virtual directory (e.g. ECP, EWS etc.) can be used if preferred or more appropriate. All have an associated healthcheck.htm that can be used in the same way. Note that healthcheck.htm is generated in-memory based on the component state of the protocol in question and does not physically exist on disk.

15. Set *Response expected* to **200 OK**.
16. Click **Update**.

b) Setting up the Real Servers

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

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Click **Update**.
6. Repeat the above steps to add your other Mailbox Server(s).

c) Configure HTTP to HTTPS OWA Redirect

If required, the load balancer can be configured to automatically redirect users who attempt to connect to **http://<URL-to-access-OWA>** to **https://<URL-to-access-OWA>**. For details on configuring this, please refer to [Configuring an HTTP to HTTPS redirect for OWA](#).

Configure VIP2 – Mailbox Server Role IMAP4/POP3 Services

a) Setting up the Virtual Service

Note

These steps show IMAP4 settings, for POP3 change the port numbers from 143 & 993 to 110 & 995.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Service* and click **Add a New Virtual Service**.

2. Enter the following details:

Label	<input type="text" value="MBOX-IMAP4"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	?
Ports	<input type="text" value="143,993"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	?

3. Enter an appropriate label for the VIP, e.g. **MBOX-IMAP4**.

4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.

5. Set the *Virtual Service Ports* field to **143,993**.

6. Leave *Protocol* set to **TCP**.

7. Leave *Forwarding Method* set to **Direct Routing**.

8. Click **Update**.

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

10. Set *Balance mode* to **Weighted Round Robin**.

Note

Microsoft recommends that 'Round Robin' rather than 'Least Connection' should be used to help prevent over loading servers when they are brought online. This could occur if Least Connection was selected, since the load balancer would try to balance the number of connections across all real servers and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

11. Un-check the *Persistence* option.

12. Click **Update**.

b) Setting up the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.

2. Enter the following details:

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Click **Update**.
6. Repeat the above steps to add your other Mailbox Server(s).

Configure VIP3 – Mailbox Server Role SMTP Services

a) Setting up the Virtual Service

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

Label	<input type="text" value="MBOX-SMTP"/>	?
Virtual Service		
IP Address	<input type="text" value="192.168.30.10"/>	?
Ports	<input type="text" value="25"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	?

3. Enter an appropriate label for the VIP, e.g. **MBOX-SMTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **25**.
6. Leave *Protocol* set to **TCP**.
7. Leave *Forwarding Method* set to **Direct Routing**.
8. Click **Update**.
9. Now click **Modify** next to the newly created VIP.

10. Un-check the *Persistence* option.

11. Click **Update**.

b) Setting up the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.

2. Enter the following details:

Label	<input type="text" value="MBOX1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.20"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.

4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.

5. Click **Update**.

6. Repeat the above steps to add your other Mailbox Server(s).

Exchange Server Configuration Steps

When using layer 4 DR mode, as mentioned in [DR Mode](#), the "ARP Problem" must be solved on each Exchange server. For full details of the steps required to do this, please refer to [Solving the ARP Problem](#).

13. Testing & Verification

Note | For additional general guidance please also refer to [Testing Load Balanced Services](#).

Useful Exchange 2016 & Other Microsoft Tools

Testing Server Health-checks using Set-ServerComponentState

The Exchange Management shell cmdlet `Set-ServerComponentState` can be used to verify that the load balancer is correctly health-checking the Exchange servers. In this guide, the health-check verifies that the owa virtual directory can be accessed.

To verify that the health-check is working correctly, the following command can be used:

```
Set-ServerComponentState <SERVER> -Component OwaProxy -Requester Maintenance -State Inactive
```

Where `<SERVER>` is the hostname of the Exchange Server.

Once run, the server specified should be marked down (shown red) in the System Overview of the loadbalancer's

WebUI.

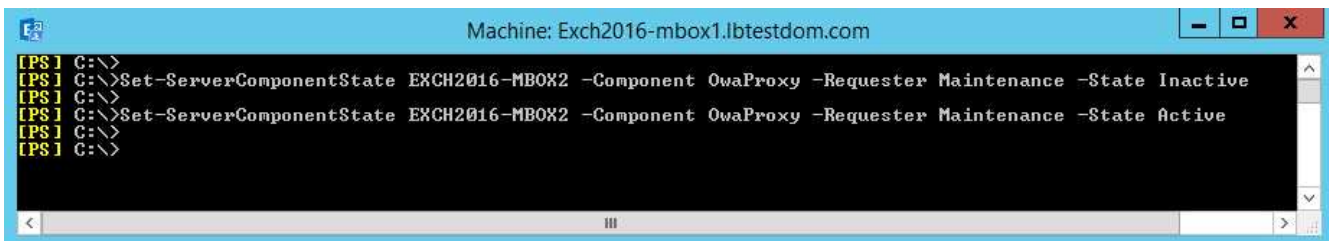
To bring it back online, use the following command:

```
Set-ServerComponentState <SERVER> -Component OwaProxy -Requester Maintenance -State Active
```

Where <SERVER> is the hostname of the Exchange Server.

Once run, the server specified should be marked up (shown green) in the System Overview of the loadbalancer's WebUI.

Exchange Management Shell:



```
Machine: Exch2016-mbox1.lbtestdom.com
[PS] C:\>
[PS] C:\>Set-ServerComponentState EXCH2016-MBOX2 -Component OwaProxy -Requester Maintenance -State Inactive
[PS] C:\>
[PS] C:\>Set-ServerComponentState EXCH2016-MBOX2 -Component OwaProxy -Requester Maintenance -State Active
[PS] C:\>
[PS] C:\>
```

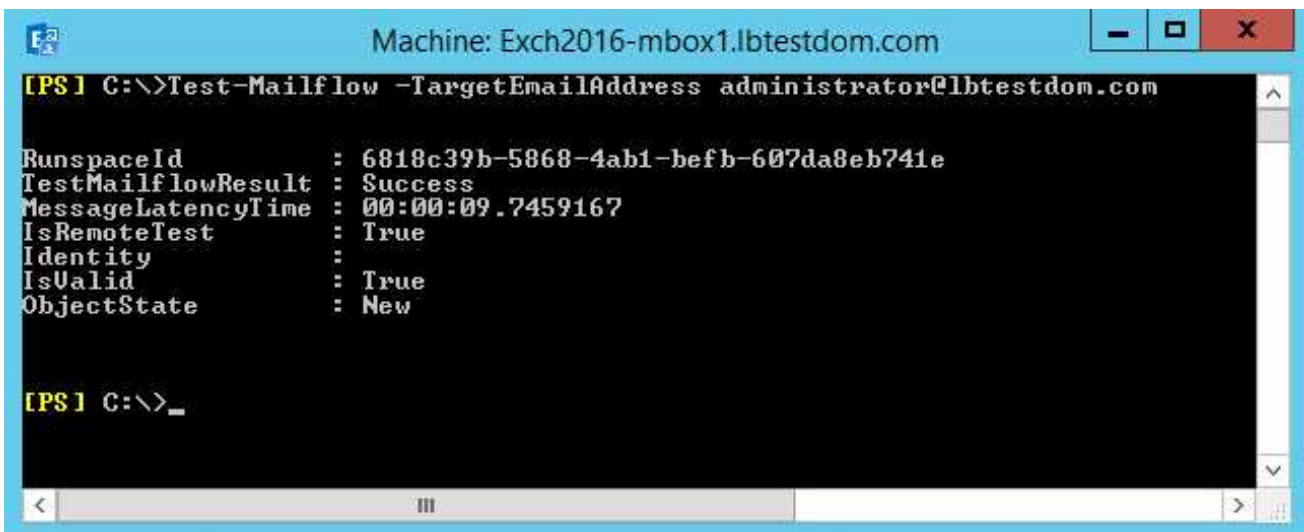
Testing Mailflow

The `Test-Mailflow` cmdlet can be used to diagnose whether mail can be successfully sent and delivered.

To send a test probe message to the administrators email address, use the following command:

```
Test-Mailflow -TargetEmailAddress \administrator@lbtestdom.com
```

Exchange Management Shell:



```
Machine: Exch2016-mbox1.lbtestdom.com
[PS] C:\>Test-Mailflow -TargetEmailAddress administrator@lbtestdom.com

RunspaceId      : 6818c39b-5868-4ab1-befb-607da8eb741e
TestMailflowResult : Success
MessageLatencyTime : 00:00:09.7459167
IsRemoteTest     : True
Identity         :
IsValid          : True
ObjectState      : New

[PS] C:\>_
```

If everything is working correctly, a new message will appear in the test users mailbox:



Testing SMTP Mail flow using Telnet

SMTP can be tested using telnet to connect to port 25, then by issuing various commands to simulate an email being sent. Using *System Overview* in the WebUI, each Mailbox Server server can be tested by 'Halting' all others then running through the tests.

To connect to port 25 of a server using Telnet, use the following command:

```
telnet <IP Address> 25
```

The following screenshot shows an example of using telnet to verify SMTP operation:

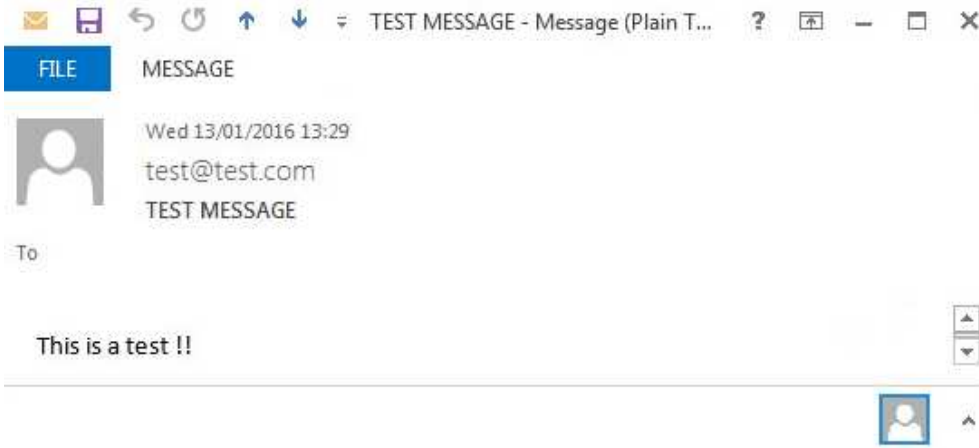
A screenshot of a terminal window titled 'root@lbmaster:~'. The terminal shows the following output:

```
[root@lbmaster ~]# telnet 192.168.112.2 25
Trying 192.168.112.2...
Connected to 192.168.112.2.
Escape character is '^]'.
220 Exch2016-mbox1.lbtestdom.com Microsoft ESMTMP MAIL Service ready at Wed, 13 Jan 2016 13:28:38 +0000
helo
250 Exch2016-mbox1.lbtestdom.com Hello [192.168.111.40]
mail from:test@test.com
250 2.1.0 Sender OK
rcpt to:administrator@lbtestdom.com
250 2.1.5 Recipient OK
data
354 Start mail input; end with <CRLF>.<CRLF>
subject:TEST MESSAGE

This is a test !!

.
250 2.6.0 <134fcfc4-dale-4ac6-a496-0483e0fd0480@Exch2016-mbox1.lbtestdom.com> [InternalId=5046586572854, Hostname=Exch2016-mbox1.lbtestdom.com] Queued mail for delivery
quit
221 2.0.0 Service closing transmission channel
Connection closed by foreign host.
[root@lbmaster ~]#
```

If everything is working correctly, a new message will appear in the test user's mailbox:



To do the same test via the load balancer, connect to the VIP rather than directly to each server, e.g.:

```
telnet mail.lbtestdom.com 25
```

Microsoft Exchange Testing Tool

The Remote Connectivity Analyzer tool available at <https://testconnectivity.microsoft.com/> is a useful Web-based Microsoft tool designed to help IT Administrators troubleshoot connectivity issues with their Exchange Server deployments. The tool simulates several client logon and mail flow scenarios. When a test fails, many of the errors have troubleshooting tips to assist the IT Administrator in correcting the problem.

Useful Appliance based Tools & Features

Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the Exchange Servers) and shows the state/health of each server as well as the state of the each cluster as a whole. The example below shows that both Mailbox Servers are healthy and available to accept connections:

SYSTEM OVERVIEW ? 2016-01-13 10:30:26 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
MBOX-HTTPS	192.168.111.100	443	0	TCP	Layer 7	Proxy
<i>REAL SERVER</i>						
rip1	192.168.112.2	443	100	0	Drain	Halt
rip2	192.168.112.3	443	100	0	Drain	Halt

The example below shows that rip2 has been put in halt mode:

SYSTEM OVERVIEW ? 2016-01-13 10:32:07 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
MBOX-HTTPS	192.168.111.100	443	0	TCP	Layer 7	Proxy
<i>REAL SERVER</i>						
rip1	192.168.112.2	443	100	0	Drain	Halt
rip2	192.168.112.3	443	0	0	Online (halt)	

Layer 4 Status Report

The Layer 4 Status report gives a summary of layer 4 configuration and running stats as shown below. This can be accessed in the WebUI using the option: *Reports > Layer 4 Status*.

LAYER 4 STATUS

Check Status

Virtual Service Real Server Forwarding Method Weight Active Connections Inactive Connections

CAS-HTTPS
192.168.111.96
port 443/tcp

Virtual Service	Real Server	Forwarding Method	Weight	Active Connections	Inactive Connections
CAS-HTTPS	rip1 192.168.111.240	Route	1	6	0
	rip2 192.168.111.241	Route	1	6	0

Layer 7 Statistics Report

The Layer 7 Statistics report gives a summary of all layer 7 configuration and running stats as shown below. This can be accessed in the WebUI using the option: *Reports > Layer 7 Status*.

HAProxy

Statistics Report for pid 8727

> General process information

pid = 8727 (process #1, nbproc = 1)
uptime = 0d 0h01m33s
system limits: memmax = unlimited; ulimit-n = 81000
maxsock = 80025; maxconn = 40000; maxpipes = 0
current conns = 12; current pipes = 0/0; conn rate = 4/sec
Running tasks: 2/17; 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

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

Display option:

- Hide DOWN servers
- Refresh now
- CSV export

External resources:

- Primary site
- Updates (v1.5)
- Online manual

Queue	Session rate		Sessions				Bytes		Denied		Errors		Warnings		Server													
	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle		
Frontend				0	5	-	10	12	40 000	22			44 059	96 458	0	0	0			OPEN								
backup	0	0	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0				1	-	Y		0	0	-	
rip1	0	0	-	0	2	5	6	-	10	10	12	906	43 441	0	0	0	0	0	1m33s UP	L4OK in 0ms	1	Y	-	0	0	0s	-	
rip2	0	0	-	0	3	5	6	-	12	12	31 183	53 017	0	0	0	0	0	0	1m33s UP	L4OK in 0ms	1	Y	-	0	0	0s	-	
Backend	0	0	0	0	5	10	12	4 000	22	22	44 059	96 458	0	0	0	0	0	0	1m33s UP		2	2	1		0	0s		

Queue	Session rate		Sessions				Bytes		Denied		Errors		Warnings		Server												
	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle	
Frontend	4	10	-	2	2	2	2 000	108			46 004	1 182 147	0	0	0			OPEN									
Backend	0	0	0	0	0	0	200	0	0	0	46 004	1 182 147	0	0	0	0	0	0	1m33s UP		0	0	0		0	0s	

Appliance Logs

Logs are available for both layer 4 and layer 7 services and can be very useful when trying to diagnose issues. Layer 4 logs are active by default and can be accessed using the WebUI option: *Logs > Layer 4*.

Layer 7 logging is not enabled by default (because its extremely verbose) and can be enabled using the WebUI option: *Cluster Configuration > Layer 7 – Advanced Configuration > Logging*, and then viewed using the WebUI option: *Logs > Layer 7*.

14. Technical Support

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

15. Further Documentation

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

16. Conclusion

Loadbalancer.org appliances provide a very cost effective solution for highly available load balanced Exchange 2016 environments.

17. Appendix

Configuring Firewall Rules to Lockdown SMTP

Because layer 7 is not transparent by default, it's not possible to filter inbound SMTP connections by IP address at the receive connector. Our recommended way to address this is to use the load balancer's built-in firewall to control which hosts can connect to the SMTP VIP on port 25. The examples below show how the rules are constructed:

Example 1 – limit inbound SMTP connections to a specific smart host:

```
VIP1="192.168.30.10"
SRC1="192.168.30.50"
iptables -A INPUT -p tcp --src $SRC1 --dst $VIP1 --destination-port 25 -j ACCEPT
iptables -A INPUT -p tcp --dport 25 -j DROP
```

These rules will only allow SMTP traffic from the host 192.168.30.50 to reach the 192.168.30.10 VIP.

Example 2 – limit inbound SMTP connections to a range of smart hosts:

```
VIP1="192.168.30.10"
SRC1="192.168.30.50-192.168.30.60"
iptables -A INPUT -p tcp -m iprange --src-range $SRC1 --destination $VIP1 --destination-port 25 -j ACCEPT
iptables -A INPUT -p tcp --dport 25 -j DROP
```

These rules will only allow SMTP traffic from hosts in the range 192.168.30.50 through 192.168.30.60 to reach the 192.168.30.10 VIP.

To add firewall rules

Note

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

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

1. Using the WebUI, navigate to: *Maintenance > Firewall Script*.
2. Scroll down to the bottom of the script, add a descriptive comment for the rules, then copy & paste the appropriate example rules as shown in the example below:

```
54
55 # Lockdown SMTP inbound connections
56 VIP1="192.168.30.10"
57 SRC1="192.168.30.50"
58 iptables -A INPUT -p tcp --src $SRC1 --dst $VIP1 --destination-port 25 -j ACCEPT
59 iptables -A INPUT -p tcp --dport 25 -j DROP
60
61 echo "Firewall Activated"
62 exit 0;
63
```

Update

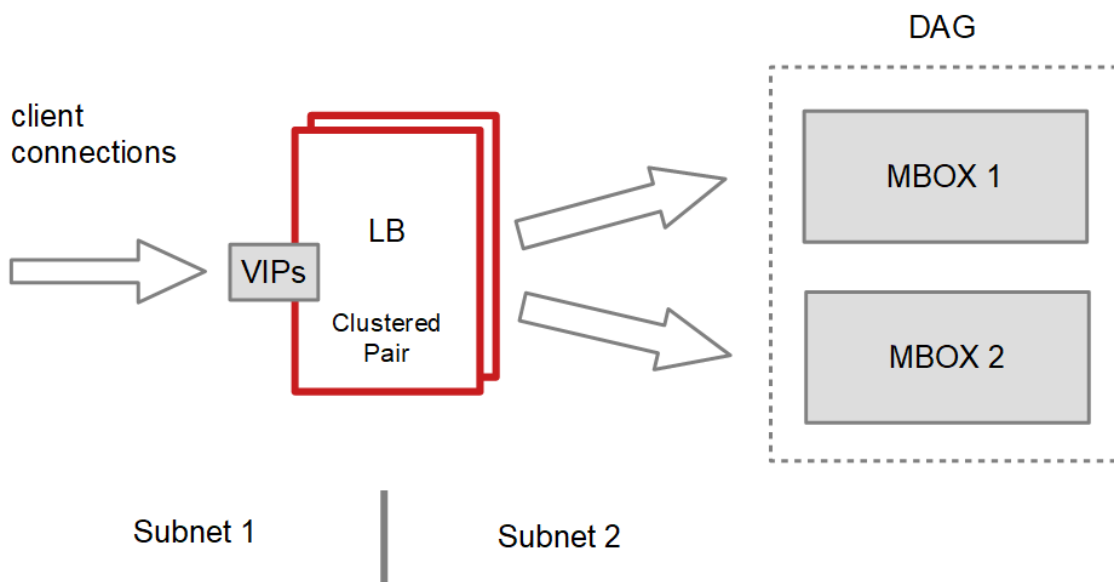
3. Insert a comment using the ' #' symbol, e.g. **# Lockdown SMTP inbound connections.**
4. Ensure that the IP addresses specified for VIP1 and SRC1 are configured for your environment.
5. Click **Update**.

Enabling Layer 7 Transparency using TPROXY

As mentioned previously, Layer 7 SNAT mode is not transparent by default. If a fully transparent configuration is required, TPROXY can be used.

Layer 7 SNAT mode with TProxy is typically used in a 2-arm configuration where the VIP is located in one subnet and the load balanced Real Servers are located in another. This can be achieved by using two network adapters, or by creating VLAN's on a single adapter. Single arm configuration is also supported under certain conditions - for more information please refer to [Transparency at Layer 7](#).

Using a 2-arm Configuration:



2-arm configuration - key points to note:

1. The Exchange Servers must be on a different subnet to the VIP – this can be achieved by using two network adapters, or by creating VLANs on a single adapter.
2. The default gateway on the Exchange Servers must be configured to be an IP address on the load balancer. For a clustered pair of load balancers, an additional floating IP should be used for this purpose to allow failover to the Secondary.

To enable TProxy:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Service*.
2. Click **Modify** next to the virtual service in question.
3. Scroll down to the *Other* section and click [**Advanced**].
4. Check the **Transparent Proxy** checkbox.
5. Click **Update**.

Note

If the load balancer has been deployed in Layer 4 DR mode, this is transparent by default so no additional steps are required. This section only applies when Layer 7 SNAT mode was initially used and transparency is now required.

Using a Layer 4 Virtual Service for SMTP

Layer 7 Virtual Services are not transparent by default which can be an issue for the HT role. One option in this case is to use a Layer 4 DR mode VIP. For more details about Layer 4 DR mode please refer to [Layer 4 DR Mode](#).

Note

If the load balancer has been deployed in Layer 4 DR mode, this is transparent by default so no additional steps are required. This section only applies when Layer 7 SNAT mode was initially used and transparency is now required.

Layer 4 DR Mode – Solving the ARP Problem:

Layer 4 DR mode works by changing the MAC address of the inbound packets to match the Real Server selected by the load balancing algorithm. To enable DR mode to operate:

- Each Real Server must be configured to accept packets destined for both the VIP address and the Real Servers IP address (RIP). 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). The service/process (e.g. IIS) must respond to both addresses.
- Each Real Server must 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 particular version of Windows being used. For detailed steps on solving the ARP problem for Windows 2012 and later please refer to [Solving the ARP Problem](#).

Configuring an HTTP to HTTPS redirect for OWA

An additional layer 7 VIP is required that listens on HTTP port 80 on the same IP address. The VIP is then configured to redirect connections to HTTPS port 443.

e.g. <http://mail.robstest.com/owa> should be redirected to <https://mail.robstest.com/owa>

1) Create another Layer 7 VIP with the following settings:

- *Label*: **HTTP-redirect**
- Virtual Service IP Address: <same as the VIP that's listening on port 443>
- Virtual Service Ports: **80**
- *Layer 7 Protocol*: **HTTP Mode**
- Persistence Mode: **None**
- Force to HTTPS: **Yes**

Note

This additional VIP will be shown purple/green to indicate that it's being used for HTTP to HTTPS redirection.

2) Apply the new settings – to apply the new settings, HAProxy must be restarted:

- Using the WebUI, navigate to: *Maintenance > Restart Services* and click **Restart HAProxy**

Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution.

We recommend that the Primary appliance should be configured first, then the Secondary should be added. Once the Primary and Secondary are paired, all load balanced services configured on the Primary are automatically replicated to the Secondary over the network using SSH/SCP.

Note

For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

The clustered HA pair uses Heartbeat to determine the state of the other appliance. Should the active device (normally the Primary) suffer a failure, the passive device (normally the Secondary) will take over.

Note

A number of settings are not replicated as part of the Primary/Secondary pairing process and therefore must be manually configured on the Secondary appliance. These are listed by WebUI menu option in the table below:

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	All network settings including IP address(es), bonding configuration and VLANs
Local Configuration	Routing	Routing configuration including default gateways and static routes
Local Configuration	System Date & time	All time and date related settings
Local Configuration	Physical – Advanced Configuration	Various settings including Internet Proxy, Management Gateway, Firewall connection tracking table size, NIC offloading, SMTP relay, logging and Syslog Server
Local Configuration	Security	Appliance security settings
Local Configuration	SNMP Configuration	Appliance SNMP settings
Local Configuration	Graphing	Appliance graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Software Updates	Appliance software update management
Maintenance	Firewall Script	Appliance firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

To add a Secondary node - i.e. create a highly available clustered pair:

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

CREATE A CLUSTERED PAIR

The screenshot shows the 'CREATE A CLUSTERED PAIR' configuration page. On the left, there is a header with a circular logo and the text 'loadbalancer.org'. The main area contains three input fields: 'Local IP address' with a dropdown menu showing '192.168.110.40', 'IP address of new peer' with a text box containing '192.168.110.41', and 'Password for loadbalancer user on peer' with a masked password field. A green button labeled 'Add new node' is positioned at the bottom right.

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

CREATE A CLUSTERED PAIR

The screenshot shows the 'CREATE A CLUSTERED PAIR' configuration page during the pairing process. The top header remains the same. Below it, a status bar shows a 'P' icon, the IP '192.168.110.40', and 'loadbalancer.org'. In the center, there is a circular loading icon and the text 'Attempting to pair..'. Below this, another status bar shows an 'S' icon, the IP '192.168.110.41', and 'loadbalancer.org'. The input fields for 'Local IP address' and 'IP address of new peer' are still visible, with the same values as in the previous screenshot. The 'Password for loadbalancer user on peer' field is also visible. At the bottom, the green button is now grey and labeled 'configuring'.

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

High Availability Configuration - primary



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

Note

Clicking the **Restart Heartbeat** button on the Primary appliance will also automatically restart heartbeat on the Secondary appliance.

Note

For more details on configuring HA with 2 appliances, please refer to [Appliance Clustering for HA](#).

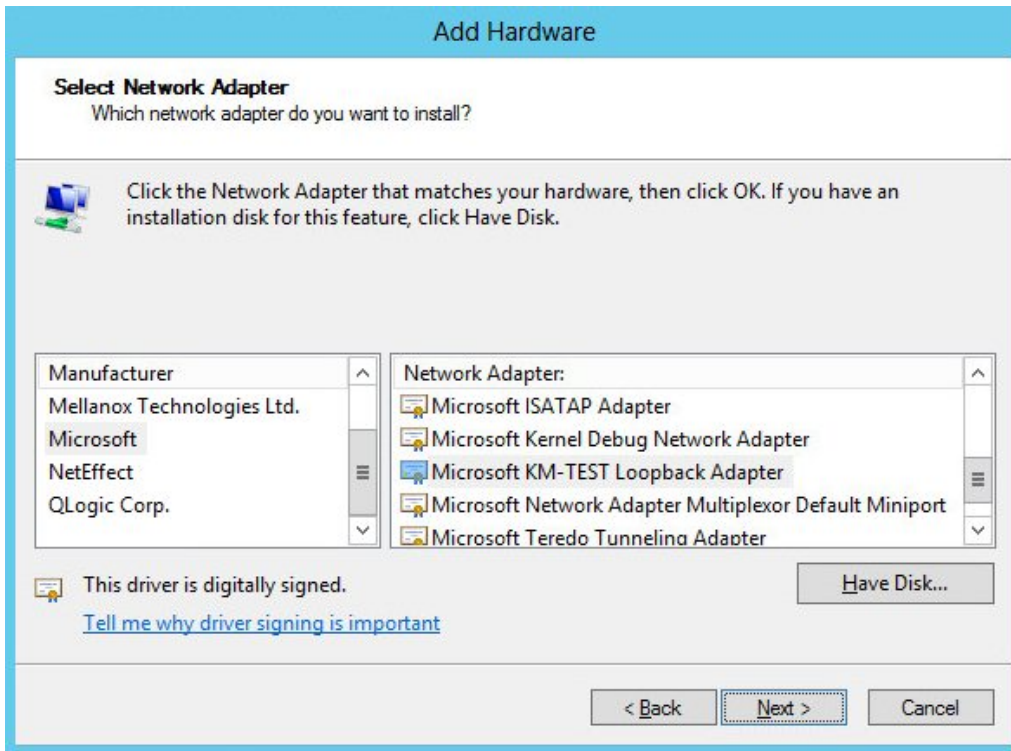
Solving the ARP Problem

Windows Server 2012 & Later

Windows Server 2012 and later support Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter. The IP address allocated to the Loopback Adapter must be the same as the Virtual Service (VIP) address. If the Real Server is included in multiple DR mode VIPs, additional IP addresses can be added to the Loopback Adapter that correspond to each VIP. In addition, steps must be taken to set the strong/weak host behavior which is used to either block or allow interfaces to receive packets destined for a different interface on the same server.

Step 1 of 3: Install the Microsoft Loopback Adapter

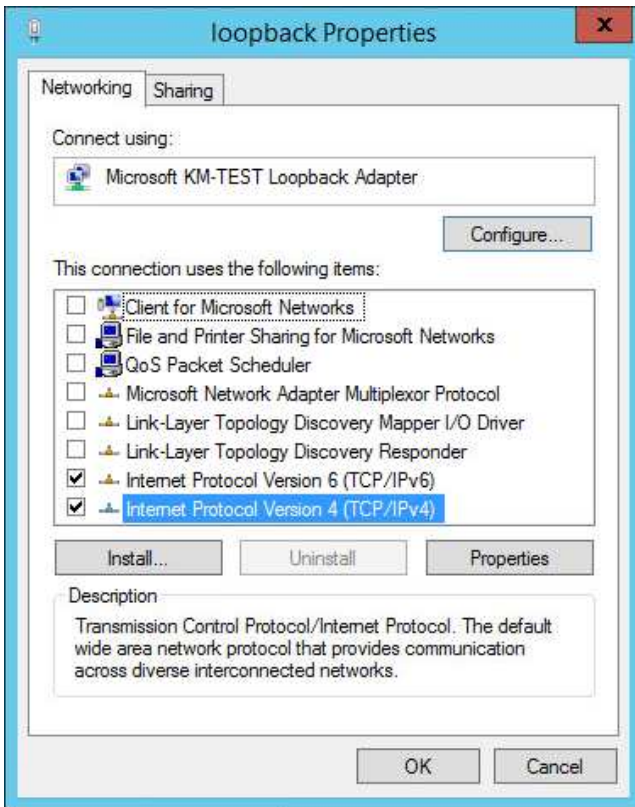
1. Click **Start**, then run `hdwwiz` to start the Hardware Installation Wizard.
2. When the Wizard has started, click **Next**.
3. Select **Install the hardware that I manually select from a list (Advanced)**, click **Next**.
4. Select **Network adapters**, click **Next**.
5. Select **Microsoft & Microsoft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



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

Step 2 of 3: Configure the Loopback Adapter

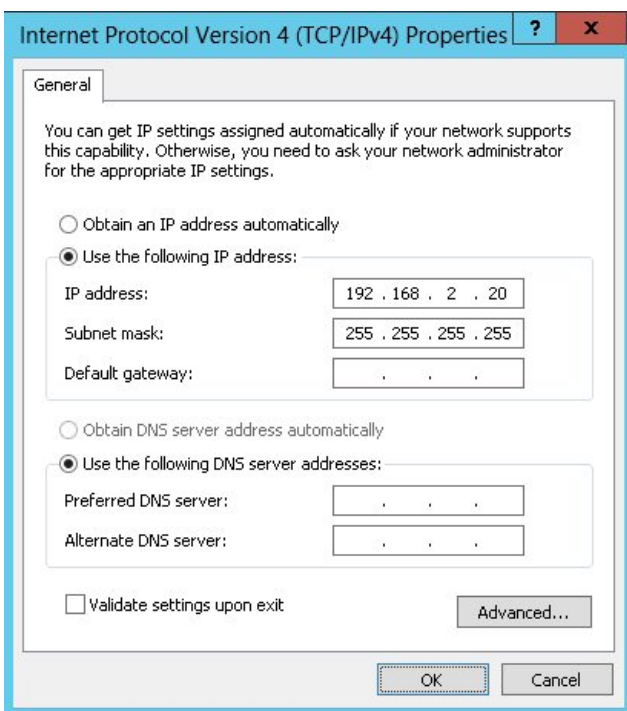
1. Open Control Panel and click **Network and Sharing Center**.
2. Click **Change adapter settings**.
3. Right-click the new Loopback Adapter and select **Properties**.
4. Uncheck all items except **Internet Protocol Version 4 (TCP/IPv4)** and **Internet Protocol Version 6 (TCP/IPv6)** as shown below:



Note

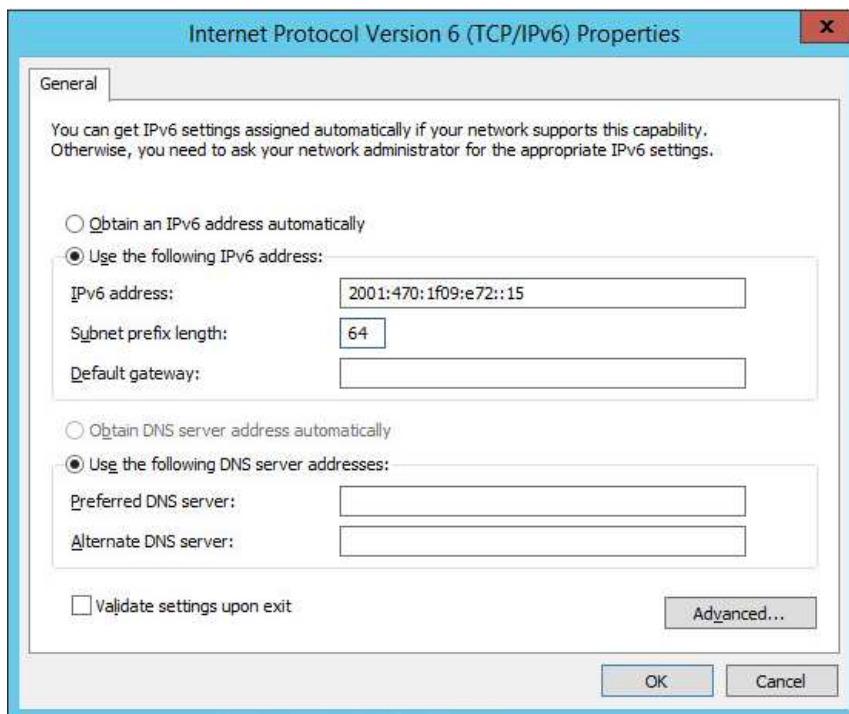
Leaving both checked ensures that both IPv4 and IPv6 are supported. Select one if preferred.

5. If configuring IPv4 addresses select **Internet Protocol Version (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:



6. If configuring IPv6 addresses select **Internet Protocol Version (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.

Note

For Windows 2012/2016/2019, it's not necessary to modify the interface metric on the advanced tab and should be left set to Automatic.

Step 3 of 3: Configure the strong/weak host behavior

To configure the correct strong/weak host behavior for Windows 2012/2016/2019, 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
```



Note

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.

- Start PowerShell or use a command window to run the appropriate netsh commands as shown in the example below:



Note

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

Repeat steps 1 - 3 on all remaining Windows 2012/2016/2019 Real Server(s).

If preferred you can also use the following PowerShell Cmdlets:

The following example configures both IPv4 and IPv6 at the same time:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled -DadTransmits 0
```

```
Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled
```

To configure just IPv4:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled  
-DadTransmits 0 -AddressFamily IPv4
```

```
Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled -AddressFamily IPv4
```

To configure just IPv6:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled  
-DadTransmits 0 -AddressFamily IPv6
```

```
Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled -AddressFamily IPv6
```

Update the Network Adapter Priority Order

To ensure that that newly added loopback adapter has no effect on which interface Windows attempts to use, it's important that the loopback adapter has the lowest priority. In Windows Server 2016 and later, you can use the interface metric to [configure the order of network interfaces](#). As mentioned [here](#), the interface metric can be viewed and configured using either PowerShell or via the Windows GUI.

To check the current interface metric for all adapters using PowerShell:

Note | Perform these steps on ALL mailbox servers.

1. Open a PowerShell command window and run the following command:

```
Get-NetIPinterface
```

Output similar to the following will be displayed:

```
Administrator: Windows PowerShell  
-----  
IfIndex InterfaceAlias AddressFamily NI\Mtu(Bytes) InterfaceMetric Dhcp ConnectionState PolicyStore  
-----  
5 Local Area Connection* 10 IPv6 1300 15 Enabled Connected Activ...  
29 loopback IPv6 1500 25 Enabled Connected Activ...  
9 net IPv6 1500 25 Enabled Connected Activ...  
1 Loopback Pseudo-Interface 1 IPv6 4294967295 75 Disabled Connected Activ...  
5 Local Area Connection* 10 IPv4 1300 15 Disabled Connected Activ...  
29 loopback IPv4 1500 25 Enabled Connected Activ...  
9 net IPv4 1500 25 Disabled Connected Activ...  
1 Loopback Pseudo-Interface 1 IPv4 4294967295 75 Disabled Connected Activ...  
PS C:\Users\administrator.LBTESTDOMAIN>
```

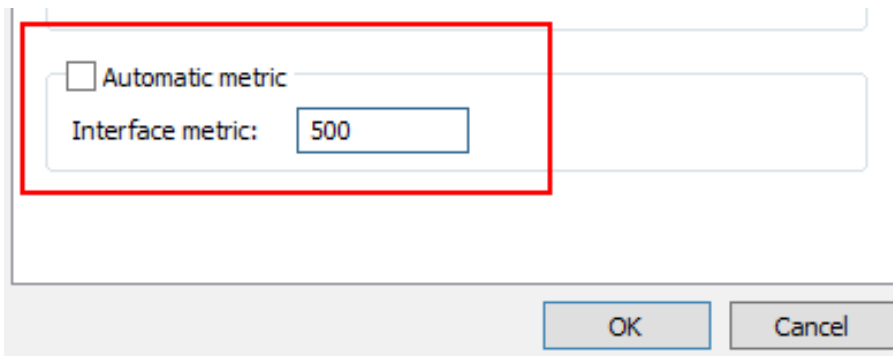
Note | The interface metric is displayed in the 5th column

In the above example, the 'loopback' and 'net' adapters have the same interface metric (25). To ensure that there is

no possibility of issues occurring, the loopback adapter should be modified so that it has a higher interface metric, and is therefore a lower priority (see below).

To configure the loopback adapter's interface metric using the Windows GUI:

1. Open the Properties of the loopback adapter, select the required IP version (if IPv4 and IPv6 are needed, repeat these steps for both), click **Properties**, then click **Advanced**.
2. Uncheck the *Automatic Metric* checkbox, then enter a suitable value to ensure that the loopback adapter has the highest value, e.g. 500 as shown below.



3. Click **OK**, **OK** and **Close** to apply the new settings.

18. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.3.0	6 August 2019	Styling and layout	General styling updates	RJC
1.3.1	17 January 2020	Added note explaining how to disable "Secure Mode" to unlock the firewall script page	Required update	RJC
1.3.2	2 May 2020	Minor correction	Required update	RJC
1.3.3	3 June 2020		New title page Updated Canadian contact details Branding update Change to Canadian contact details	AH
1.3.4	9 February 2021	Minor update	Required update	RJC
1.3.5	25 June 2021	Minor updates	Required update	RJC
1.4.0	1 December 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.4.1	14 April 2022	Updated TPROXY instructions	Changes to the appliance WebUI	AH
1.4.2	22 April 2022	Updated SSL related content to reflect latest software version	New software release	RJC

About Loadbalancer.org

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



United Kingdom

Loadbalancer.org Ltd.
Compass House, North Harbour
Business Park, Portsmouth, PO6 4PS
UK: +44 (0) 330 380 1064
sales@loadbalancer.org
support@loadbalancer.org

Canada

Loadbalancer.org Appliances Ltd.
300-422 Richards Street, Vancouver,
BC, V6B 2Z4, Canada
TEL: +1 866 998 0508
sales@loadbalancer.org
support@loadbalancer.org

United States

Loadbalancer.org, Inc.
4550 Linden Hill Road, Suite 201
Wilmington, DE 19808, USA
TEL: +1 833.274.2566
sales@loadbalancer.org
support@loadbalancer.org

Germany

Loadbalancer.org GmbH
Tengstraße 2780798,
München, Germany
TEL: +49 (0)89 2000 2179
sales@loadbalancer.org
support@loadbalancer.org