LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing Microsoft Exchange 2016

Version 1.5.1

Table of Contents

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/enterprise.

Some features may not be available or fully supported in all cloud platforms due to platform specific limitations. For more details, please refer to the "Main Differences to our Standard (Non-Cloud) Product" section in the appropriate cloud platform Quick Start Guide or check with Loadbalancer.org support.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.9.1 and later

> **Note** The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If you're using an older version, or the very latest, the screenshots presented here may not match your WebUI exactly.

3.2. Microsoft Exchange

- All versions

4. 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 **Note** Exchange 2016 has since been superseded by Exchange 2019.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.

5. 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

6. Load Balancing Exchange 2016

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

6.1. 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.

6.1.1. 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.

6.2. Persistence (aka Server Affinity)

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

6.3. 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) *Note: Outlook Web App has been renamed as Outlook on the Web in Exchange 2016*
993	MBOX	Secure IMAP4 clients
995	MBOX	Secure POP3 clients

6.4. 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-encypted 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).

6.5. 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

Namespace	Purpose
autodiscover.lbtestdom.com	Auto Discover

Notes

1. In this case a single VIP is used for all HTTPS namespaces/services.
2. Both DNS entries should then point at the same VIP.
3. 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

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

6.6. 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.

6.7. 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.

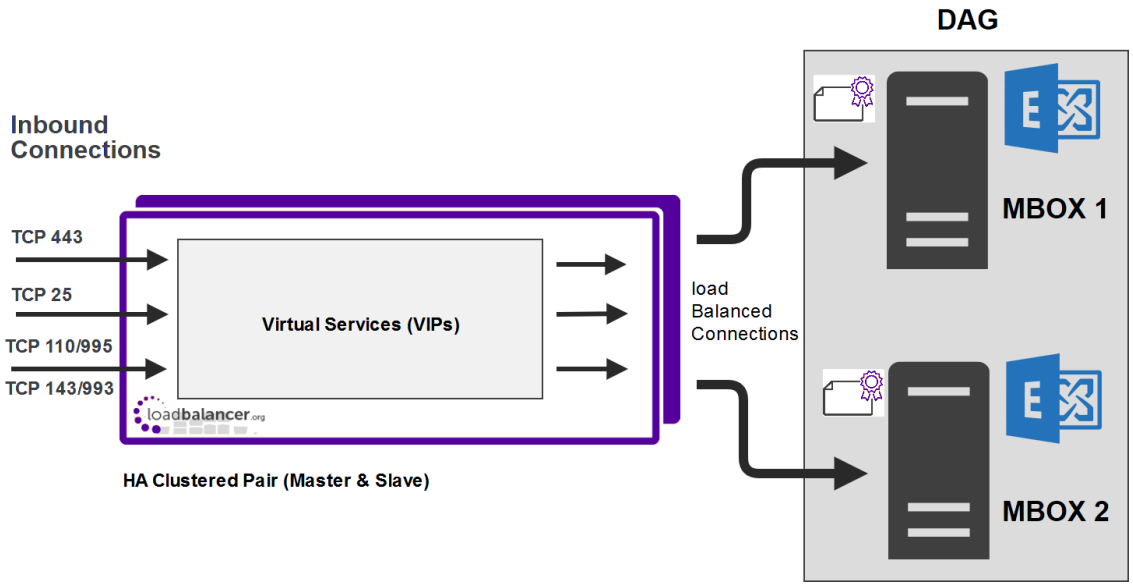

> **Note** The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to Configuring HA - Adding a Secondary Appliance for more details on configuring a clustered pair.

6.8. 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.

6.9. 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.

6.9.1. 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. The image below shows an example network diagram for this mode.

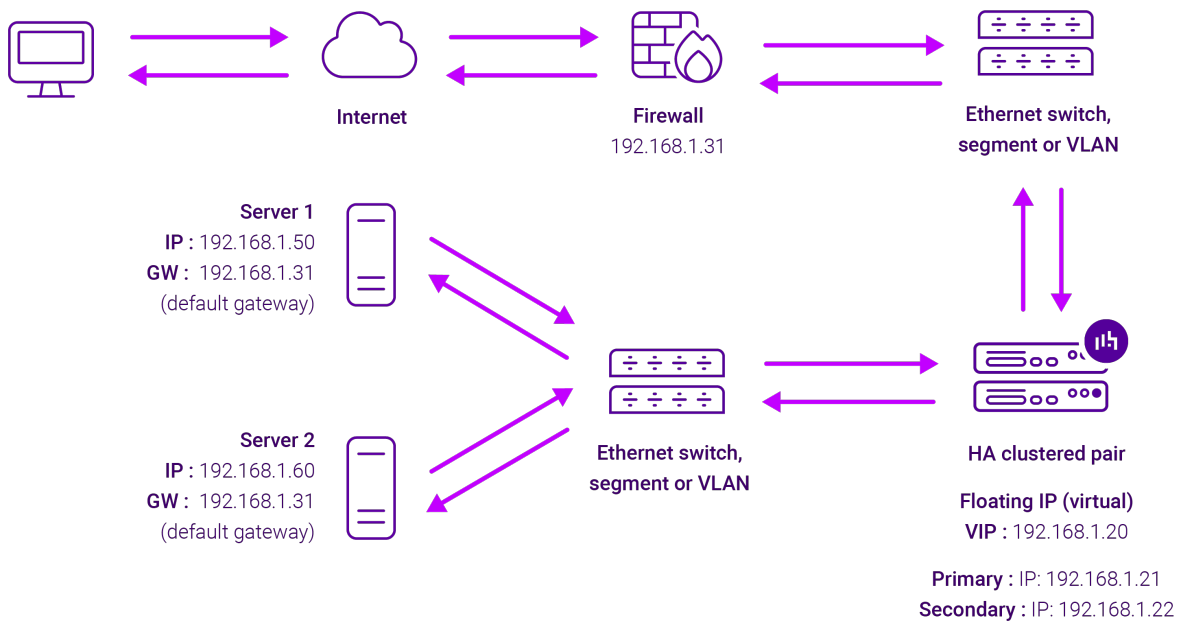

- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's 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, **eth1** is typically used for client side connections and **eth0** is used for Real Server connections, although this is not mandatory since any interface can be used for any purpose.
- Requires no mode-specific 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.

6.9.2. Layer 4 DR Mode

Layer 4 DR (Direct Routing) mode is a very high performance solution that requires little change to your existing

infrastructure. The image below shows an example network diagram for this mode.

> **Note** Kemp, Brocade, Barracuda & A10 Networks call this ***Direct Server Return*** and F5 call it ***nPath***.

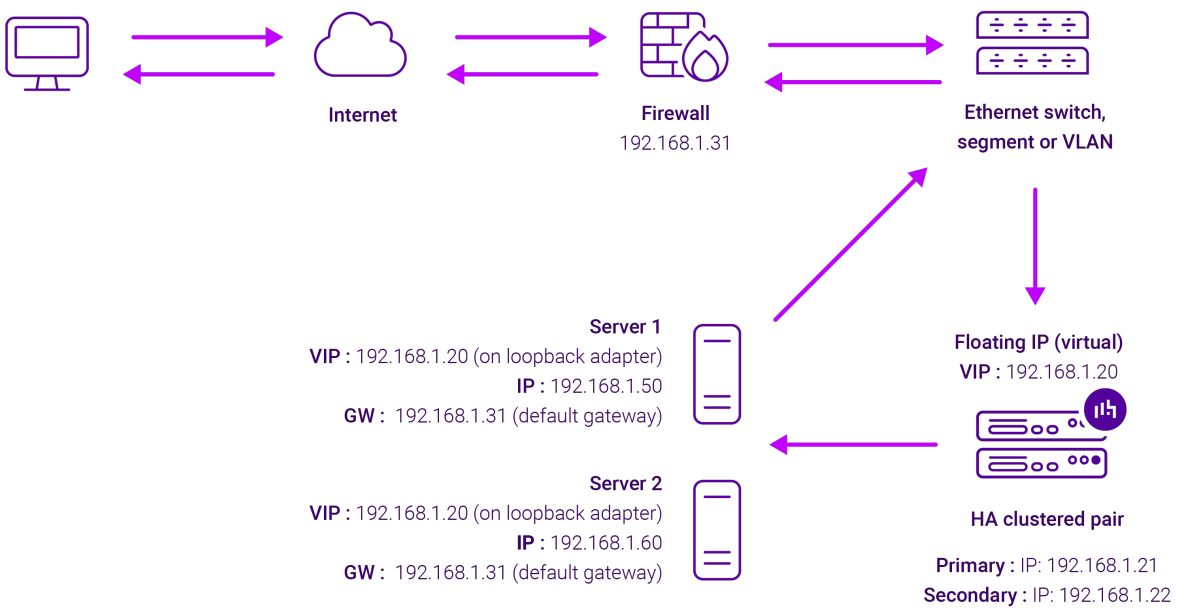

- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that each Real Server (and the load balanced application) must respond to both the Real Server's own IP address and the VIP.
- The Real Server should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Server 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 mode for HTTP and much faster for other applications such as Remote Desktop Services, streaming media and FTP.
- The load balancer must have an interface in the same subnet as the Real Servers to ensure layer 2 connectivity which is required for DR mode to operate.
- 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.

6.10. 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.

6.10.1. 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-encypted 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).

Note

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.

7. Configuring Exchange 2016 for Load Balancing

7.1. 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**

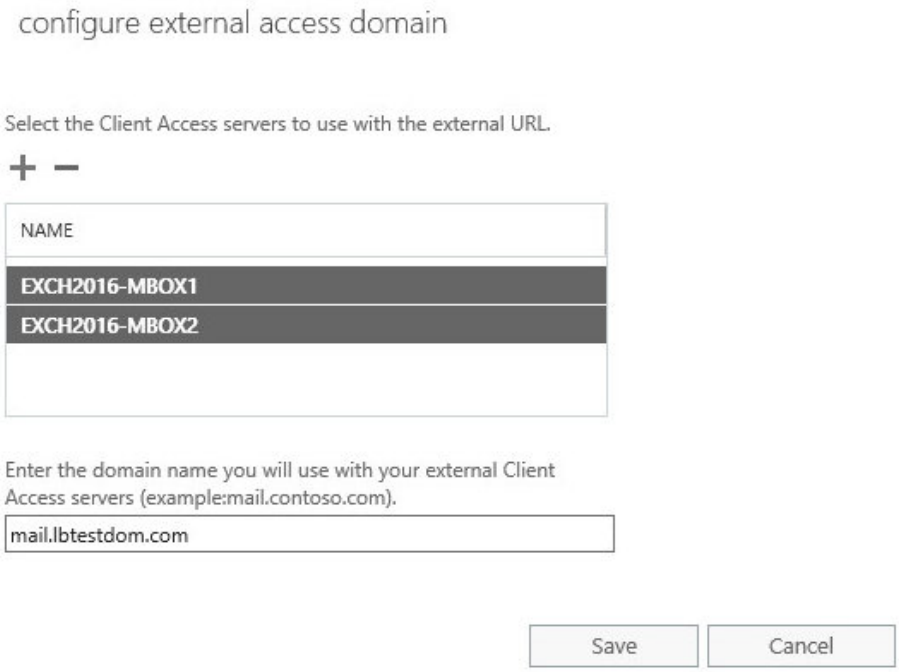

7.2. Virtual Directories

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:

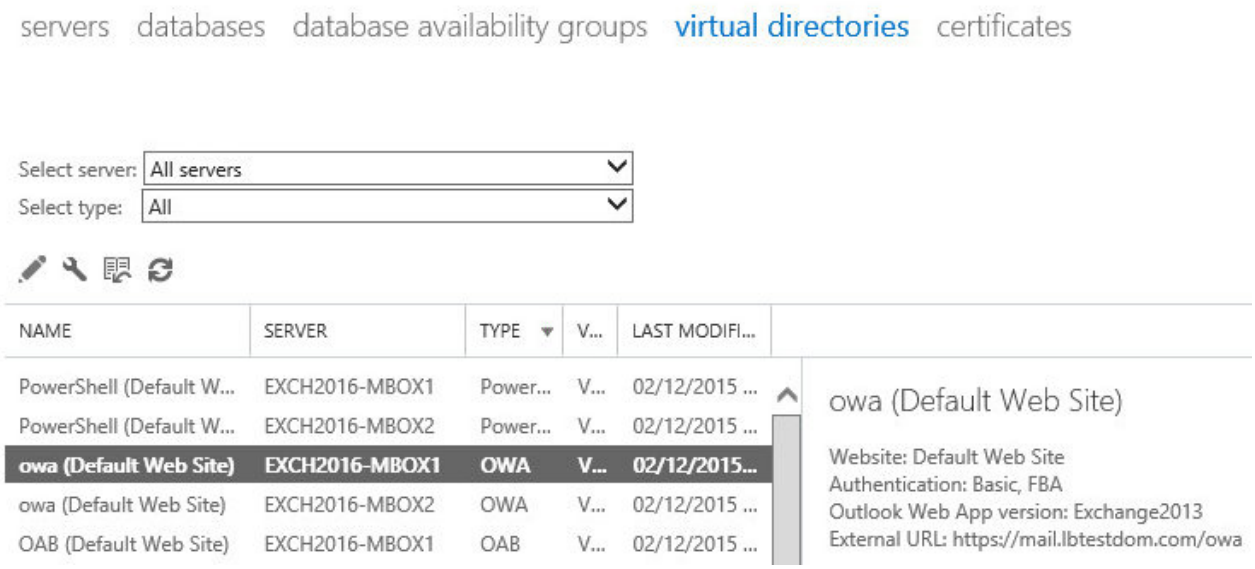

7.3. Outlook Anywhere

This is configured using the EAC. Select ***servers > servers*** and then click the edit (pen) icon next to each server, 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**.

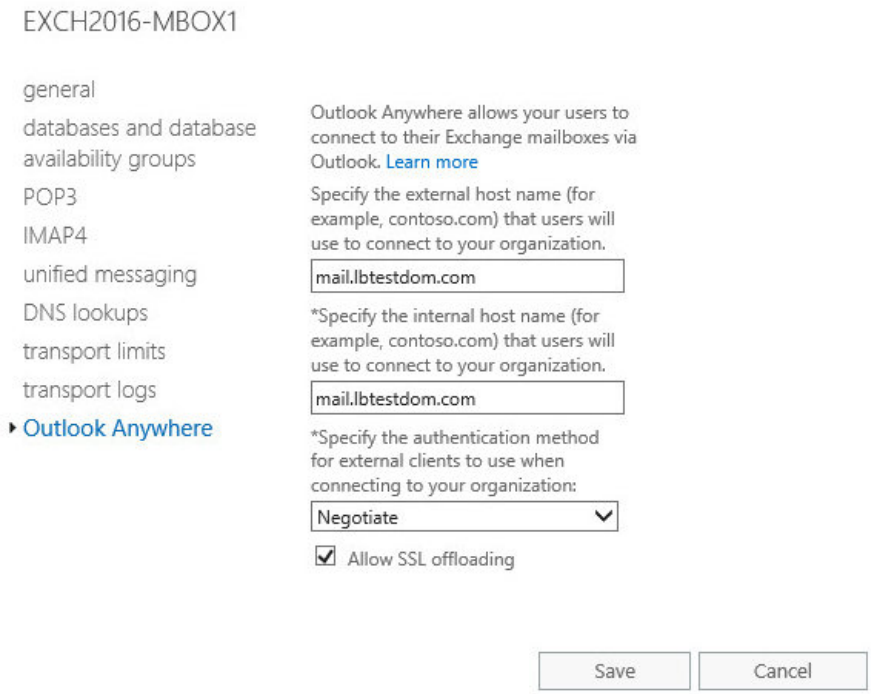

7.4. Autodiscover

7.4.1. 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.

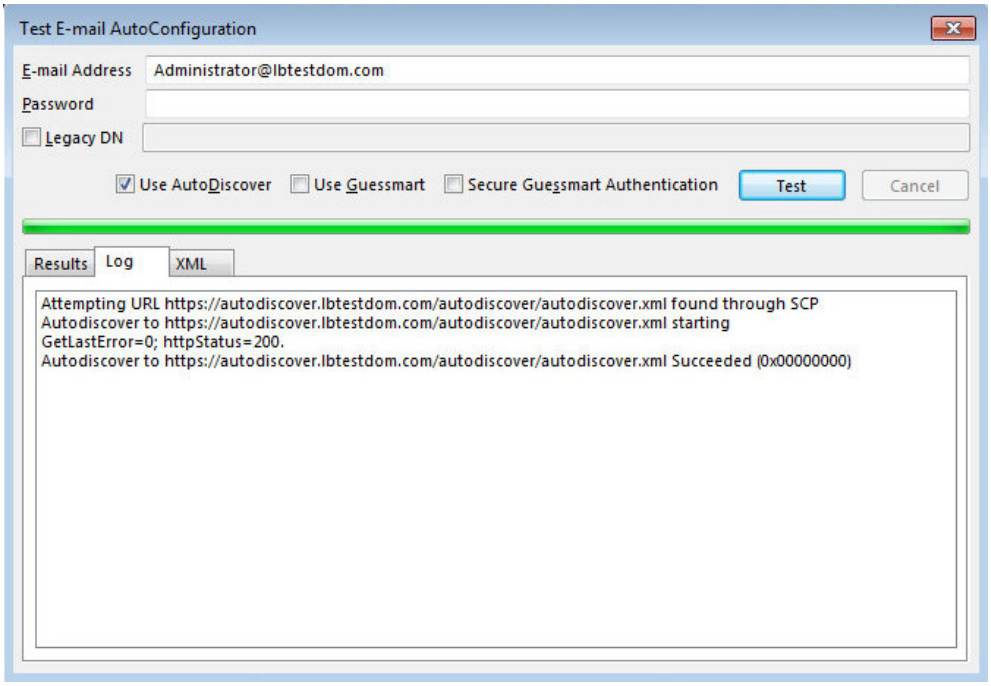

7.4.2. External

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:

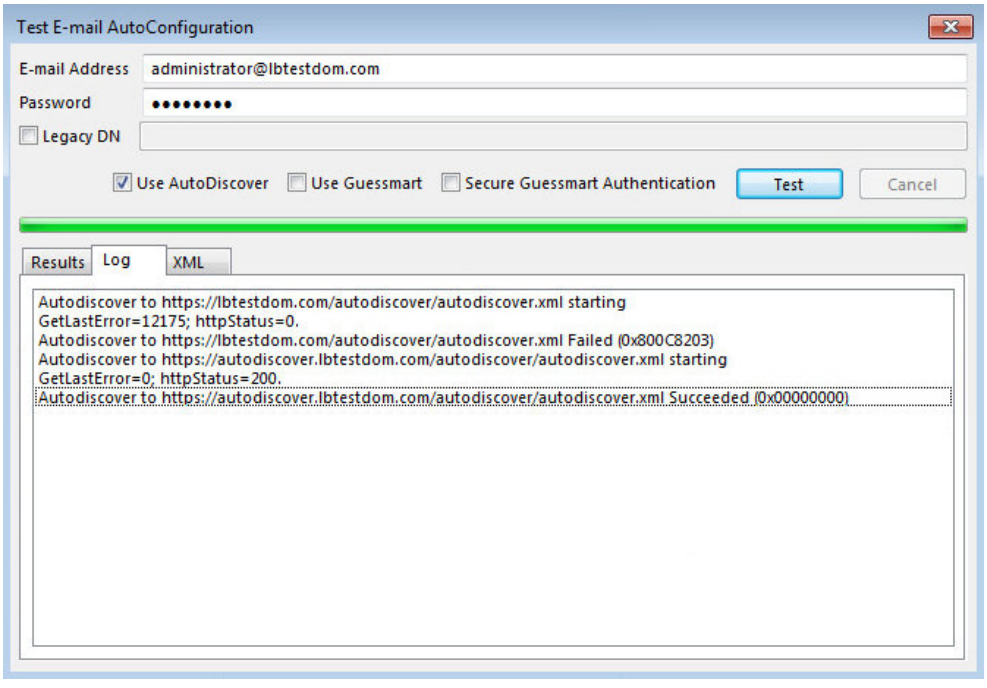

7.5. Certificates

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.***

7.6. 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:

https://technet.microsoft.com/en-us/library/jj657461(v=exchg.160).aspx

7.6.1. 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 created, this is typically done to enables multiple outbound email routes to be specified that have different costs.

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

7.7. 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.

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

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

7.8.1. SNAT Mode

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

7.8.2. 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.

7.9. 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
```

8. Loadbalancer.org Appliance – the Basics

8.1. Virtual Appliance

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

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

> **Note** Please refer to Virtual Appliance Installation and the ReadMe.txt text file included in the VA

download for additional information on deploying the VA using the various Hypervisors.

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

8.2. Initial Network Configuration

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

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

8.3. Accessing the Appliance WebUI

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

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

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

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

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

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

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

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

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

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

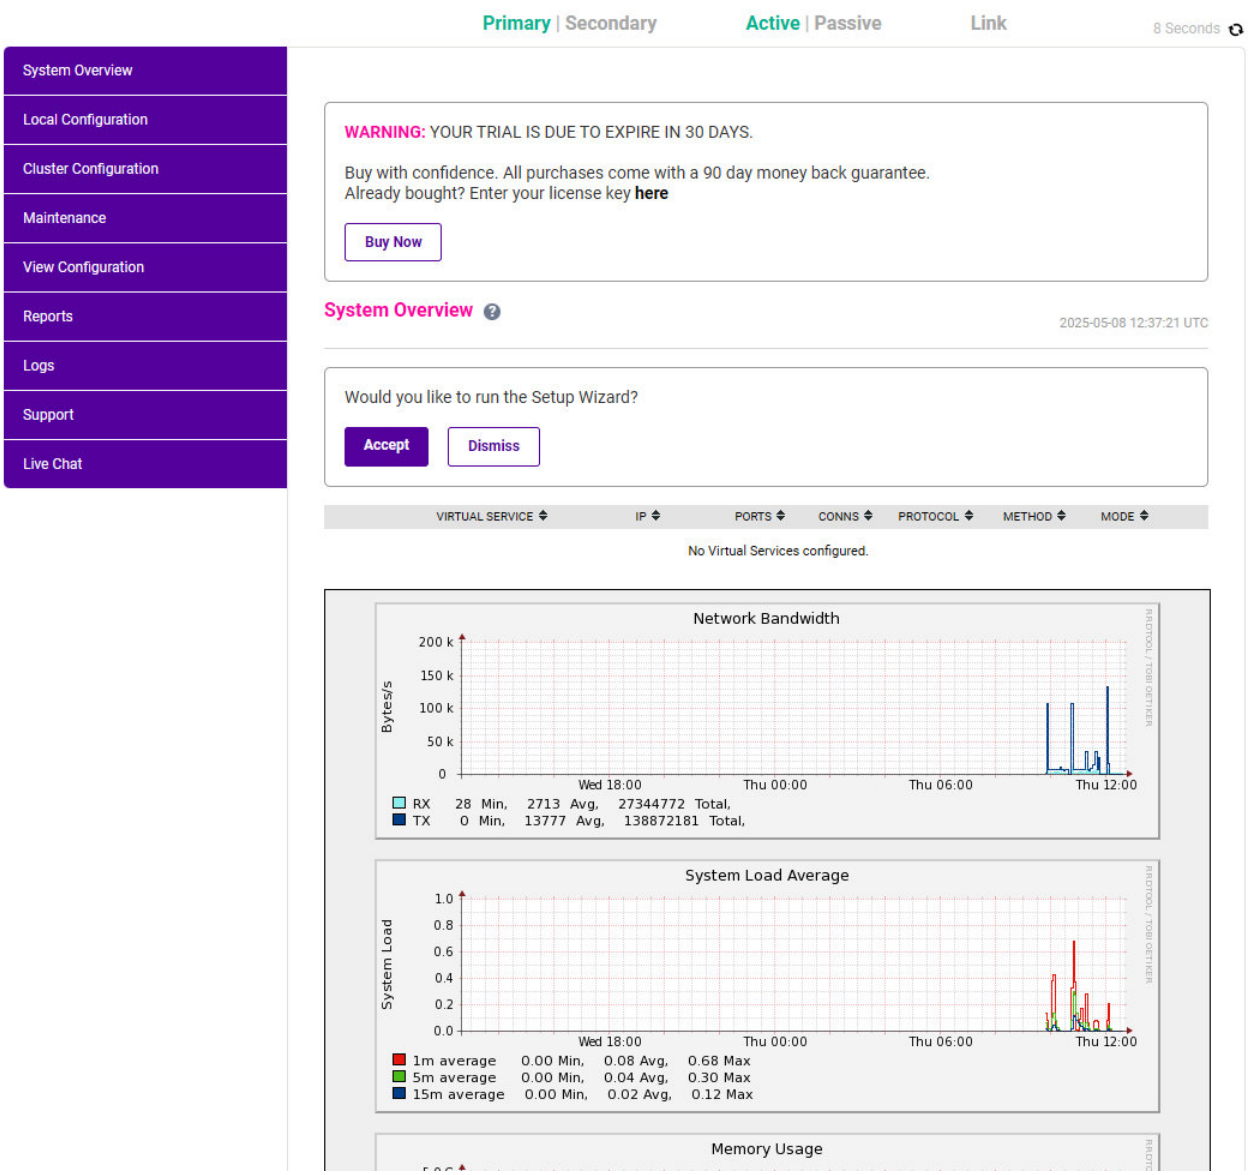

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

> **Note** The Setup Wizard can only be used to configure Layer 7 services.

8.3.1. 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 creating 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

8.4. Appliance Software Update

We recommend that the appliance is kept up to date to ensure that you benefit from the latest bug fixes, security updates and feature improvements. Both online and offline update are supported.

> **Note** For full details, please refer to Appliance Software Update in the Administration Manual.

> **Note** Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

8.4.1. Online Update

The appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. This is the default behavior and can be disabled if preferred. If an update is found, a notification similar to the example below will be displayed at the top of the WebUI:

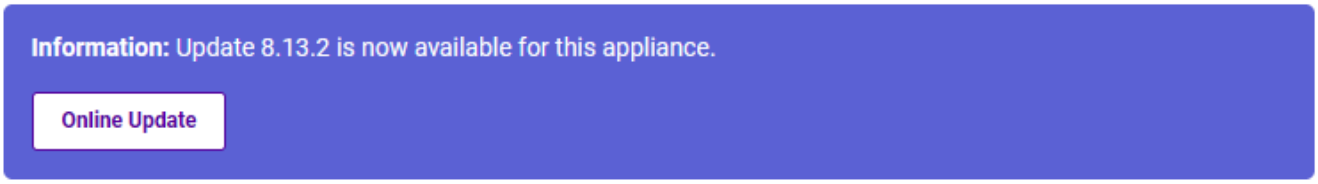

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

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

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

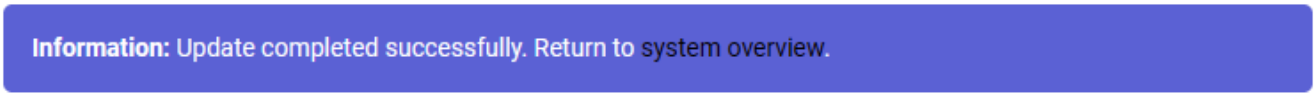

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

8.4.2. Offline Update

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

To check for the latest version, please refer to our product roadmap page available here. To obtain the latest offline update files contact support@loadbalancer.org.

To perform an offline update:

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

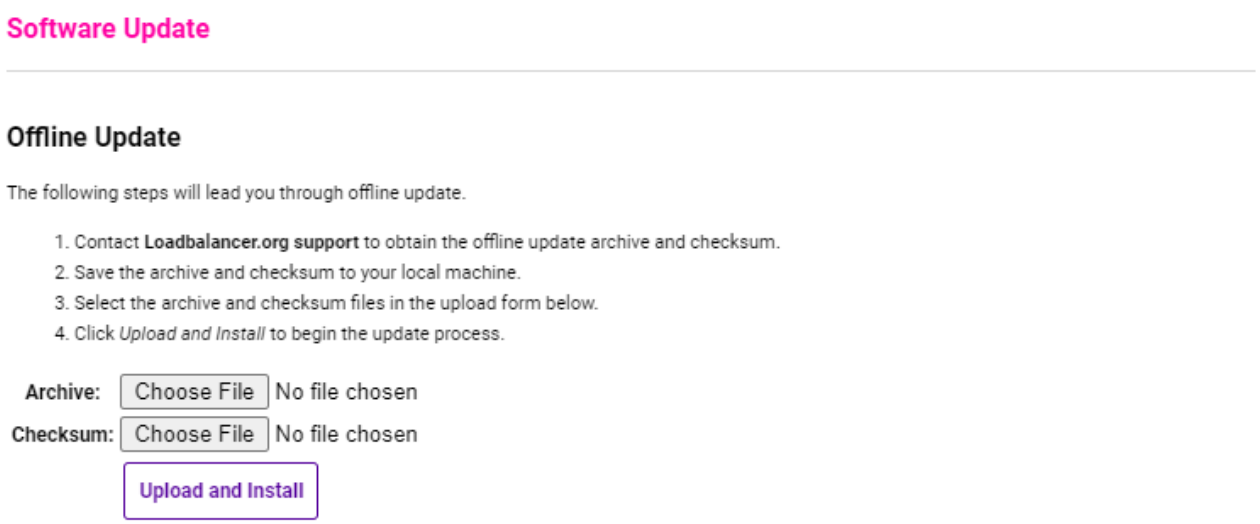

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

8.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22 *	SSH
TCP & UDP	53 *	DNS / GSLB
TCP & UDP	123	NTP
TCP & UDP	161 *	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000 *	Gateway service (Centralized/Portal Management)
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

> **Note** The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to Service Socket

Addresses.

8.6. HA Clustered Pair Configuration

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

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

9.1. Load Balancer Deployment Overview

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

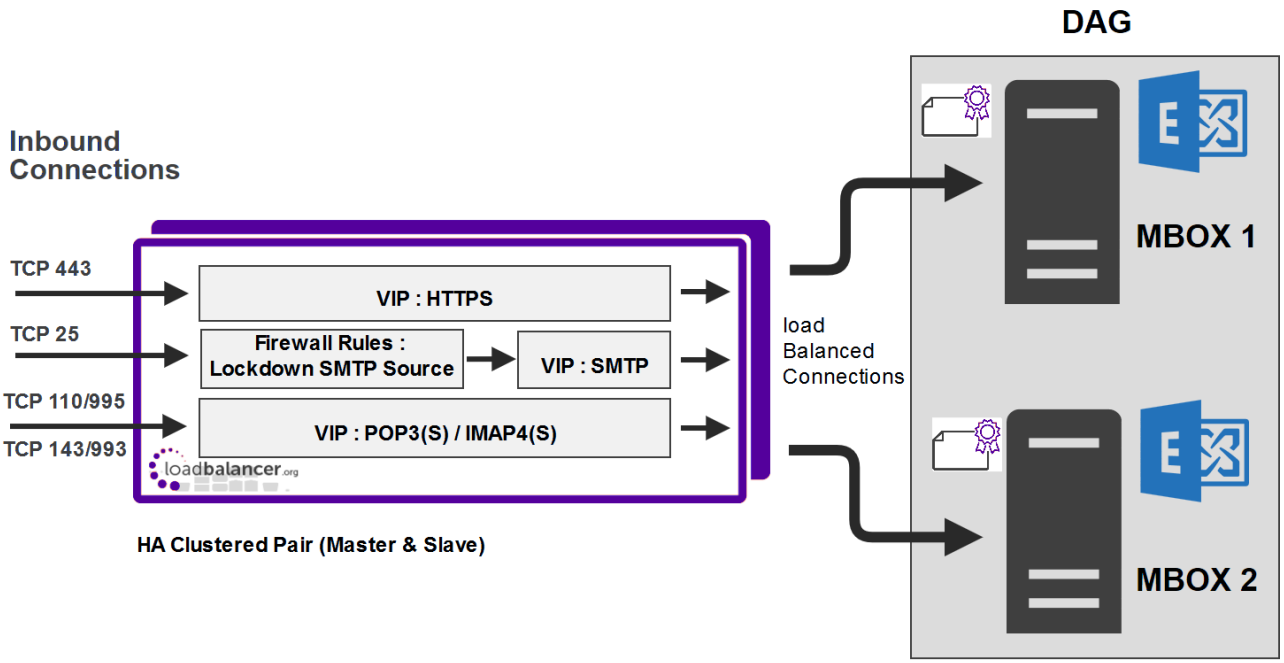

Notes

- 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.

9.2. Load Balancer Configuration

9.2.1. 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:

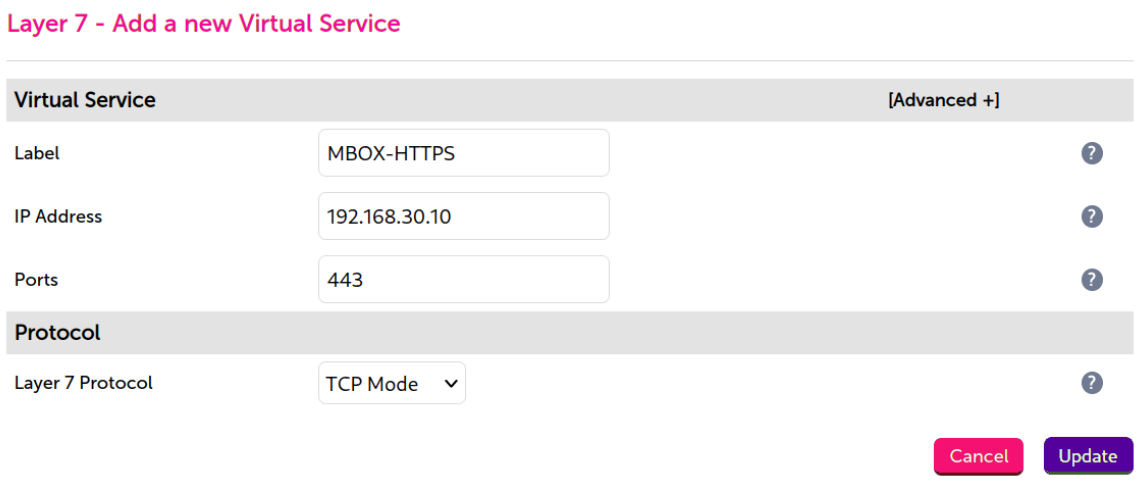

3. Enter an appropriate label for the VIP, e.g. **MBOX-HTTPS**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **443**.
6. Set *Layer 7 Protocol* set to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** 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**.
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. Set the *Response expected* to **200 OK**.
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. 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:

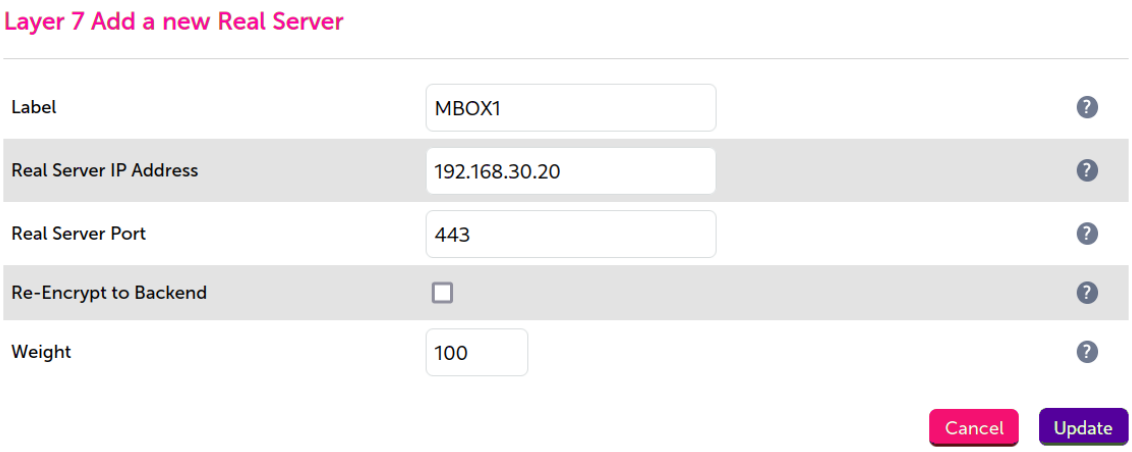

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **443**.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

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.

9.2.2. 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:

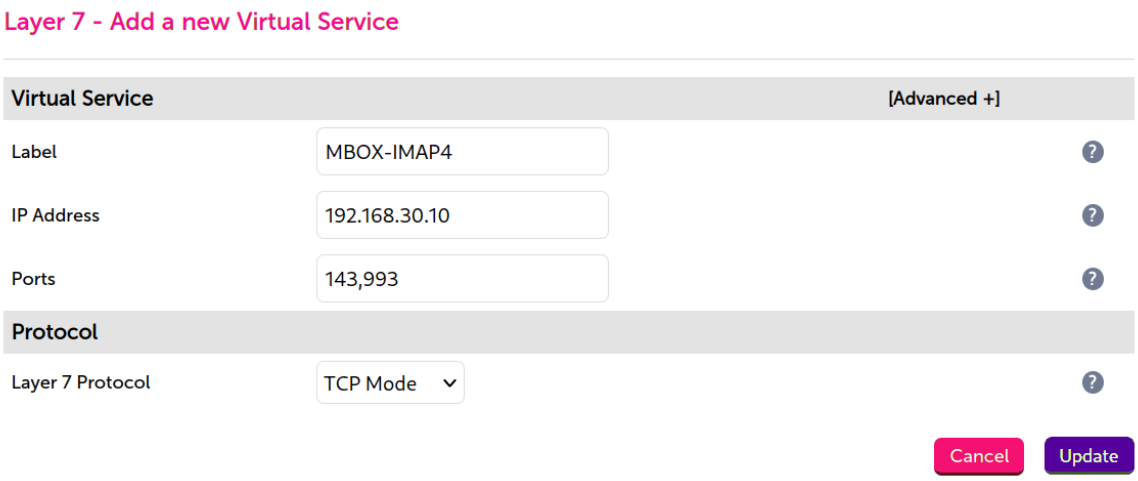

3. Enter an appropriate label for the VIP, e.g. **MBOX-IMAP4**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **143,993**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify**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

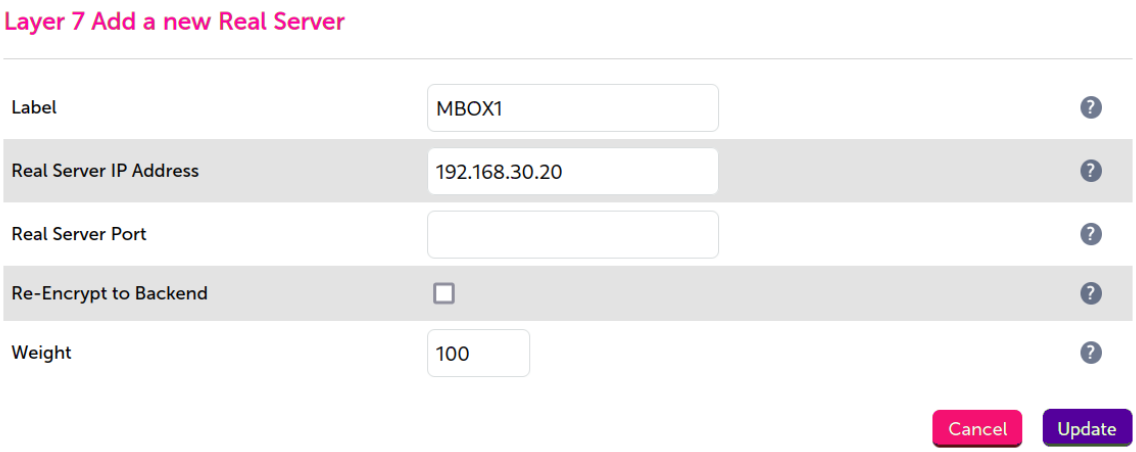

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

9.2.3. 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:

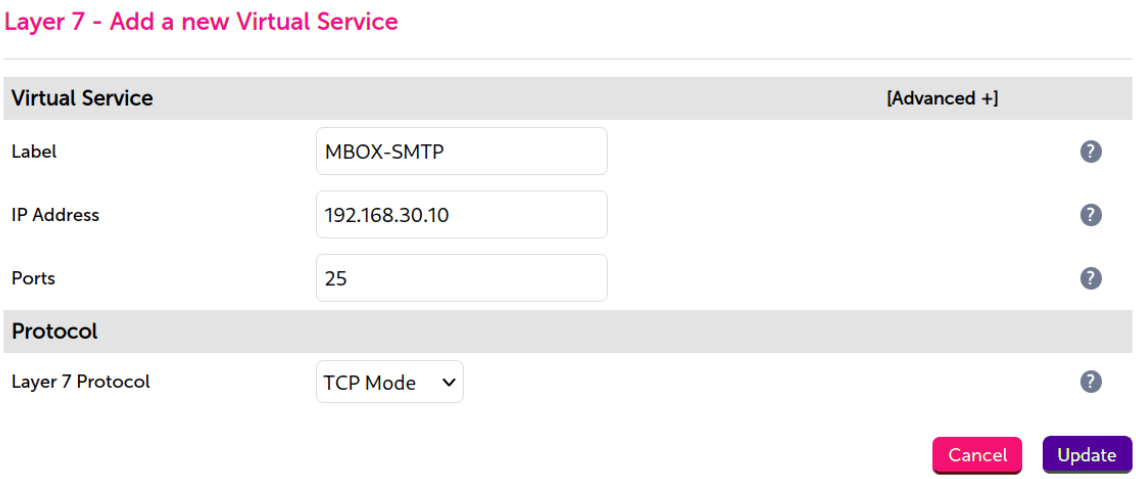

3. Enter an appropriate label for the VIP, e.g. **MBOX-SMTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **25**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** 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:

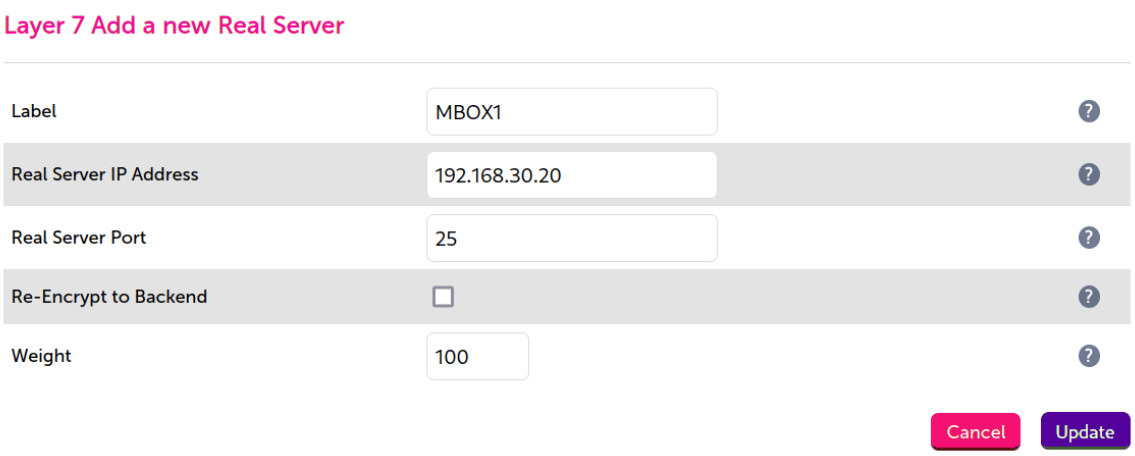

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **25**.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

9.2.4. 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.

9.2.5. 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).

9.2.6. Finalizing the Configuration

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the "Commit changes" 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**.

9.3. Exchange Server Configuration Steps

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

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

10.1. 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.

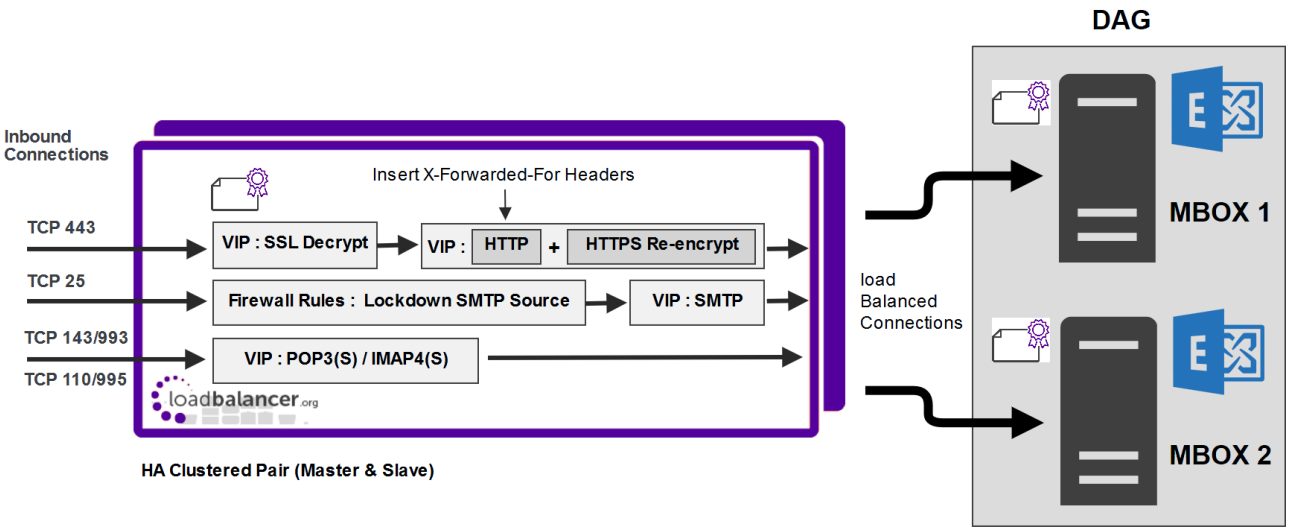

Notes

- 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.

10.2. Load Balancer Configuration

10.2.1. 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:

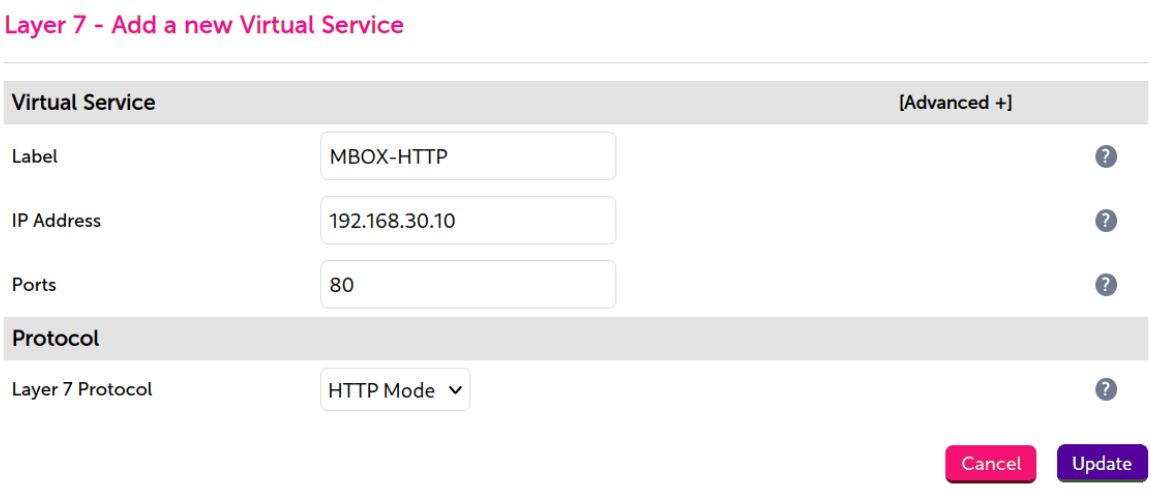

3. Enter an appropriate label for the VIP, e.g. **MBOX-HTTP**.

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. Set the *Response expected* to **200 OK**.
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:

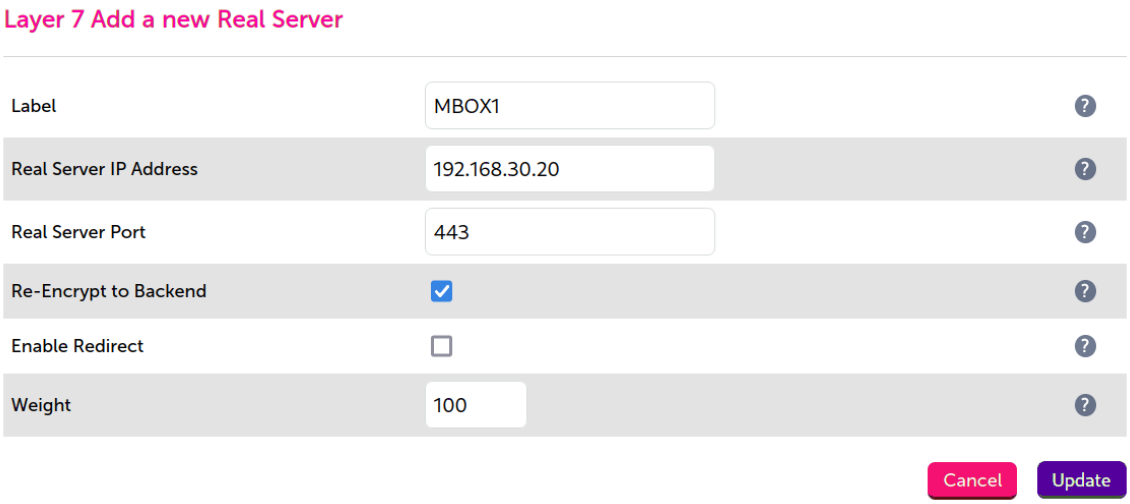

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **443**.
6. Enable (check) the *Re-Encrypt to Backend* checkbox.
7. Click **Update**.
8. Repeat the above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

c) Export Your SSL Certificate

When you export your certificate from Exchange, make sure that your 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*.

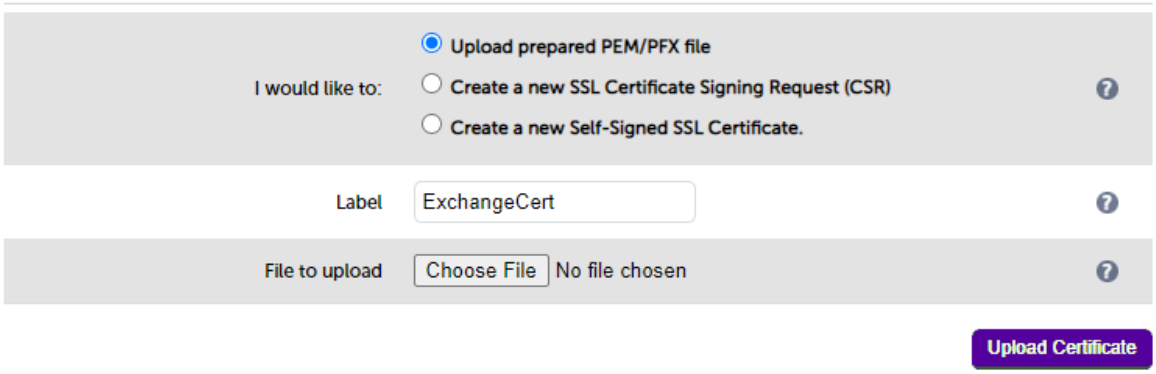

3. Enter a suitable *Label* (name) for the certificate, e.g. **ExchangeCert**.
4. Browse to and select the certificate file to upload (PEM or PFX format).

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**.

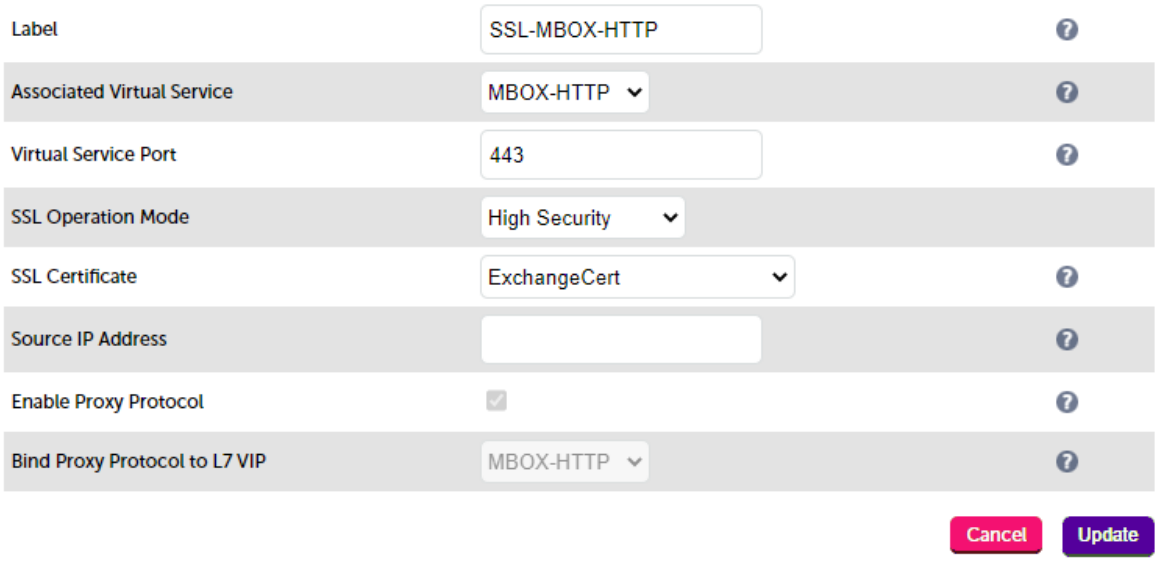

2. Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **MBOX-HTTP**.

> **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.

10.2.2. 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:

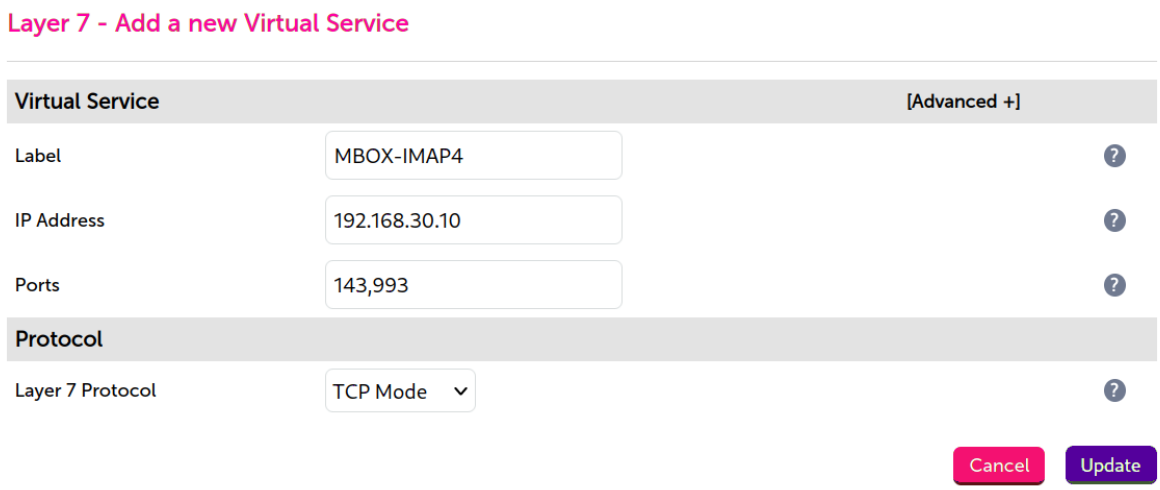

3. Enter an appropriate label for the VIP, e.g. **MBOX-IMAP4**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **143,993**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify**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

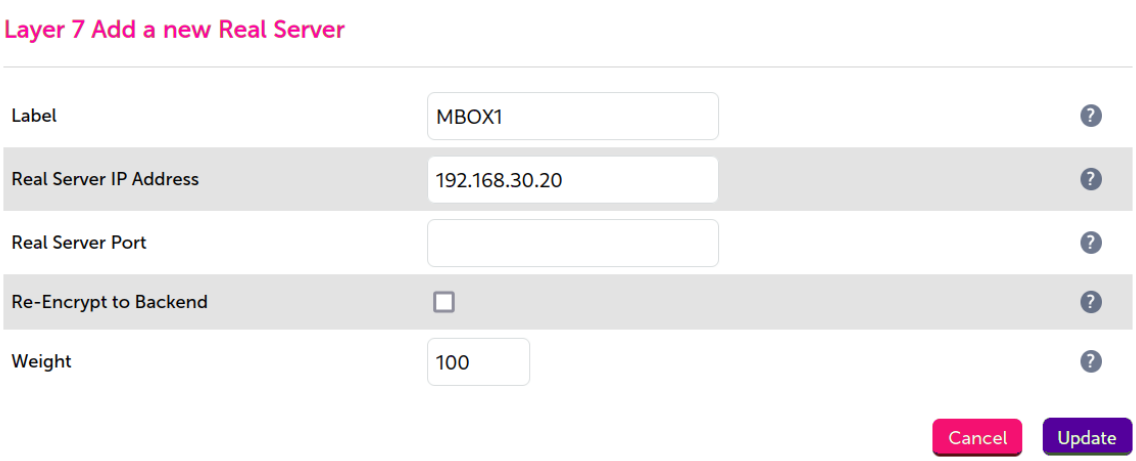

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Leave the *Real Server Port* field blank.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

10.2.3. 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:

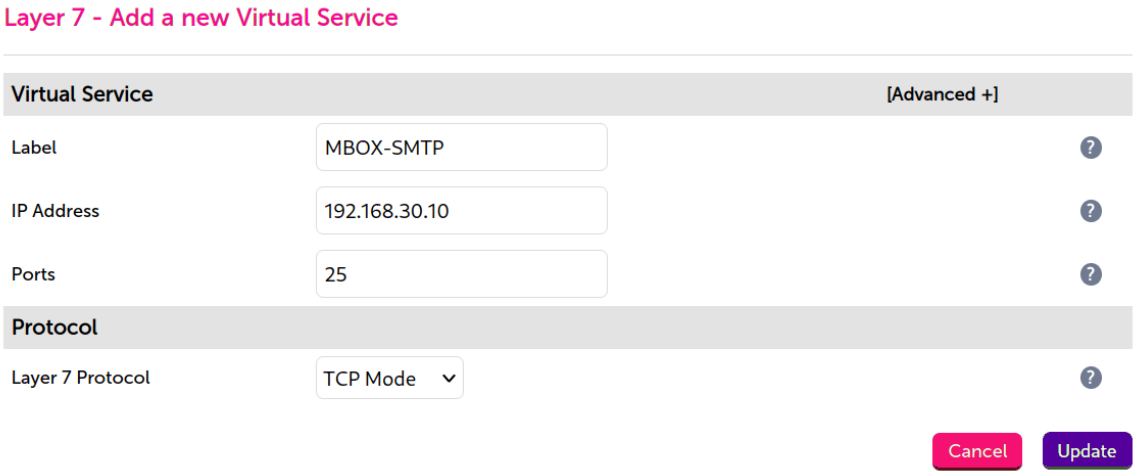

3. Enter an appropriate label for the VIP, e.g. **MBOX-SMTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **25**.
6. Set *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** 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:

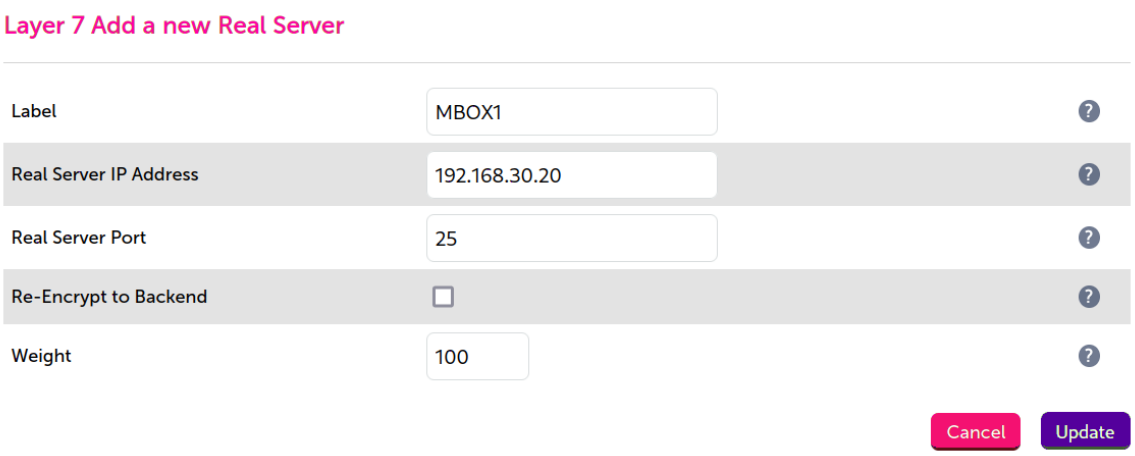

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Change the *Real Server Port* field to **25**.
6. Click **Update**.
7. Repeat the above steps to add your other Mailbox Server(s).

10.2.4. 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.

10.2.5. 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).

10.2.6. Finalizing the Configuration

To apply the new settings, HAProxy and STunnel must both be reloaded. This can be done using the buttons in the "Commit changes" 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**.

10.3. Exchange Server Configuration Steps

10.3.1. Configure IIS logging to Capture XFF Header IP Addresses

Please refer to this Microsoft article for configuration steps.

11. Appliance Configuration – Using Layer 4 DR Mode

11.1. Load Balancer Deployment Overview

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

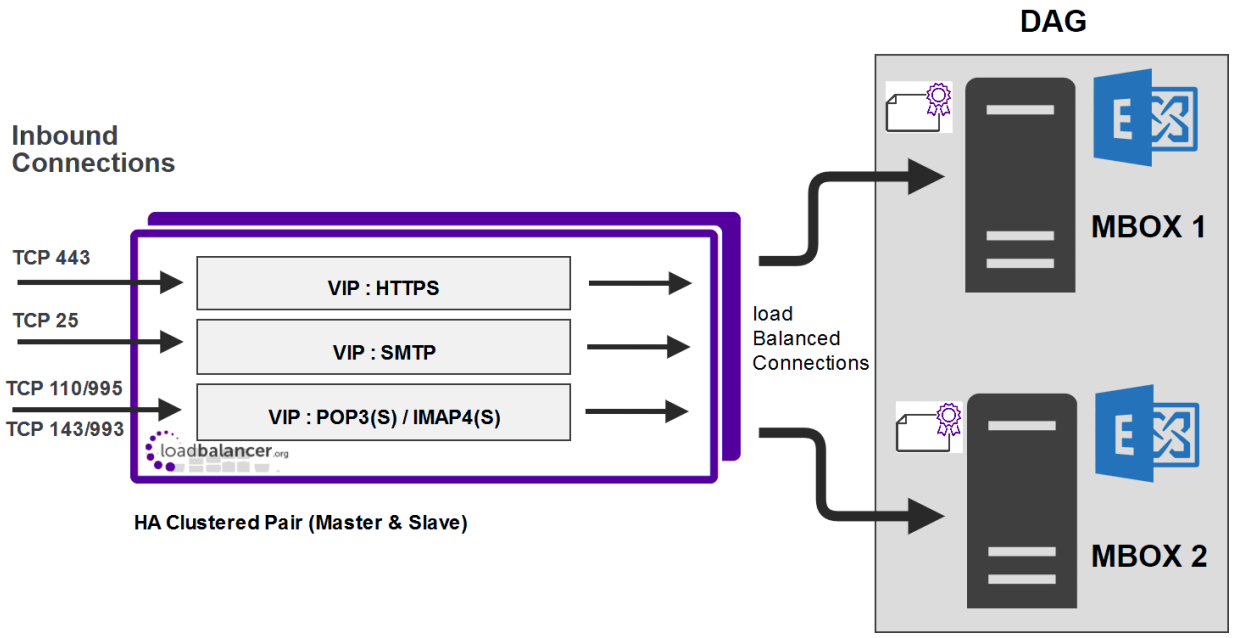

Notes

- 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.

11.2. Load Balancer Configuration

11.2.1. 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:

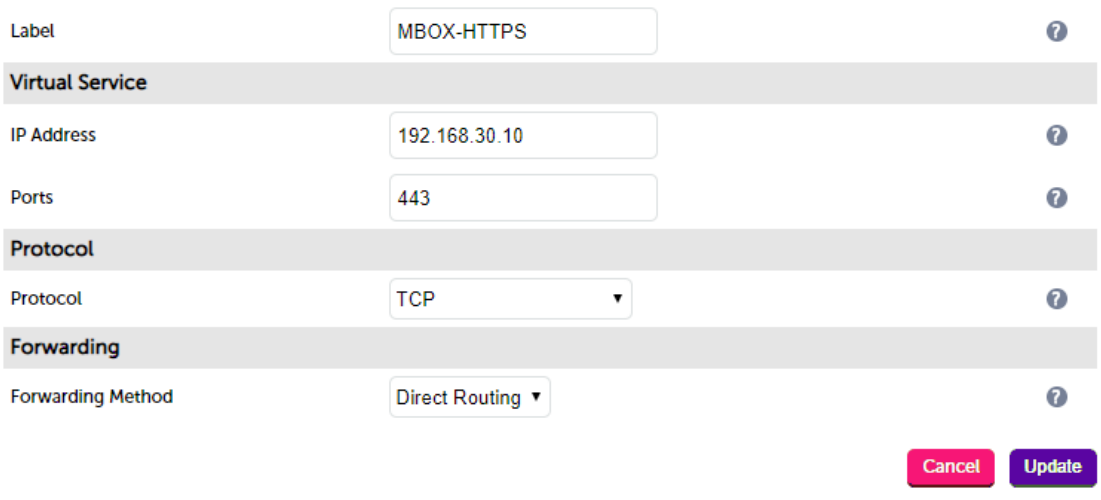

3. Enter an appropriate label for the VIP, e.g. **MBOX-HTTPS**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **443**.
6. Leave *Protocol* set to **TCP**.
7. Leave *Forwarding Method* set to **Direct Routing**.
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 the *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:

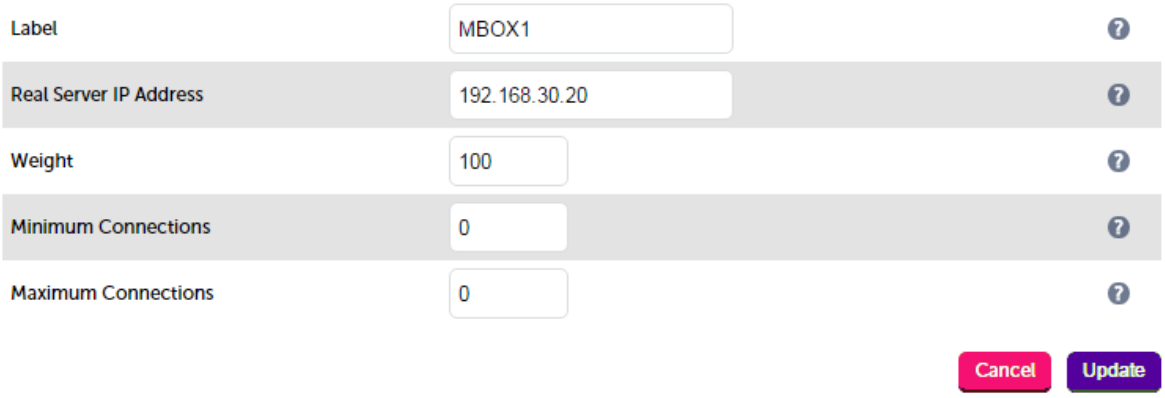

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Click **Update**.
6. Repeat the above steps to add your other Mailbox Server(s).

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.

11.2.2. 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:

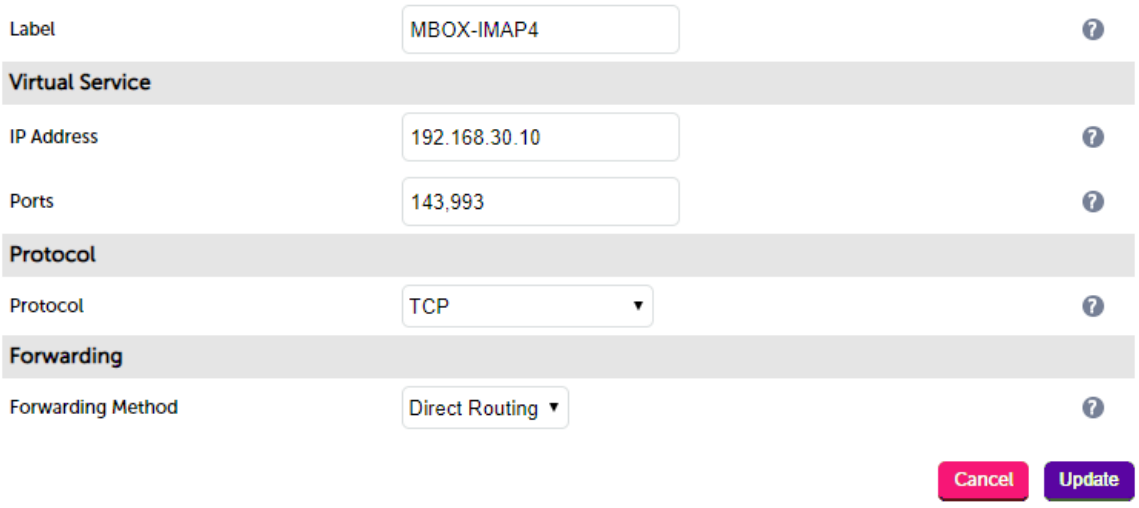

3. Enter an appropriate label for the VIP, e.g. **MBOX-IMAP4**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **143,993**.
6. Leave *Protocol* set to **TCP**.
7. Leave *Forwarding Method* set to **Direct Routing**.
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:

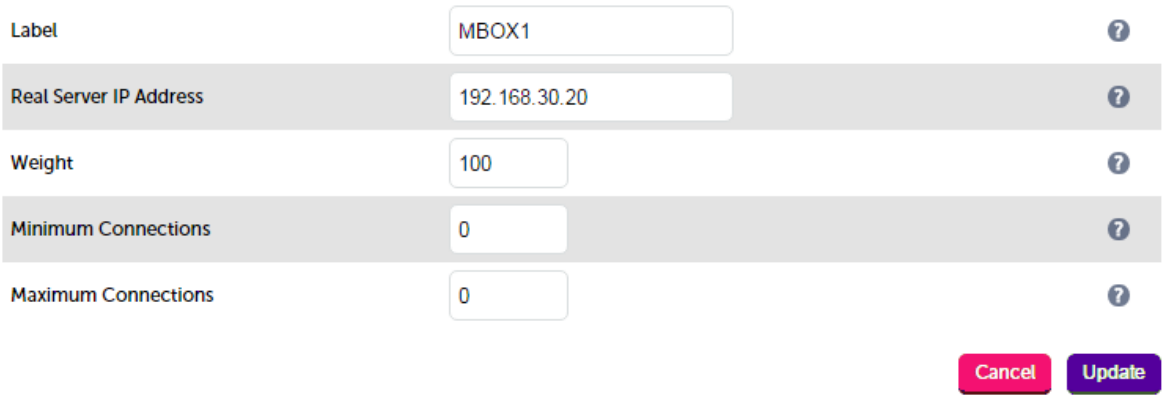

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**.
5. Click **Update**.
6. Repeat the above steps to add your other Mailbox Server(s).

11.2.3. 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:

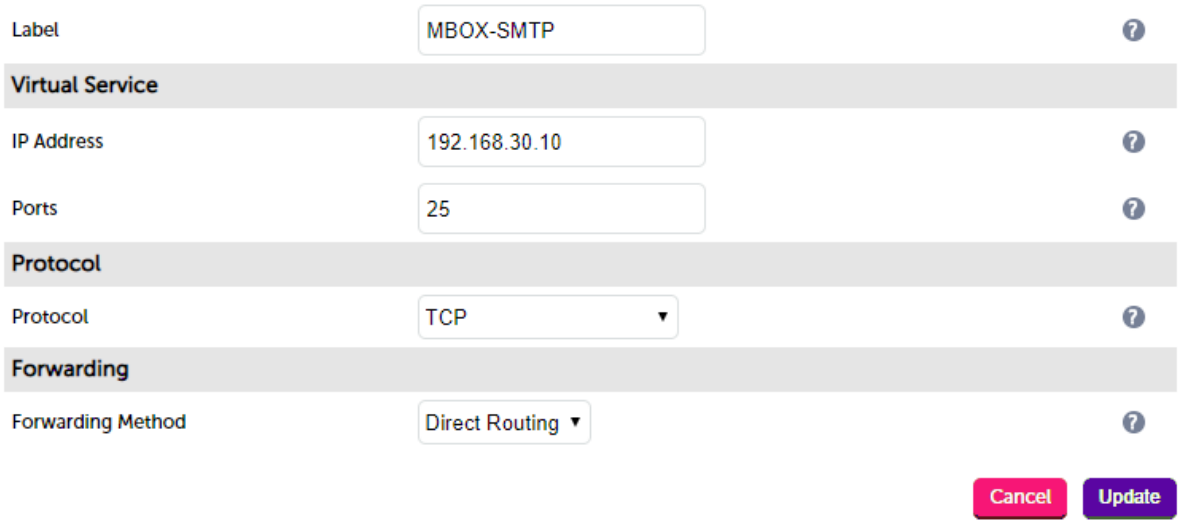

3. Enter an appropriate label for the VIP, e.g. **MBOX-SMTP**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**.
5. Set the *Virtual Service Ports* field to **25**.
6. Leave *Protocol* set to **TCP**.
7. Leave *Forwarding Method* set to **Direct Routing**.
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:

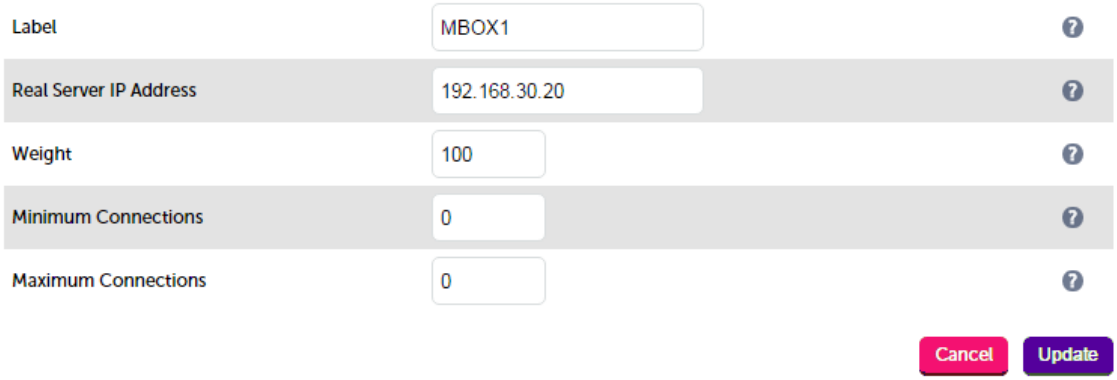

3. Enter an appropriate label for the RIP, e.g. **MBOX1**.
4. Change the ***Real Server IP Address*** field to the required IP address, e.g. **192.168.30.20**.
5. Click **Update**.
6. Repeat the above steps to add your other Mailbox Server(s).

11.3. 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.

12. Testing & Verification

> **Note** For additional guidance on diagnosing and resolving any issues you may have, please also refer to Diagnostics & Troubleshooting.

12.1. Useful Exchange 2016 & Other Microsoft Tools

12.1.1. 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:

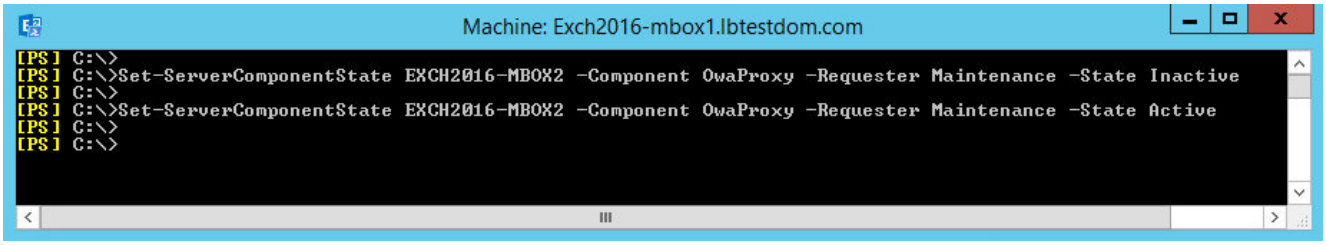

12.1.2. Testing Mailflow

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:

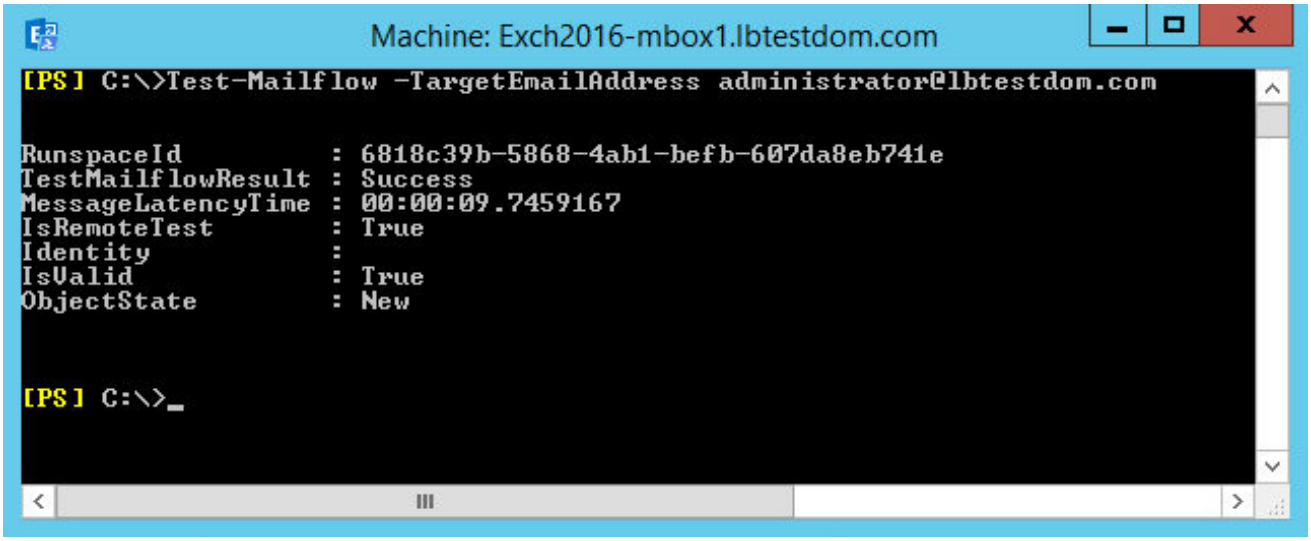

If everything is working correctly, a new message will appear in the test users mailbox:

12.1.3. 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:

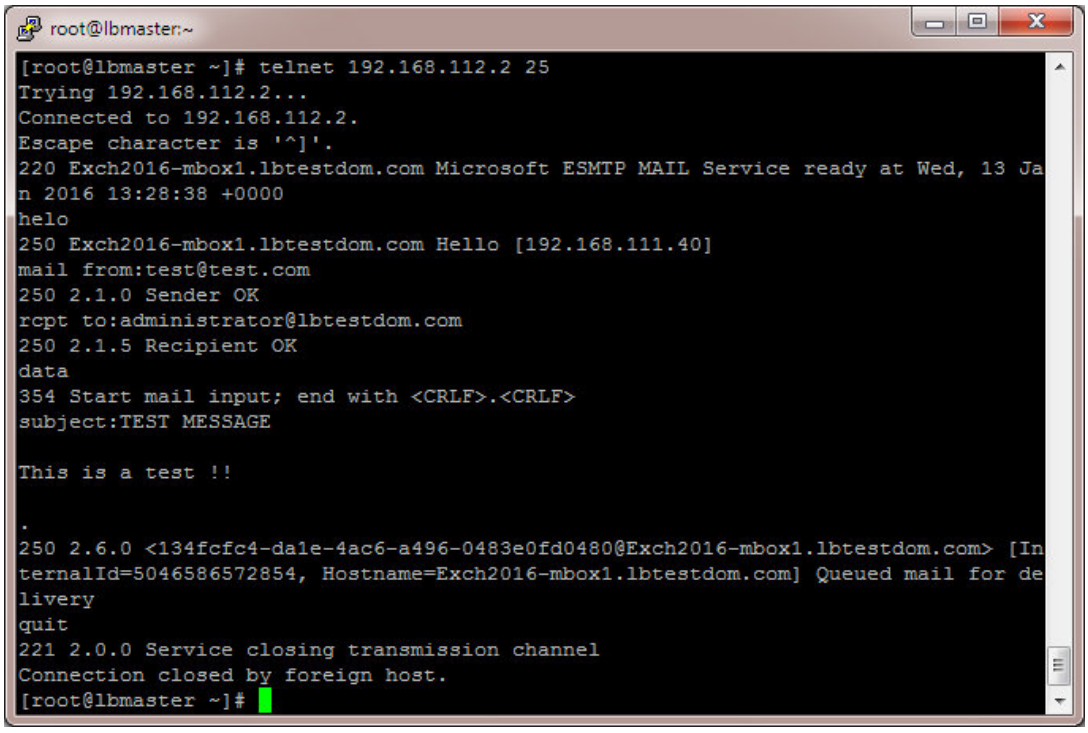

If everything is working correctly, a new message will appear in the test user'ss mailbox:

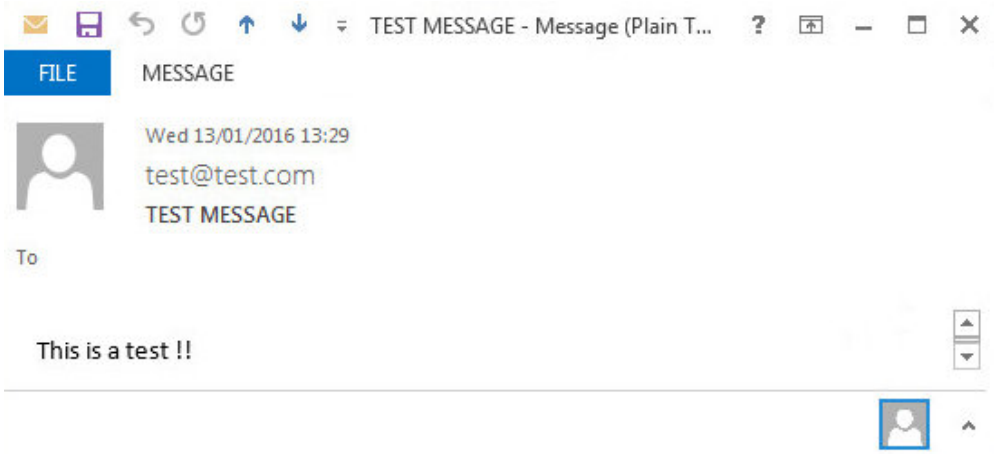

To do the same test via the load balancer, connect to the VIP rather than directly to each server, e.g.:

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

12.1.4. 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.

12.2. Useful Appliance based Tools & Features

12.2.1. 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:

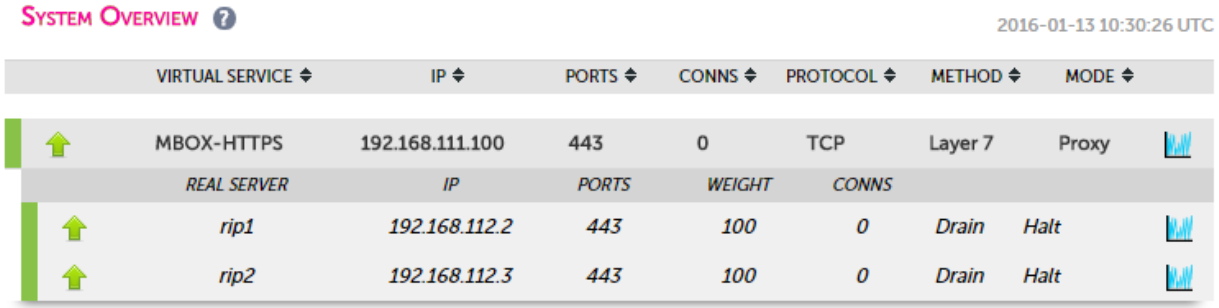

The example below shows that rip2 has been put in halt mode:

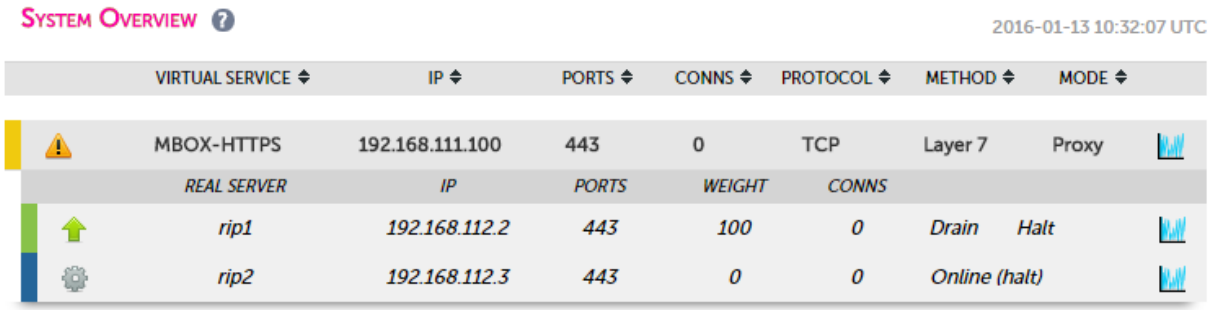

12.2.2. Layer 4 Status Report

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***.

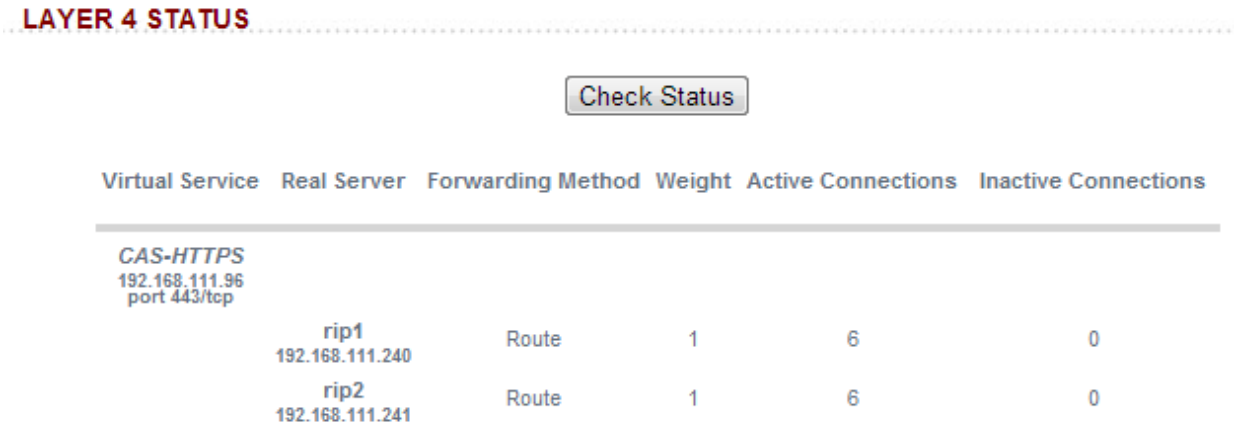

12.2.3. Layer 7 Statistics Report

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***.

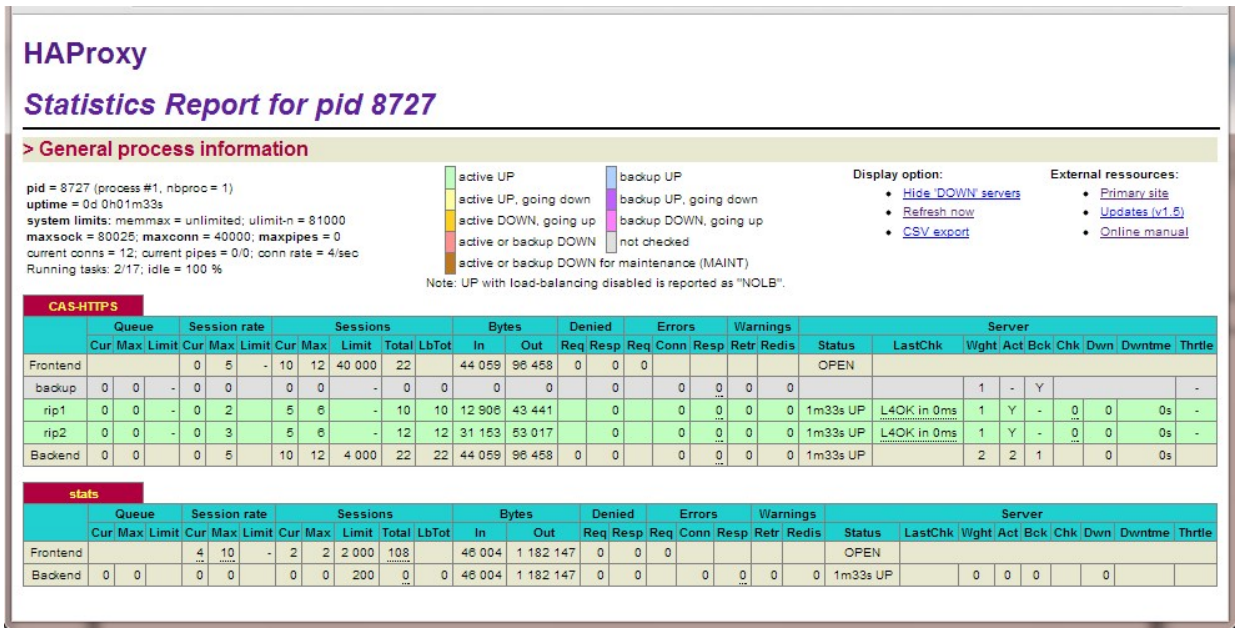

12.2.4. Appliance Logs

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***.

13. Technical Support

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

14. Further Documentation

For additional information, please refer to the Administration Manual.

15. Appendix

15.1. Configuring 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.

15.1.1. 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:

```
# Lockdown SMTP inbound connections
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

echo "Firewall Activated"
exit 0;
```

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**.

15.2. 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:

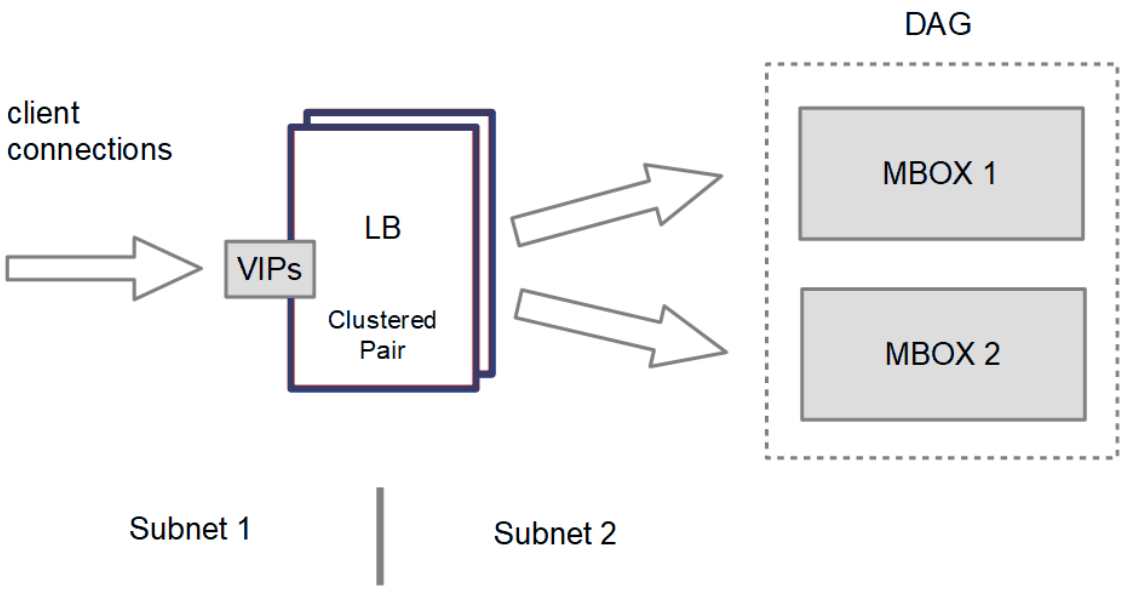

2-arm configuration - key points to note:

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.

15.3. 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.

15.3.1. 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.

15.4. 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**

15.5. Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution. We recommend that the Primary appliance is fully configured first, then the Secondary appliance can be added to create an HA pair. Once the HA pair is configured, load balanced services must be configured and modified on the Primary appliance. The Secondary appliance will be automatically kept in sync.

> **Note** For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the documentation library

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

15.5.1. Non-Replicated Settings

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

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	Interface IP addresses, bonding configuration and VLANs
Local Configuration	Routing	Default gateways and static routes
Local Configuration	System Date & time	Time and date related settings

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Physical – Advanced Configuration	Various appliance settings
Local Configuration	Portal Management	Portal management settings
Local Configuration	Security	Security settings
Local Configuration	SNMP Configuration	SNMP settings
Local Configuration	Graphing	Graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Backup & Restore	Local XML backups
Maintenance	Software Updates	Appliance software updates
Maintenance	Firewall Script	Firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

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

15.5.2. Configuring the HA Clustered Pair

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

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

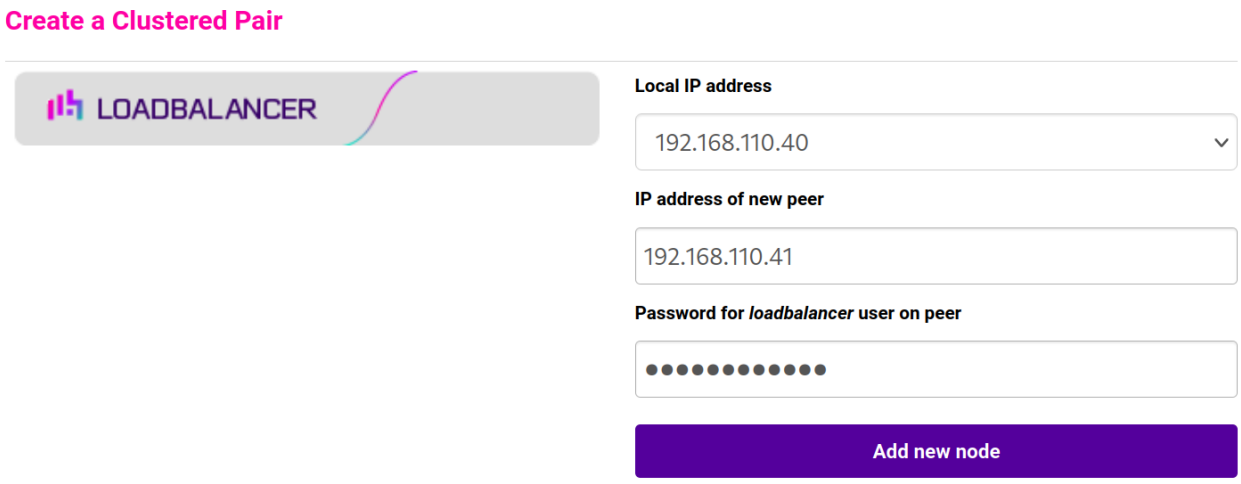

3. Specify the IP address and the ***loadbalancer*** user's password for the Secondary (peer) appliance as shown in the example above.

4. Click **Add new node**.
5. The pairing process now commences as shown below:

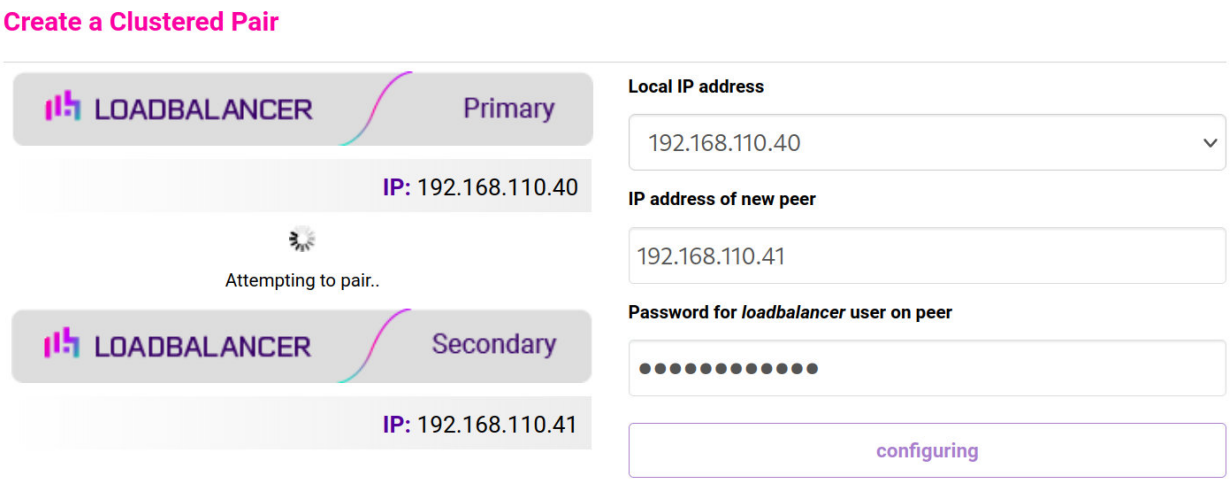

6. Once complete, the following will be displayed on the Primary appliance:

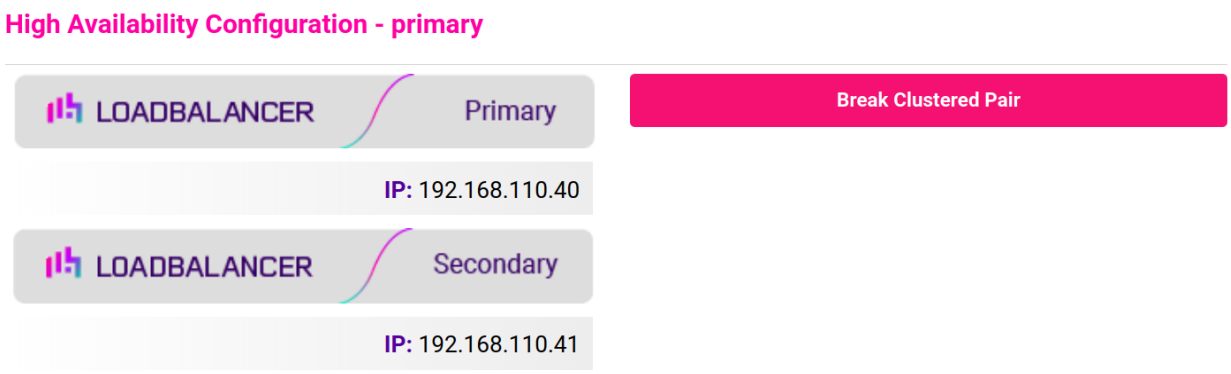

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

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

> **Note** For more details on configuring HA with 2 appliances, please refer to Appliance Clustering for HA.

> **Note** For details on testing and verifying HA, please refer to Clustered Pair Diagnostics.

15.6. Solving the ARP Problem

15.6.1. Windows Server 2012 & Later

Windows Server 2012 and later support Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter that must be installed and configured on each load balanced (Real) Server. The IP address configured on

the Loopback Adapter must be the same as the Virtual Service (VIP) address. This enables the server to receive packets that have their destination set as the VIP address. If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

In addition, the strong/weak host behavior must be configured on each Real Server. The weak host model allows packets with any IP to be sent or received via an interface. The strong host model only allows packets with an IP belonging to the interface to be sent or received.

> **Important** The following 3 steps must be completed on ***all*** Real Servers associated with the VIP.

Step 1 of 3: Install the Microsoft Loopback Adapter

1. Click **Start**, then run **hdwwiz** to start the Hardware Installation Wizard.
2. Once 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**.

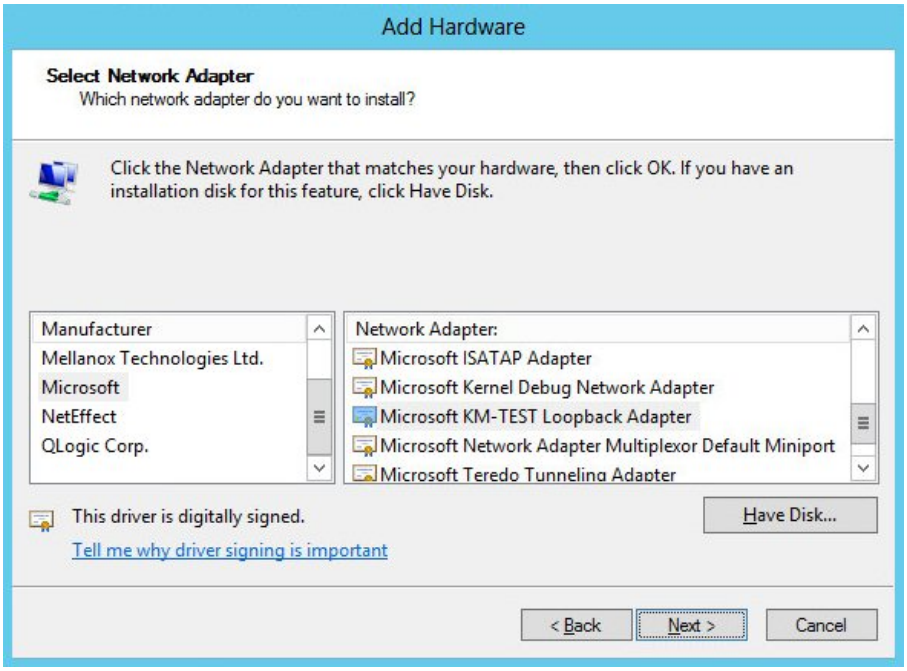

5. Select **Microsoft** & **Microsoft KM-Test Loopback Adapter**, click **Next**.
6. Click **Next** to start the installation, when complete click **Finish**.

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**.

> **Note** You can configure IPv4 or IPv6 addresses or both depending on your requirements.

IPv4 Addresses

1. Uncheck all items except **Internet Protocol Version 4 (TCP/IPv4)** as shown below:

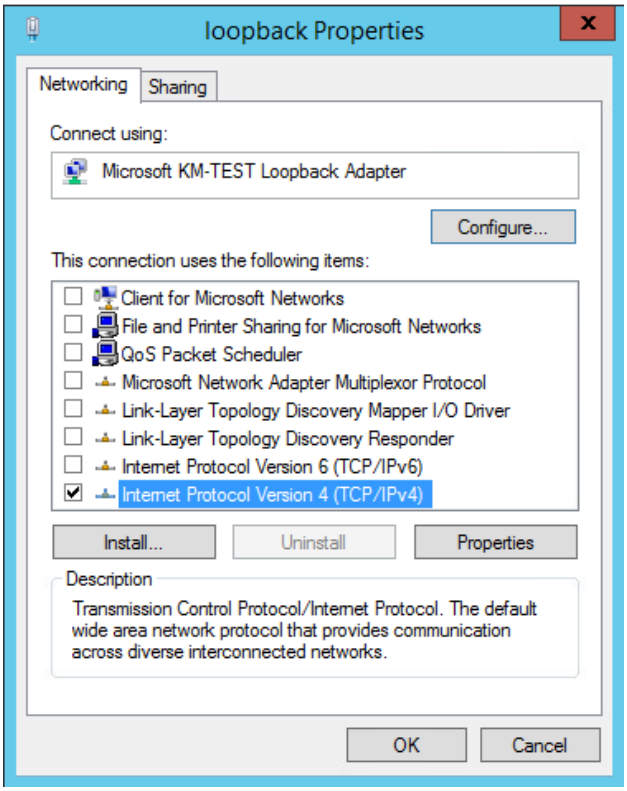

2. Ensure that **Internet Protocol Version (TCP/IPv4)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service address (VIP) with a subnet mask of **255.255.255.255**, e.g. **192.168.2.20/255.255.255.255** as shown below:

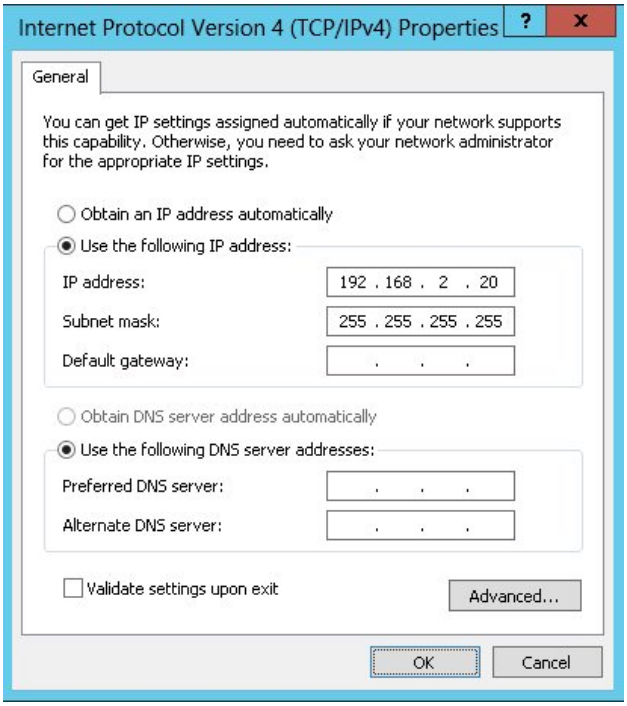

> **Note** **192.168.2.20** is an example, make sure you specify the correct VIP address.

> **Note** If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

3. Click **OK** then click **Close** to save and apply the new settings.

IPv6 Addresses

1. Uncheck all items except **Internet Protocol Version 6 (TCP/IPv6)** as shown below:

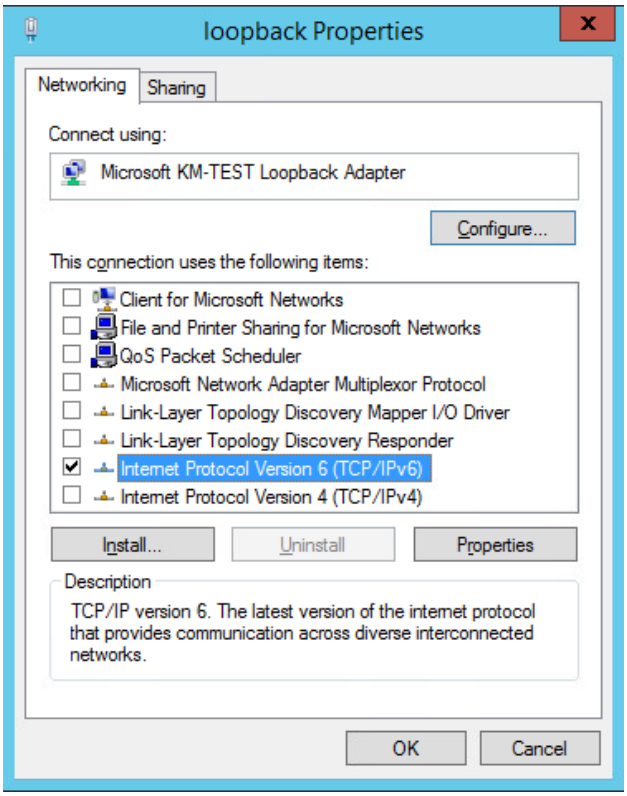

2. Ensure that **Internet Protocol Version (TCP/IPv6)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) and set the ***Subnet Prefix Length*** to be the same as your network setting, e.g. **2001:470:1f09:e72::15/64** as shown below:

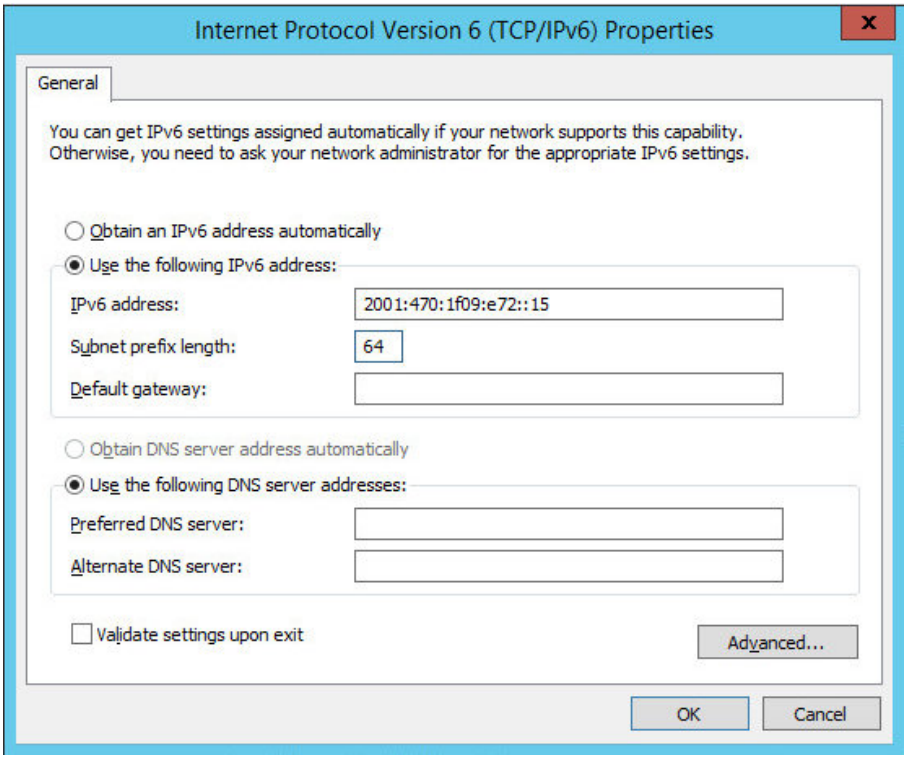

> **Note** **2001:470:1f09:e72::15/64** is an example, make sure you specify the correct VIP address.

> **Note** If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

3. Click **OK** then click **Close** to save and apply the new settings.

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

The strong/weak host behavior can be configured using either of the following 2 methods:

- Option 1 - Using network shell (netsh) commands
- Option 2 - Using PowerShell cmdlets

The commands in this section assume that the LAN Adapter is named "**net**" and the Loopback Adapter is named "**loopback**" as shown in the example below:

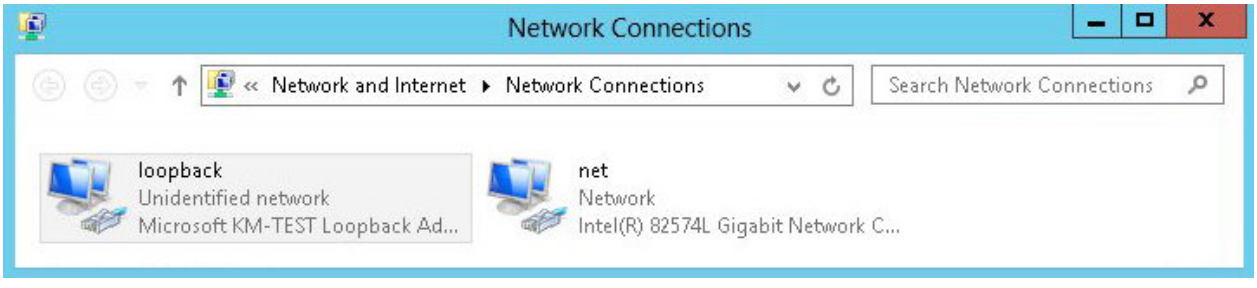

> **Important** Either adjust the commands to use the names allocated to your LAN and loopback adapters, or rename the adapters before running the commands. Names are case sensitive so make sure

that the interface names used in the commands match the adapter names exactly.

Option 1 - Using Network Shell (netsh) Commands

To configure the correct strong/weak host behavior run the following commands:

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 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
```

Option 2 - Using PowerShell Cmdlets

For IPv4 addresses:

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

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

For IPv6 Addresses:

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

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

15.6.2. 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:

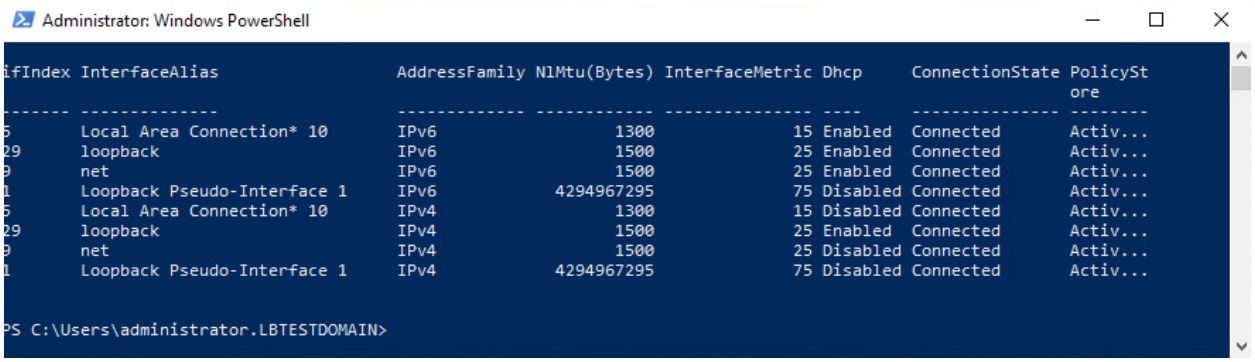

> **Note** The interface metric is displayed in the 5th column

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.

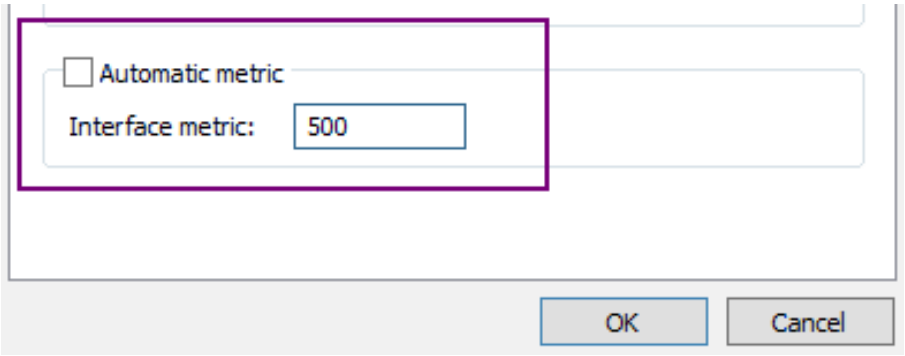

3. Click **OK**, **OK** and **Close** to apply the new settings.

16. 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
1.4.3	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.4.4	5 January 2023	Combined software version information into one section Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section Removed references to the colour of certain UI elements	Housekeeping across all documentation	AH
1.4.5	2 February 2023	Updated screenshots	Branding update	AH

Version	Date	Change	Reason for Change	Changed By
1.4.6	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.5.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH
1.5.1	19 June 2024	Changed the health check for the HTTPS mailbox service so the load balancer explicitly looks for "200 OK" in the response.	Service pack update compatibility	RJC

LOADBALANCER

Visit us: www.loadbalancer.org

Phone us: +44 (0)330 380 1064

Phone us: +1 833 274 2566

Email us: info@loadbalancer.org

Follow us: @loadbalancer.org

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.