Load Balancing Microsoft Exchange 2013

v1.5.3

Deployment Guide

NOTE: This guide has been archived and is no longer being maintained. While the content is still valid for the particular software versions mentioned, it may refer to outdated software that has now reached end-of-life. For more information please contact support@loadbalancer.org.

loadbalancer.org

Contents

1. About this Guide

This guide details the steps required to configure a load balanced Microsoft Exchange 2013 environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Microsoft Exchange 2013 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 2013. For full specifications of available models please refer to: https://www.loadbalancer.org/products.

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

3. Loadbalancer.org Software Versions Supported

- V8.3.7 and later

4. Microsoft Exchange Software Versions Supported

- Microsoft Exchange 2013 CU2 (15.0.712.24) and later

5. Exchange Server 2013

Exchange 2013 is Microsoft's latest enterprise level messaging and collaboration server. Exchange 2013 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 Note: Exchange 2013 has since been superseded by Exchange 2016 and 2019. Deployment guides for both are available here.

6. Exchange 2013 Server Roles

Exchange 2013 has been consolidated into two roles, these are: the **Client Access Server** role and the **Mailbox Server** role. The functionality of the Hub Transport server role has split between the CAS role (the *Front End Transport Service*) and the Mailbox Server role (the *Transport Service* and the *Mailbox Transport Service*). The Edge Transport server role has been removed.

Role	Purpose
Client Access Server	The Client Access Server role is comprised of three components, client protocols, SMTP, and a UM Call Router. The CAS role is a thin, protocol session stateless server that is organized into a load balanced configuration. Unlike previous

	versions, session affinity is not required at the load balancer. This is because logic now exists in CAS to authenticate the request, and then route the request to the Mailbox server that hosts the active copy of the mailbox database. *Note: A number of issues have been seen with IOS-7 on the iPhone when used with ActiveSync. Upgrading to IOS-8 resolved these issues.*
Mailbox Server	The Mailbox Server role now hosts all the components and/or protocols that process, render and store the data. No clients will ever connect directly to the Mailbox server role; all client connections are handled by the Client Access Server role. Mailbox Servers can be added to a Database Availability Group, thereby forming a high available unit that can be deployed in one or more datacenters.

CAS Array Object

This concept has been removed and there is no longer any need to define a CAS array object.

Client Access Protocols

Outlook clients no longer use RPC to access their mailbox. This is now handled only by RPC over HTTPS (aka Outlook Anywhere). Native RPC is only used for server to sever communication. POP3 and IMAP4 continue to be supported as with previous versions.

External SMTP Mail flow

External SMTP communication is now handled by the *Front End Transport Service* on the CAS role.

Exchange Administration

The Exchange Admin Center (EAC) is the new web-based management console in Microsoft Exchange Server 2013. The EAC replaces the Exchange Management Console (EMC) and the Exchange Control Panel (ECP), which were the two interfaces used to manage Exchange Server 2010. Note that "ECP" is still the name of the virtual directory used by the EAC.

7. Load Balancing Exchange 2013

> Note: It's highly recommended that you have a working Exchange 2013 environment first before implementing the load balancer.

Load Balancing & HA Requirements

In Exchange Server 2013, there are two basic building blocks – the Client Access Array and the Database Availability Group (DAG). Each provides a unit of high availability and fault tolerance that are decoupled from one another. Multiple Client Access Servers make up the Client Access Array, while multiple Mailbox Servers form the DAG.

Client Access Array

As mentioned earlier, the 2010 concept of a CAS Array no longer exists. In 2013, a Client Access Array is simply a group of two or more Client Access Servers. The load balancer then enables resilience and HA.

Database Availability Group (DAG)

A DAG is a group of up to 16 Mailbox Servers that hosts a set of databases and provides automatic database-level recovery from failures that affect individual servers or databases.

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

Persistence (aka Server Affinity)

Due to Exchange 2013's new architecture, all sessions to the CAS servers are stateless and therefore persistence/affinity is no longer required on the load balancer.

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	CAS	Inbound SMTP
110	CAS	POP3 clients
143	CAS	IMAP4 clients
443	CAS	HTTPS (Outlook Web App, AutoDiscovery, Web Services, ActiveSync, Outlook Anywhere, Offline Address Book, Exchange Administration Center)
993	CAS	Secure IMAP4 clients
995	CAS	Secure POP3 clients

SSL Termination

SSL offloading for Exchange 2013 is supported from SP1 as detailed in this Microsoft article. However, for scalability and effective load sharing we recommend terminating SSL on the Exchange Servers rather than on the load balancer.

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.robstest.com	Outlook Web App, ActiveSync, Outlook Anywhere, Offline Address Book, Exchange Web Services
autodiscover.robstest.com	Auto Discover

Notes:

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

Example 2 – expanded namespace configuration

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

Notes:

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

Health-Checks

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

Load Balancer Deployment

There are multiple ways to deploy Exchange, but in this example two servers are used. Each server hosts the CAS & Mailbox roles in a DAG configuration. This provides high availability and uses a minimum number of Exchange Servers.

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

according to the load balancing algorithm selected.

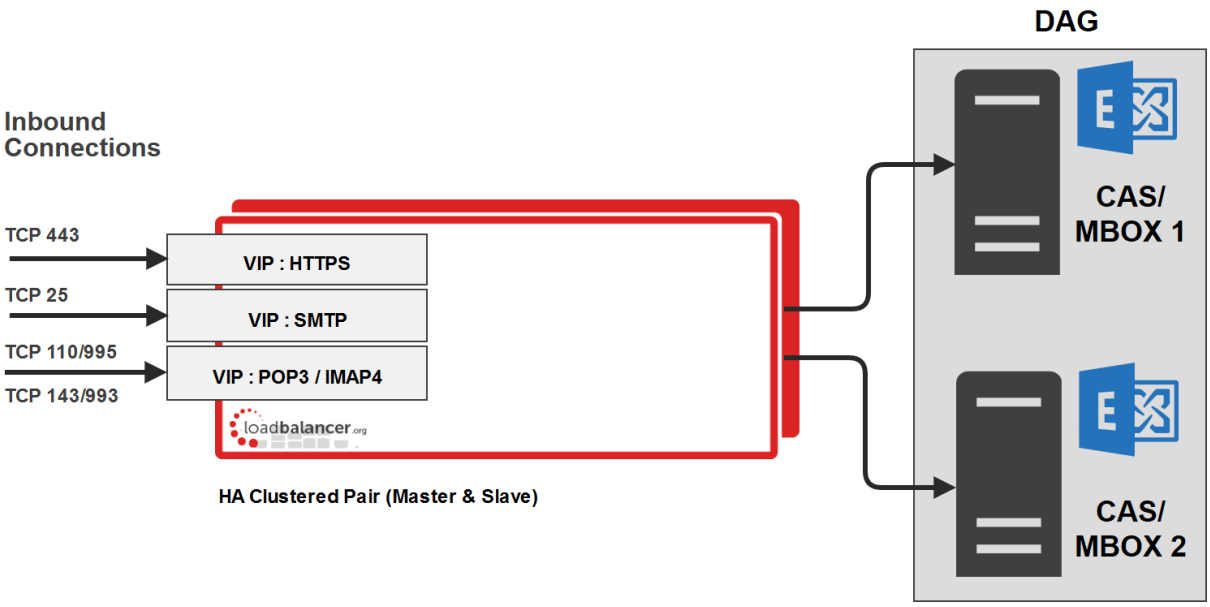

VIP = Virtual IP Addresses

> Note: The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to section 5 in the appendix on page 39 for more details on configuring a clustered pair.

Virtual Service (VIP) Requirements

To provide load balancing and HA for Exchange 2013, 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, see page 38 in the appendix for more details)
- IMAP4
- POP3

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

Load Balancer Deployment Modes

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

For Exchange 2013, layer 4 DR mode or layer 7 SNAT mode is recommended. These modes are described below and are used for the configurations presented in this guide. For configuring using DR mode please refer to the section starting on page 20, for configuring using layer 7 SNAT mode, refer to the section starting on page 25.

Layer 4 DR Mode

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

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

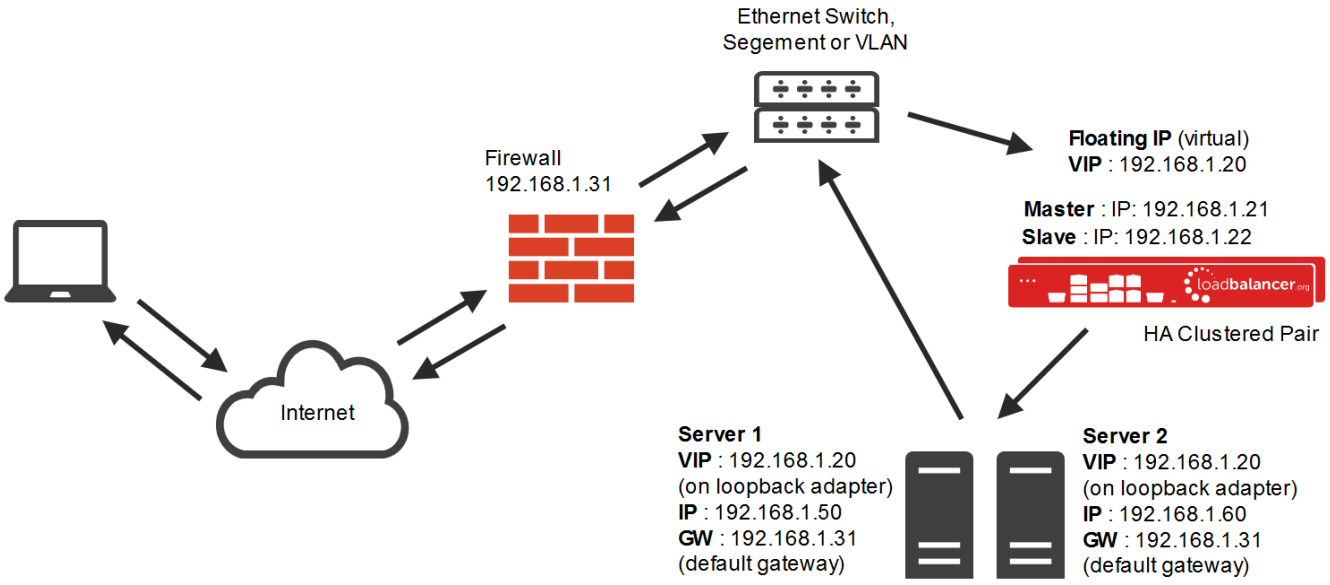

- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that you need to ensure that the Real Server (and the load balanced application) respond to both the Real Servers own IP address and the VIP
- The Real Server should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Servers in this way is referred to as ***Solving the ARP Problem***. Please refer to page 41 for more information

- On average, DR mode is 8 times quicker than NAT for HTTP, 50 times quicker for Terminal Services and much, much faster for streaming media or FTP
- The load balancer must have an Interface in the same subnet as the Real Servers to ensure layer 2 connectivity required for DR mode to work
- The VIP can be brought up on the same subnet as the Real Servers, or on a different subnet provided that the load balancer has an interface in that subnet
- Port translation is not possible in DR mode i.e. having a different RIP port than the VIP port
- DR mode is transparent, i.e. the Real Server will see the source IP address of the client

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 request to the chosen Real Server. As a result, Layer 7 is a slower technique than DR or NAT mode at Layer 4. 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.

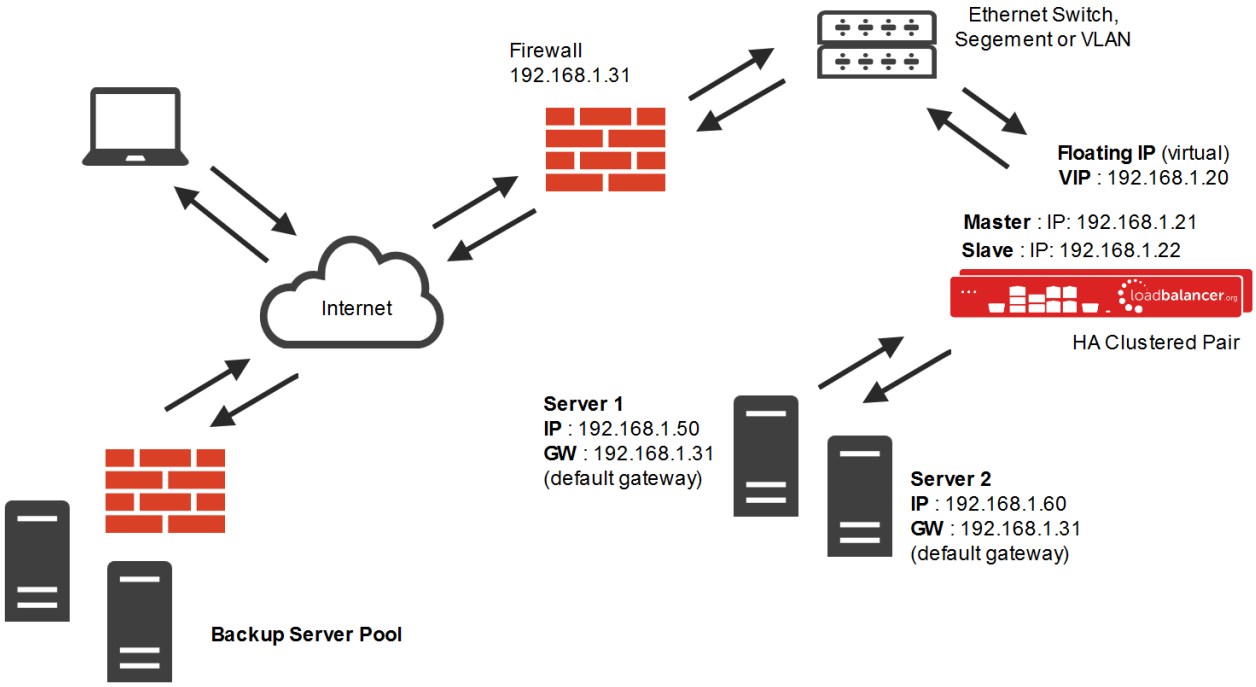

This mode can be deployed in a one-arm or two-arm configuration and does not require any changes to the Real Servers. However, since the load balancer is acting as a full proxy it doesn't have the same raw throughput as the layer 4 methods.

The load balancer proxies the application traffic to the servers so that the source of all traffic becomes the load balancer.

- SNAT mode is a full proxy and therefore load balanced Real Servers do not need to be changed in any way

- Because SNAT mode is a full proxy any server in the cluster can be on any accessible subnet including across the Internet or WAN
- 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 balancers 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 clients IP address can be passed through either by enabling TProxy on the load balancer, or for HTTP, using X-forwarded-For headers. Please refer to chapter 6 in the Administration Manual for more details
- SNAT mode can be deployed using either a 1-arm or 2-arm configuration

Our Recommendation

Where possible we recommend that Layer 4 Direct Routing (DR) mode is used. This mode offers the best possible performance since replies go directly from the Real Servers to the client, not via the load balancer. It's also relatively simple to implement. Ultimately, the final choice does depend on your specific requirements and infrastructure.

If DR mode cannot be used, for example if the Exchange servers are located in remote routed networks then SNAT mode is recommended.

Note: It's important to remember that when using SNAT mode (HAProxy), the source IP address of packets reaching the Exchange Servers will be the IP address of the load balancer and ***not*** the source IP address of the client.

Transparency is normally only an issue for SMTP traffic at the receive connector. System Administrators typically want to lock down receive connectors to accept SMTP connections only from a controlled set of devices such as external smart mail hosts, printers, networked photocopiers etc.

If transparency for SMTP is the only issue, there are a number of options available to address this:

Option 1 – Enable full layer 7 transparency using TProxy. This is covered in section 1 of the Appendix on page 36.

Option 2 – Use the load balancers on-board firewall to lock down inbound SMTP connections rather than doing this at the receive connector. This is covered in section 2 of the Appendix on page 36.

Option 3 – Configure a layer 4 Virtual Service for SMTP rather than a layer 7 (HAProxy) based Virtual Service. Layer 4 is transparent by default so the source IP address is maintained. This is covered in section 3 of the Appendix on page 37.

8. Configuring Exchange 2013 for Load Balancing

1) External Access Domain

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

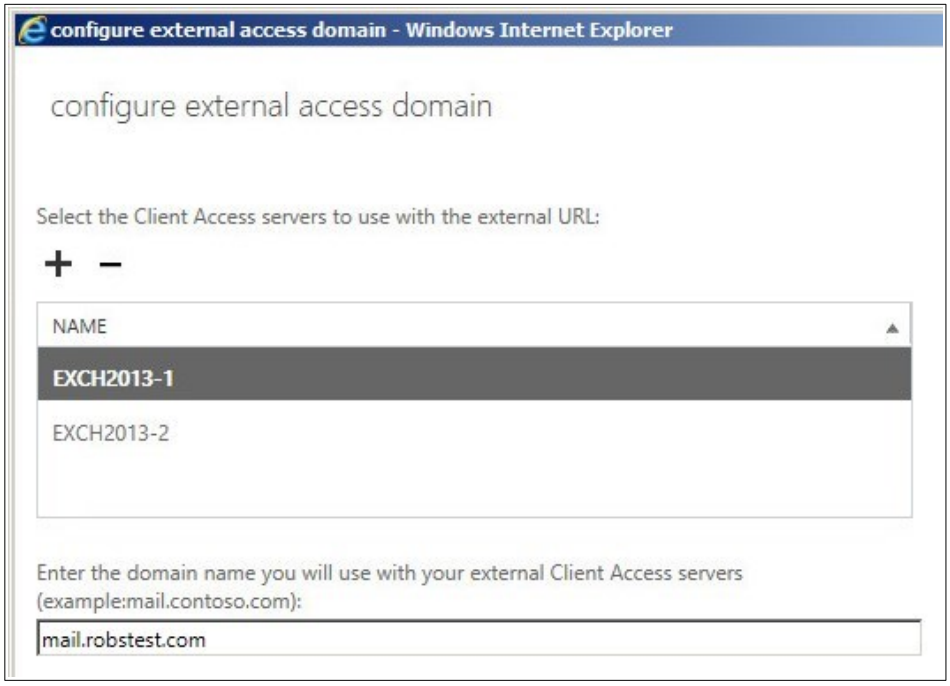

2) Virtual Directories

The Internal and External URL's for the various virtual directories need to be configured to suit your environment. The External URL's are automatically set to be the same as the external access domain when this is configured, but can be changed if needed. The Internal URL's 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:

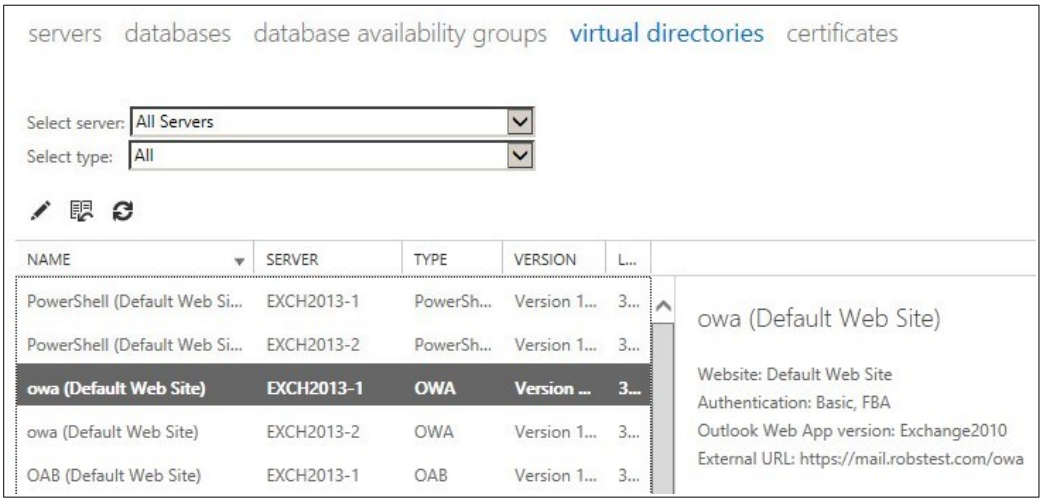

3) Outlook Anywhere

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

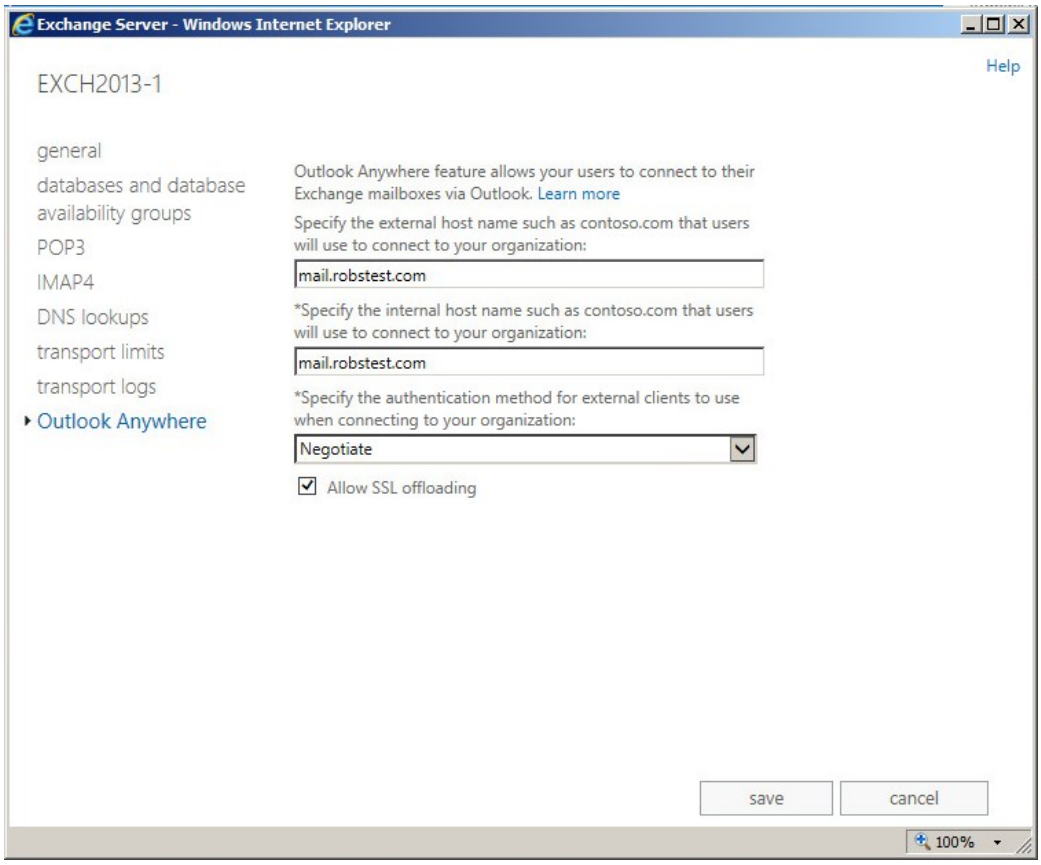

4) Autodiscover

Internal

A new Active Directory object named the service connection point (SCP) is created on the server where you install the Client Access Server role. The SCP object contains the authoritative list of Autodiscover service URLs for the forest. The Set-ClientAccessServer cmdlet is used to update the SCP object as shown in the following example:

```
Set-ClientAccessServer -Identity "EXCH01" -
AutoDiscoverServiceInternalUri
"https://autodiscover.robstest.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 2013 is Outlook 2007, 2003 is NOT supported.

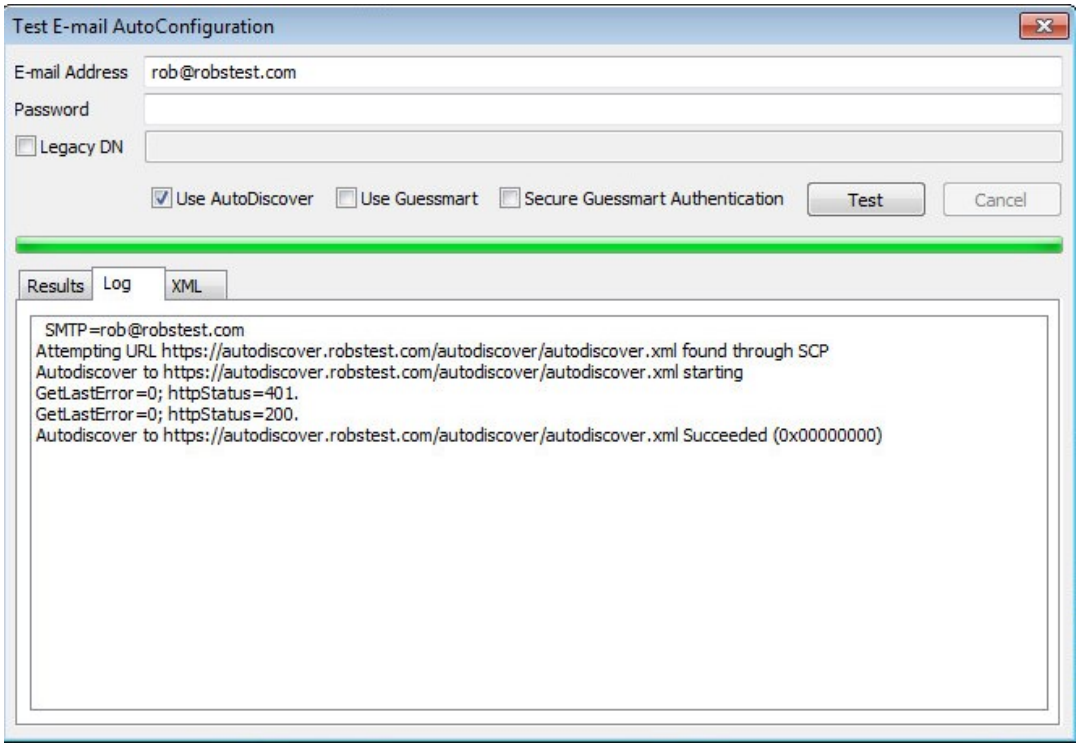

External

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 robstest.com, Outlook will try the following two URLs to try to connect to the Autodiscover service:

https://robstest.com/autodiscover/autodiscover.xml

https://autodiscover.robstest.com/autodiscover/autodiscover.xml

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

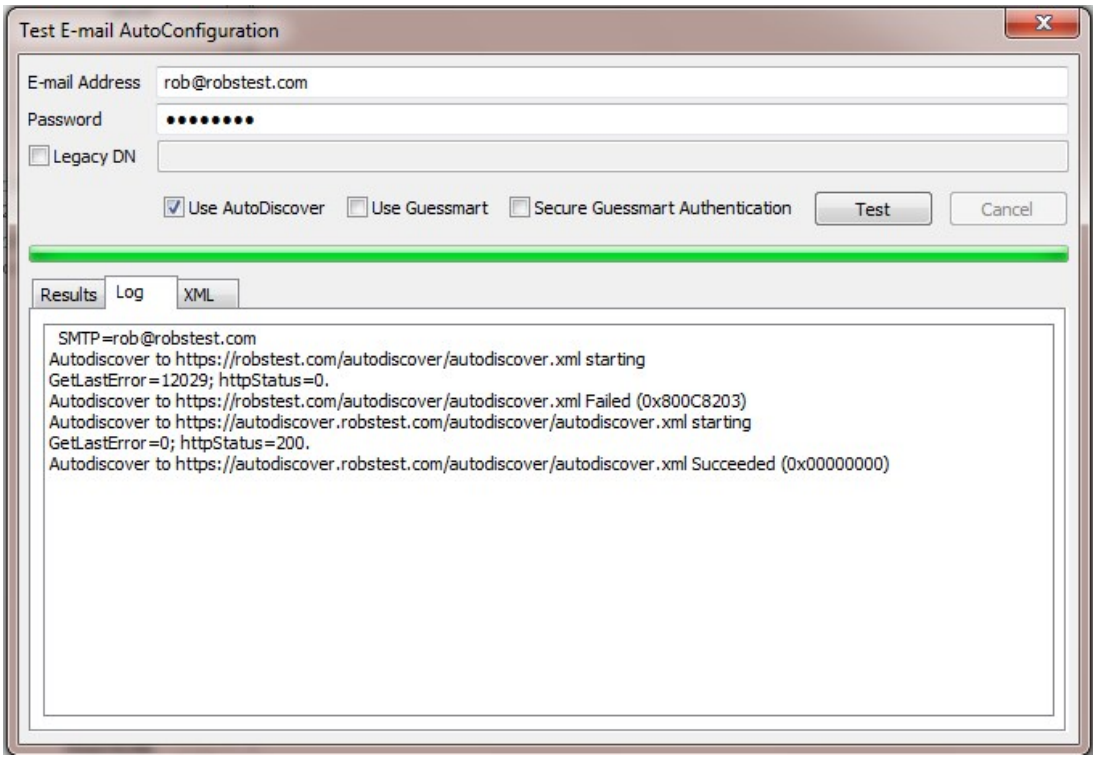

5) Certificates

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

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

*** IMPORTANT!! - the same certificate and private key must be deployed on all Exchange Servers ***

> Note: SSL offloading for Exchange 2013 is supported from SP1 as detailed in this Microsoft article. However, for scalability and effective load sharing we recommend terminating SSL on the Exchange Servers rather than on the load balancer.

6) Send & Receive Connectors

By default no send connectors are created when Exchange 2013 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.

For a dual role server that has both the CAS and Mailbox roles, five receive connectors are automatically created by default. The table below lists these connectors:

Receive Connector	Role	Purpose
Default <server name>	Mailbox	Accepts connections from Mailbox servers running

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

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

http://technet.microsoft.com/en-us/library/jj657461(v=exchg.150).aspx

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 specified that have different costs.

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

7) DNS Configuration

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

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

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

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

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

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 2012/2016, please refer to section 6 of this appendix on page 41.

NAT Mode

When using Layer 4 NAT mode, the default gateway on each Exchange Server MUST be set to be the loadbalancer. It's recommended that a floating IP address is used rather than the interface IP address. This makes it possible for the load balancer to failover to a slave unit and successfully bring up the gateway address.

SNAT Mode

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

9) IIS Restart (** Important **)

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

```
iisreset /restart
```

9. Loadbalancer.org Appliance – the Basics

Virtual Appliance Download & Deployment

A fully featured, fully supported 30 day trial is available if you are conducting a PoC (Proof of Concept) deployment. The VA is currently available for VMware, Virtual Box, Hyper-V, KVM and XEN and has been optimized for each Hypervisor. By default, the VA is allocated 1 CPU, 2GB of RAM and has an 8GB virtual disk. The Virtual Appliance can be downloaded here.

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

Note: Please refer to the Administration Manual and the ReadMe.txt text file included in the VA download for more detailed information on deploying the VA using various Hypervisors.

Initial Network Configuration

The IP address, subnet mask, default gateway and DNS settings can be configured in several ways as detailed below:

Method 1 - Using the Network Setup Wizard at the console

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

Method 2 - Using the WebUI

Using a browser, connect to the WebUI on the default IP address/port: **https://192.168.2.21:9443**

To set the IP address & subnet mask, use: *Local Configuration > Network Interface Configuration*

To set the default gateway, use: *Local Configuration > Routing*

To configure DNS settings, use: *Local Configuration > Hostname & DNS*

Accessing the Web User Interface (WebUI)

1. Browse to the following URL: **https://192.168.2.21:9443/lbadmin/**
   (replace with your IP address if it's been changed)
   ** Note the port number → **9443***

2. Login to the WebUI:

   Username: loadbalancer
   Password: loadbalancer

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

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

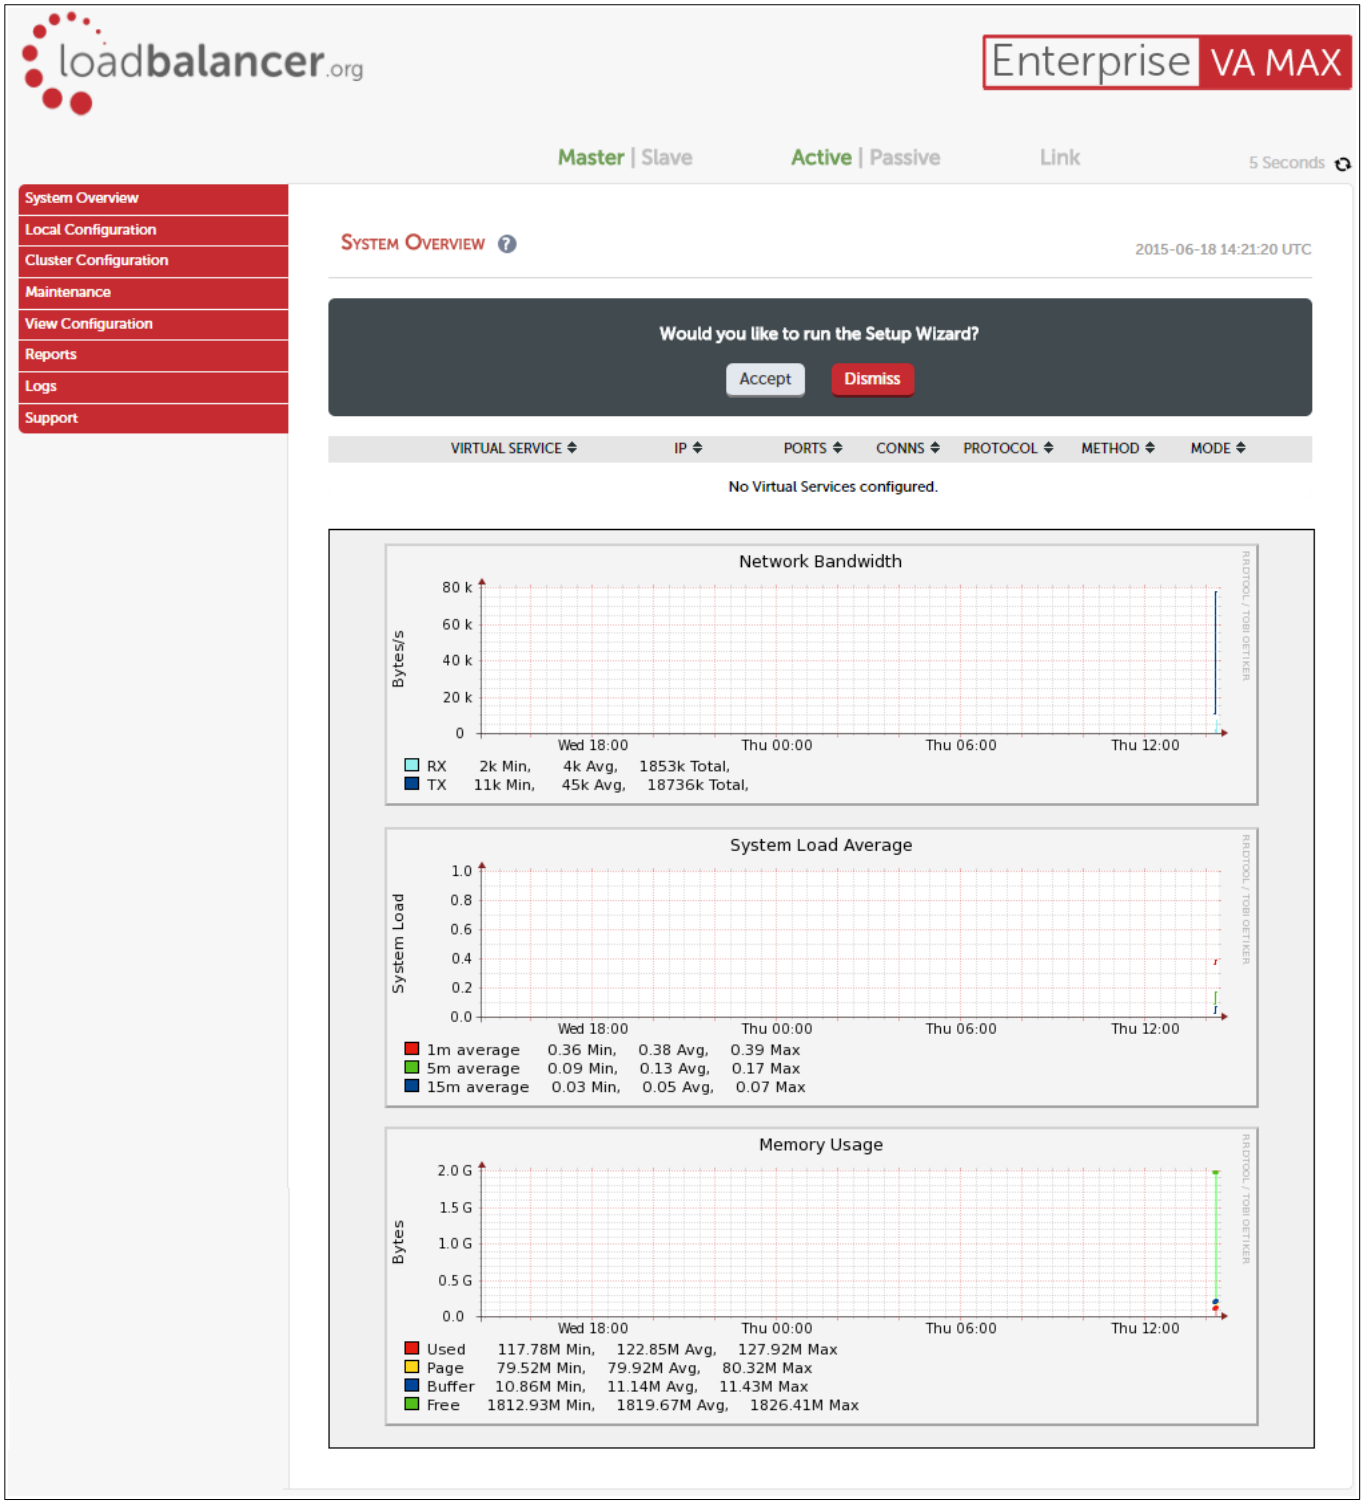

HA Clustered Pair Configuration

Loadbalancer.org recommend that load balancer appliances are deployed in pairs for high availability. In this guide a single unit is deployed first, adding a secondary slave unit is covered in section 5 of the Appendix on page 39.

10. Appliance Configuration for Exchange 2013 – Using DR Mode

Configure VIP1 – CAS 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:

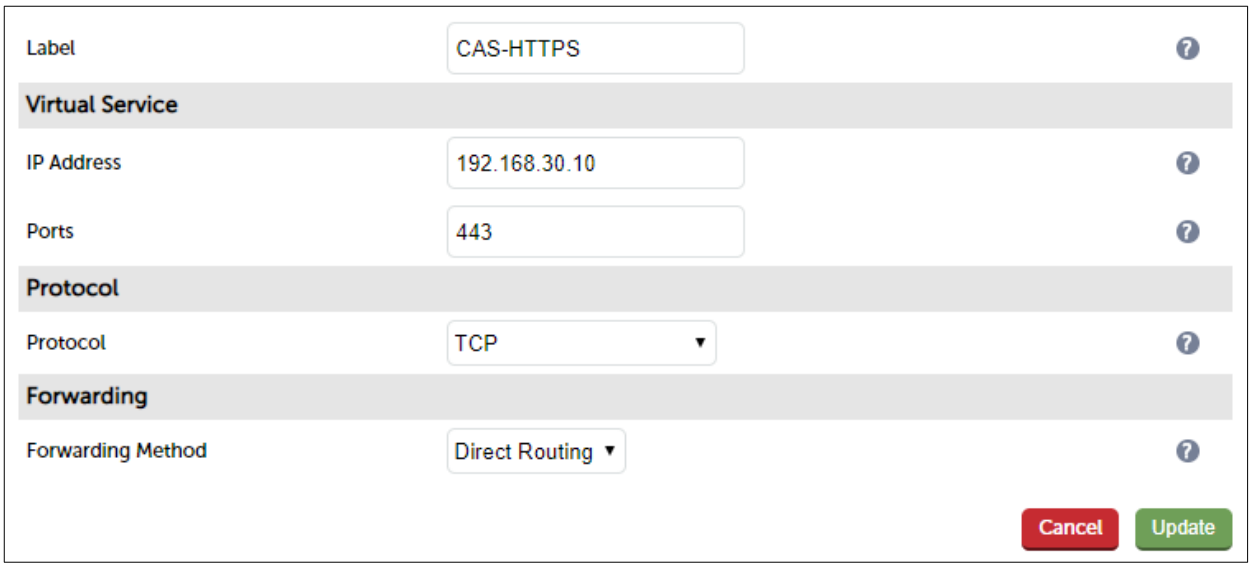

3. Enter an appropriate label for the VIP, e.g. **CAS-HTTPS**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**
5. Set the *Virtual Service Ports* field to **443**
6. Leave *Protocol* set to **TCP**
7. Leave *Forwarding Method* set to **Direct Routing**
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 severs and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server into the active pool.

11. Un-check the *Persistence* option
12. Set *Check Type* to **Negotiate**
13. Set *Protocol* to **HTTPS**
14. Set *Request to send* to **owa/healthcheck.htm**

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

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

b) Setting up the Real Servers

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

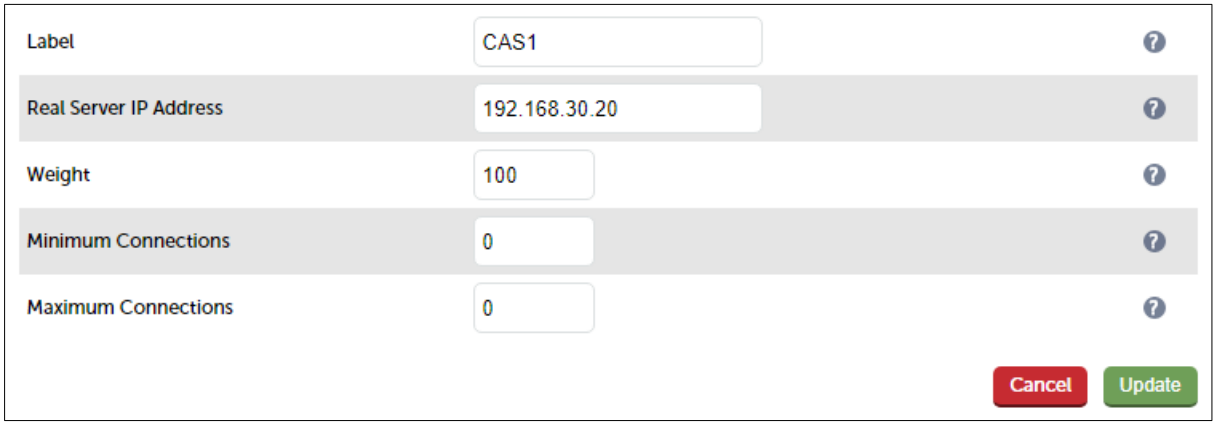

3. Enter an appropriate label for the RIP, e.g. **CAS1**
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**
5. Click **Update**
6. Repeat the above steps to add your other CAS Server(s)

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 section 4 in the Appendix on page 38.

Configure VIP2 – CAS 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:

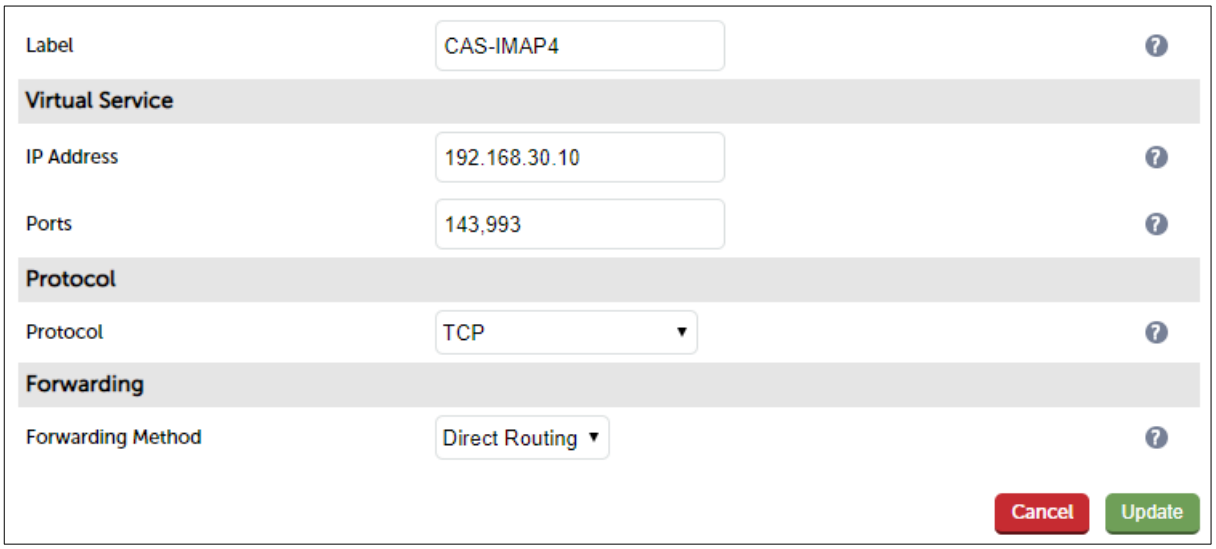

3. Enter an appropriate label for the VIP, e.g. **CAS-IMAP4**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**
5. Set the *Virtual Service Ports* field to **143,993**
6. Leave *Protocol* set to **TCP**
7. Leave *Forwarding Method* set to **Direct Routing**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

severs 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:

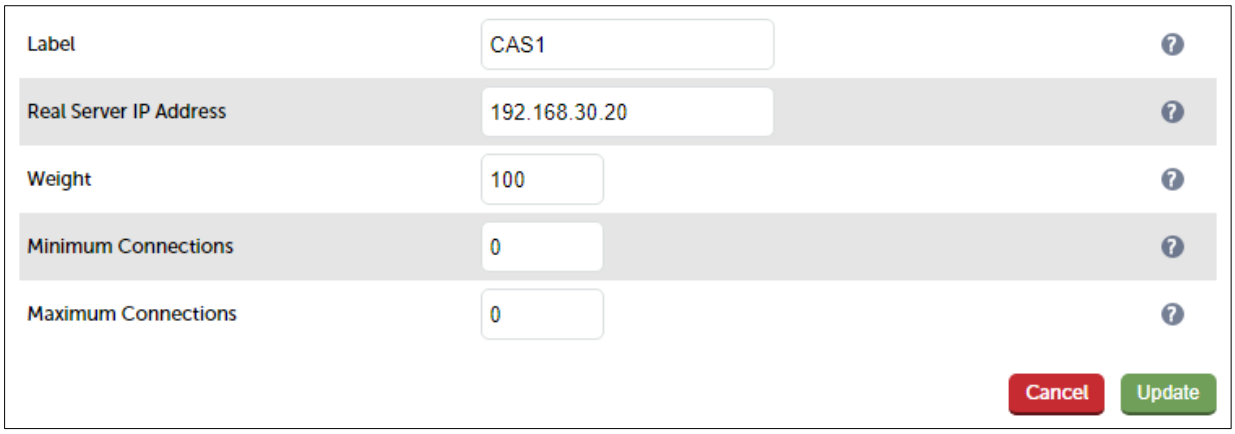

3. Enter an appropriate label for the RIP, e.g. **CAS1**
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**
5. Click **Update**
6. Repeat the above steps to add your other CAS Server(s)

Configure VIP3 – CAS 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:

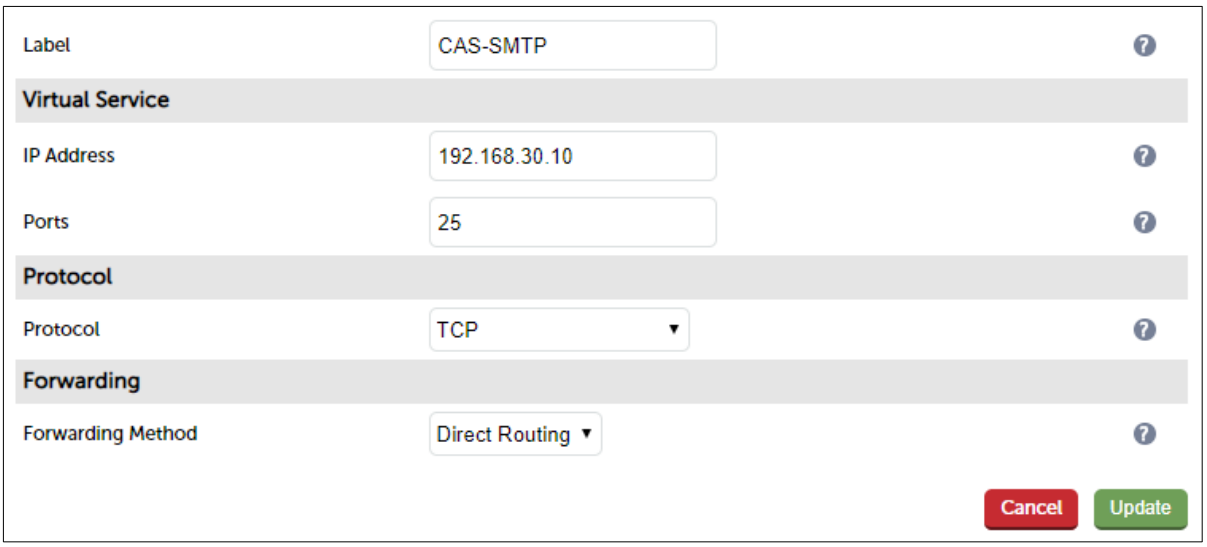

3. Enter an appropriate label for the VIP, e.g. **CAS-SMTP**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**
5. Set the *Virtual Service Ports* field to **25**
6. Leave *Protocol* set to **TCP**
7. Leave *Forwarding Method* set to **Direct Routing**
8. 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:

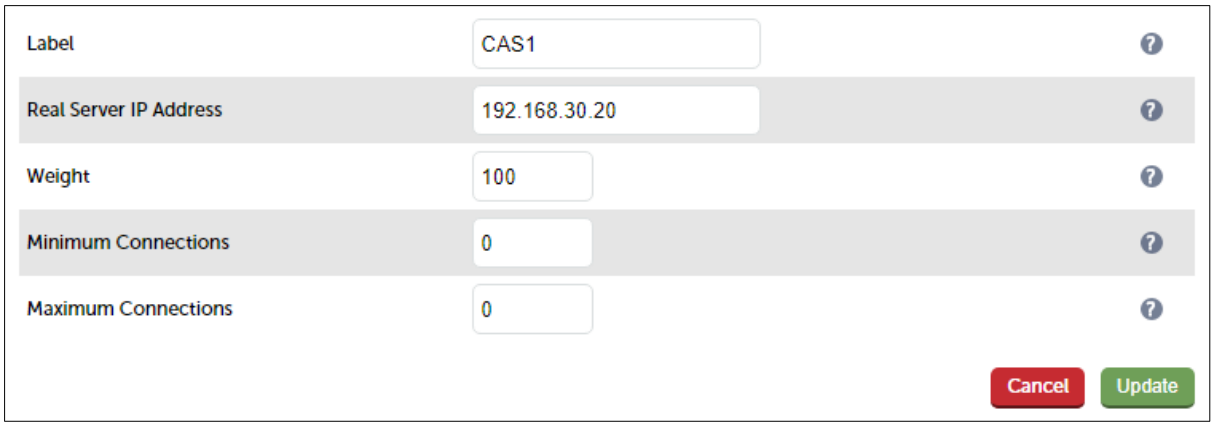

3. Enter an appropriate label for the RIP, e.g. **CAS1**

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 CAS Server(s)

11. Appliance Configuration for Exchange 2013 – Using SNAT Mode

Configure VIP1 – CAS 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:

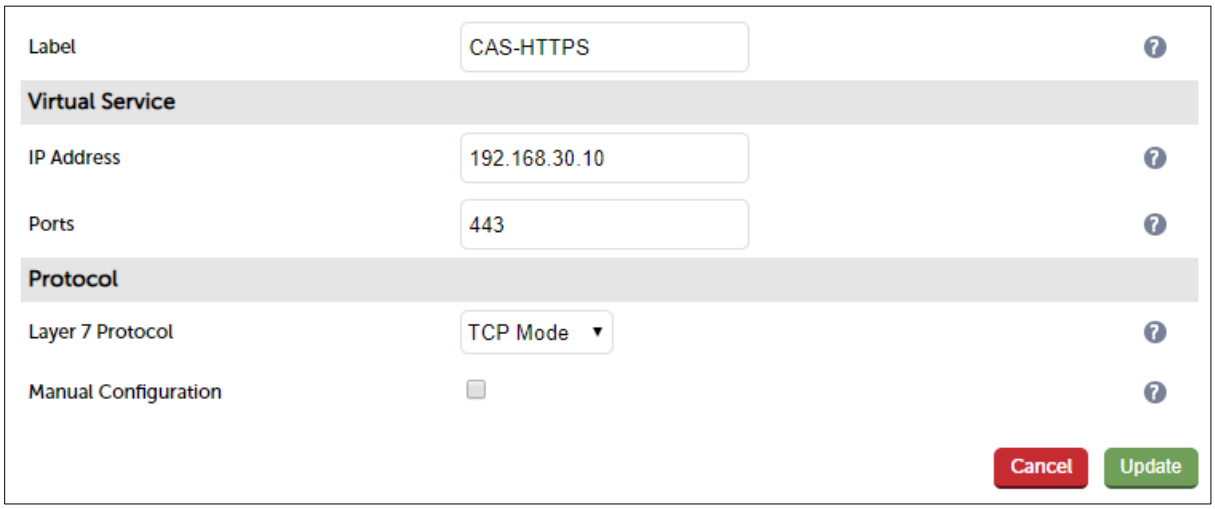

3. Enter an appropriate label for the VIP, e.g. **CAS-HTTPS**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**
5. Set the *Virtual Service Ports* field to **443**
6. Set *Layer 7 Protocol* set to **TCP Mode**
7. Click **Update**
8. Now click **Modify** 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 severs and therefore send all new requests to the new server. The trade off here is that using Round Robin will mean that server load may remain unbalanced for some time after bringing a new server

into the active pool.

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

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

13. Leave *Response expected* blank, this will configure the load balancer to look for a '200 OK' response
14. Scroll down to the *Other* section and click **[Advanced]**
15. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **30m** (i.e. 30 minutes)
16. 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:

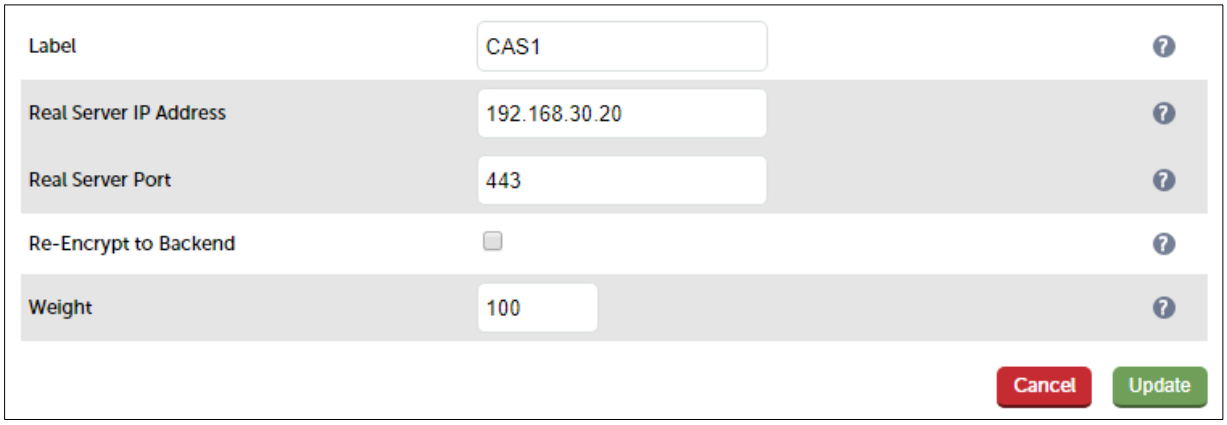

3. Enter an appropriate label for the RIP, e.g. **CAS1**
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**
5. Change the *Real Server Port* field to **443**
6. Click **Update**

7. Repeat the above steps to add your other CAS Server(s)

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 section 4 in the Appendix on page 38.

Configure VIP2 – CAS 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:

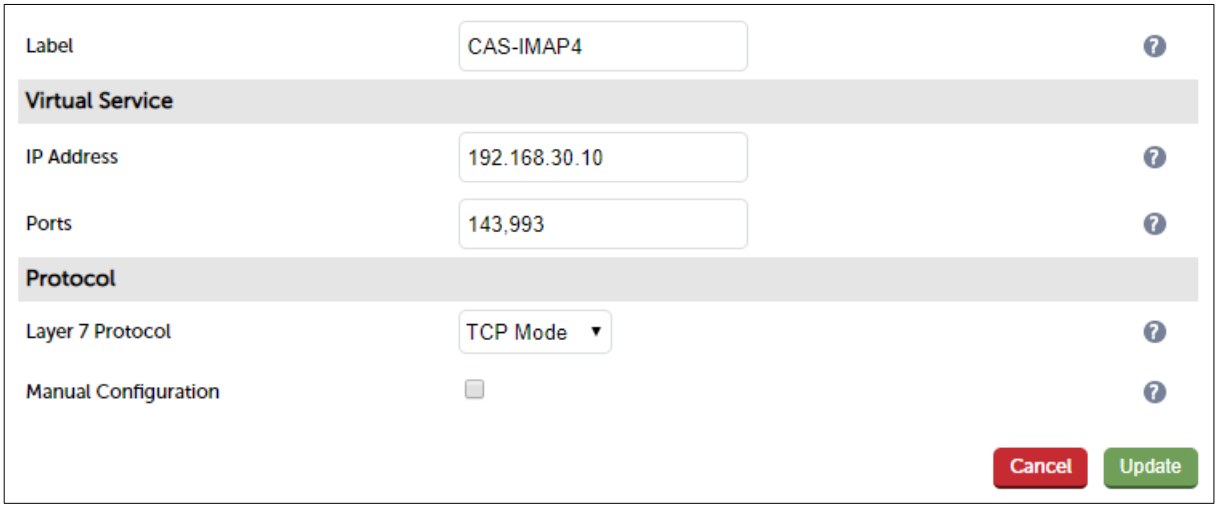

3. Enter an appropriate label for the VIP, e.g. **CAS-IMAP4**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**
5. Set the *Virtual Service Ports* field to **143,993**
6. Set *Layer 7 Protocol* to **TCP Mode**
7. Click **Update**
8. Now click **Modify** 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 severs 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. **CAS1**

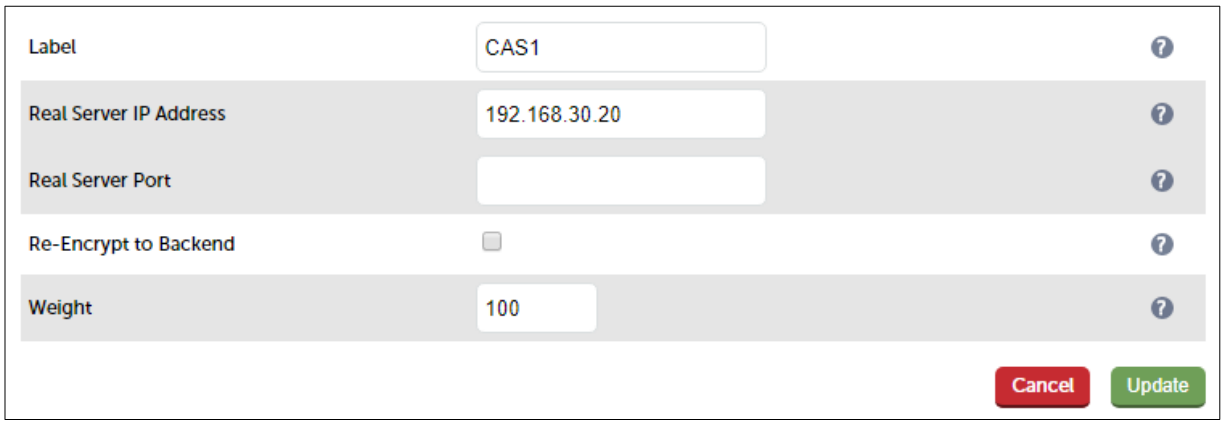

4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**
5. Leave the *Real Server Port* field blank
6. Click **Update**
7. Repeat the above steps to add your other CAS Server(s)

Configure VIP3 – CAS Role SMTP Services

Note: It's important to remember that when using SNAT mode (HAProxy), the source IP address of

packets reaching the Exchange Servers will be the IP address of the load balancer and ***not*** the source IP address of the client.

Transparency is normally only an issue for SMTP traffic at the receive connector. System Administrators typically want to lock down receive connectors to accept SMTP connections only from a controlled set of devices such as external smart mail hosts, printers, networked photocopiers etc.

If transparency for SMTP is the only issue, there are a number of options available to address this:

Option 1 – Enable full layer 7 transparency using TProxy. This is covered in section 1 of the Appendix on page 36.
Option 2 – Use the load balancers on-board firewall to lock down inbound SMTP connections rather than doing this at the receive connector. This is covered in section 2 of the Appendix on page 36.

Option 3 – Configure a layer 4 Virtual Service for SMTP rather than a layer 7 (HAProxy) based Virtual Service. Layer 4 is transparent by default so the source IP address is maintained. This is covered in section 3 of the Appendix on page 37.

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:

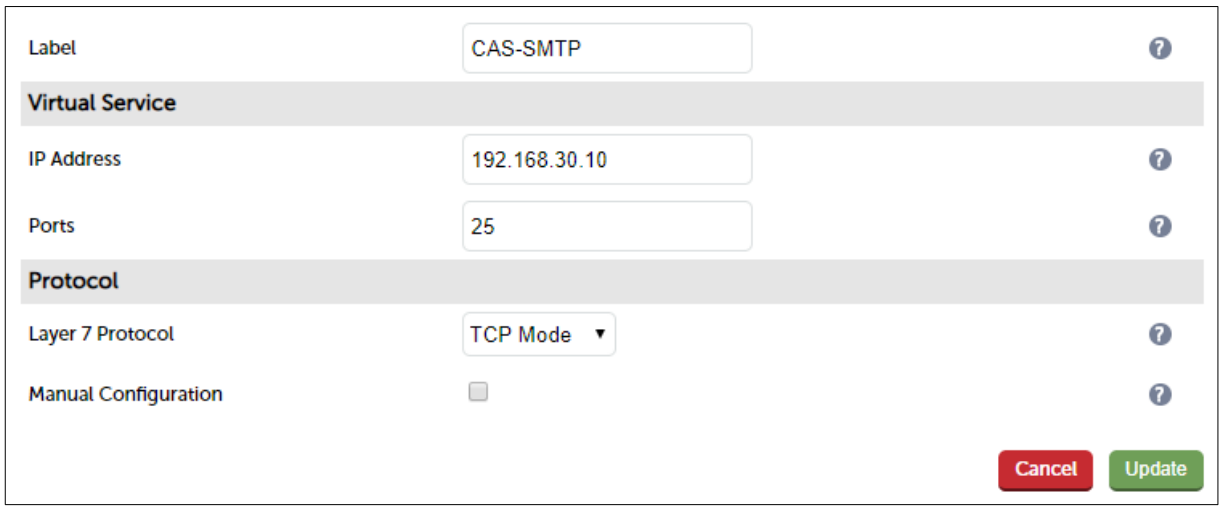

3. Enter an appropriate label for the VIP, e.g. **CAS-SMTP**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.30.10**
5. Set the *Virtual Service Ports* field to **25**
6. Set *Layer 7 Protocol* to **TCP Mode**
7. Click **Update**
8. Now click **Modify** next to 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
2. **Add a new Real Server** next to the newly created VIP
3. Enter the following details:

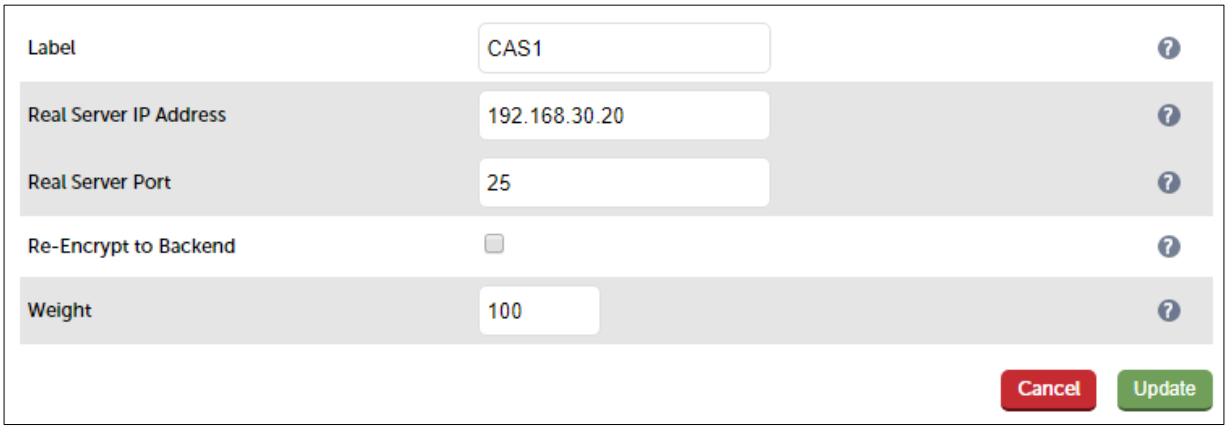

4. Enter an appropriate label for the RIP, e.g. **CAS1**
5. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.30.20**
6. Change the *Real Server Port* field to **25**
7. Click **Update**
8. Repeat the above steps to add your other CAS Server(s)

Additional Settings If Using Kerberos Authentication

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

To increase the Request buffer size:

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

Finalizing The Configuration

To apply the new settings, HAProxy must be restarted as follows:

3. Go to *Maintenance* > *Restart Services* and click **Restart HAProxy**

12. Testing & Verification

Useful Exchange 2013 & Other Microsoft Tools

Testing Server Health-checks Using Set-ServerComponentState

The Exchange Management shell cmdlet Set-ServerComponentState can be used to verify that the load balancer is correctly health-checking the Exchange servers.

In this guide, the health-check verifies that the owa virtual directory can be accessed.

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

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

Where <SERVER> is the hostname of the Exchange Server

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

To bring it back online, use the following command:

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

Where <SERVER> is the hostname of the Exchange Server

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

Exchange Management Shell:

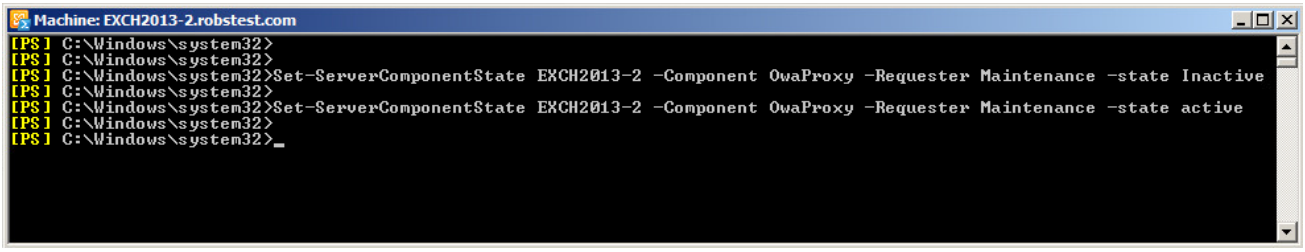

Testing Mailflow

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 rob@robstest.com
```

Exchange Management Shell:

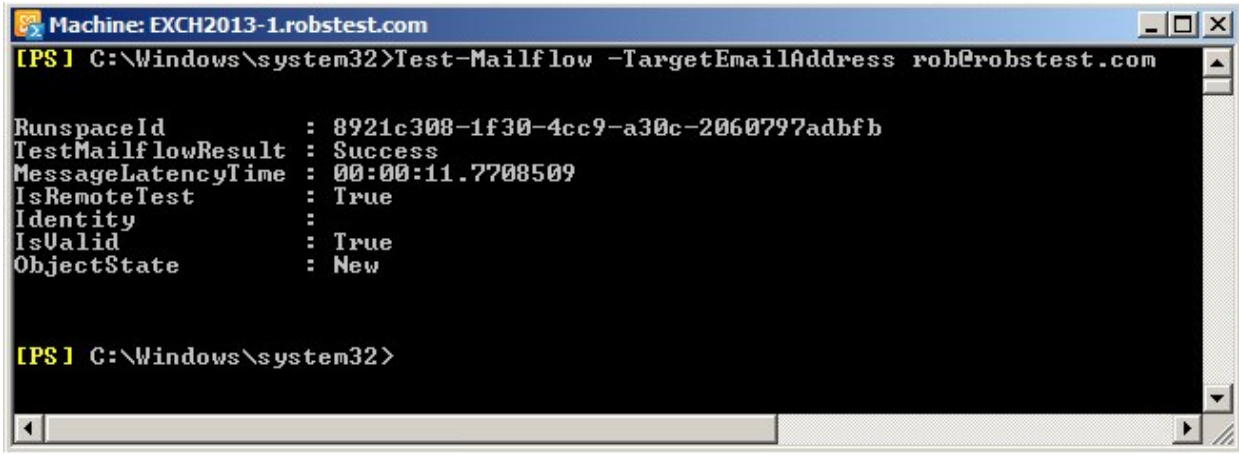

If everything is working correctly, a new message will appear in the test users mailbox:

Testing SMTP Mail Flow Using Telnet

SMTP can be tested using telnet to connect to port 25, then by issuing various commands to simulate an email being sent. Using *System Overview* in the WebUI, each CAS 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:

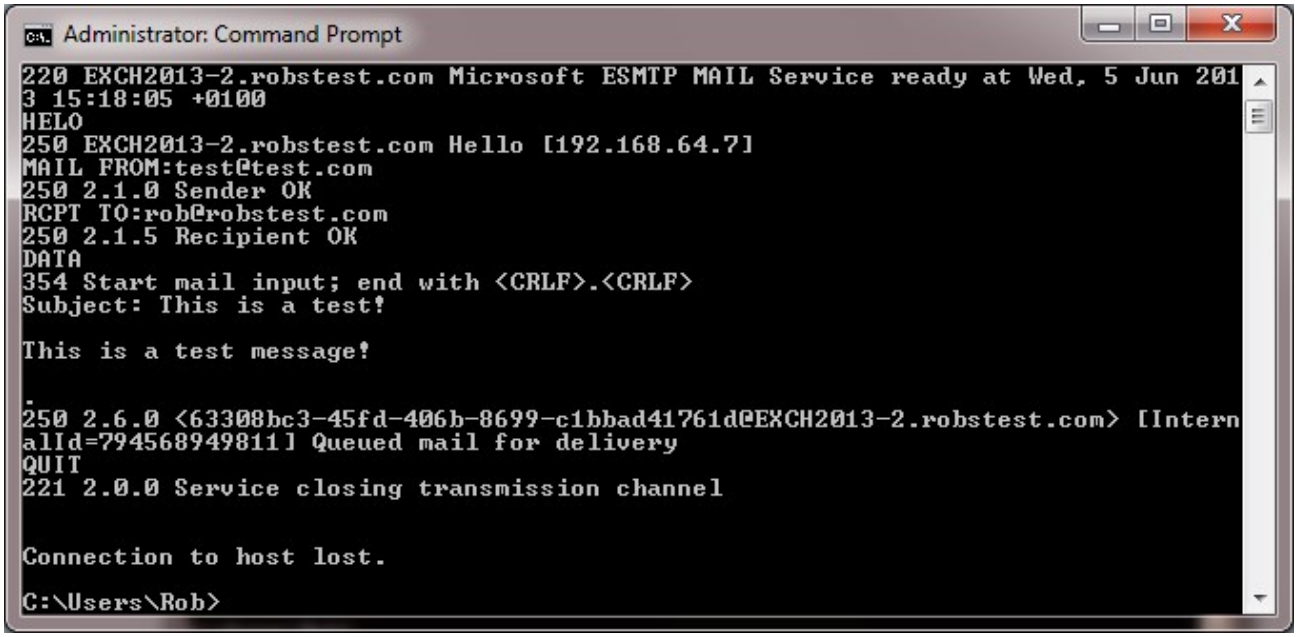

If everything is working correctly, a new message will appear in the test users mailbox:

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

```
telnet mail.robstest.com 25
```

Microsoft Exchange Testing Tool

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

Useful Appliance based Tools & Features

Using System Overview

The System Overview is accessed using 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. Th example below shows that both CAS servers are healthy and available to accept connections.

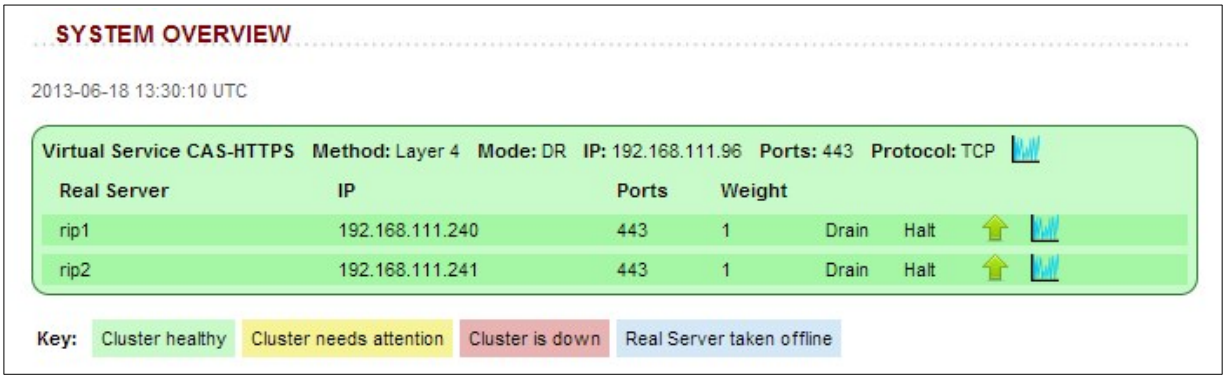

The example below shows that rip2 has been put in halt mode:

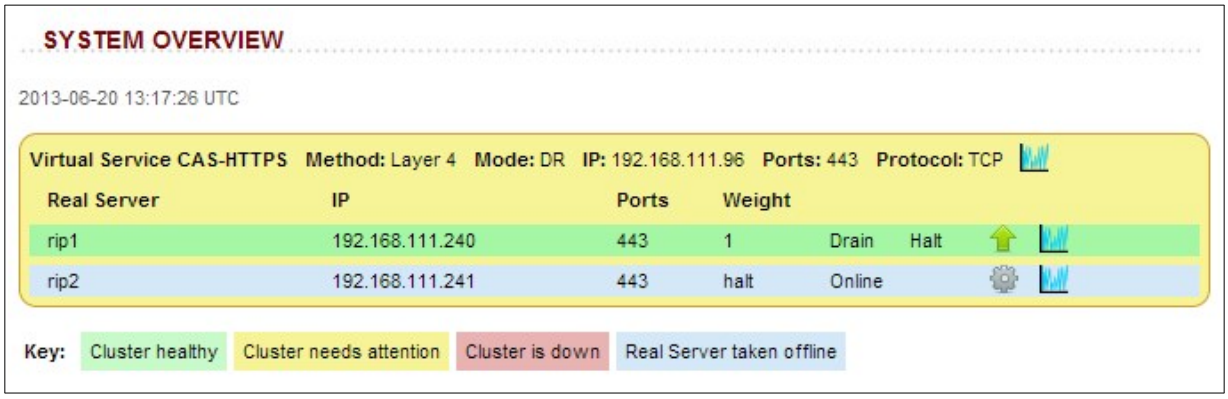

Layer 4 Status Report

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

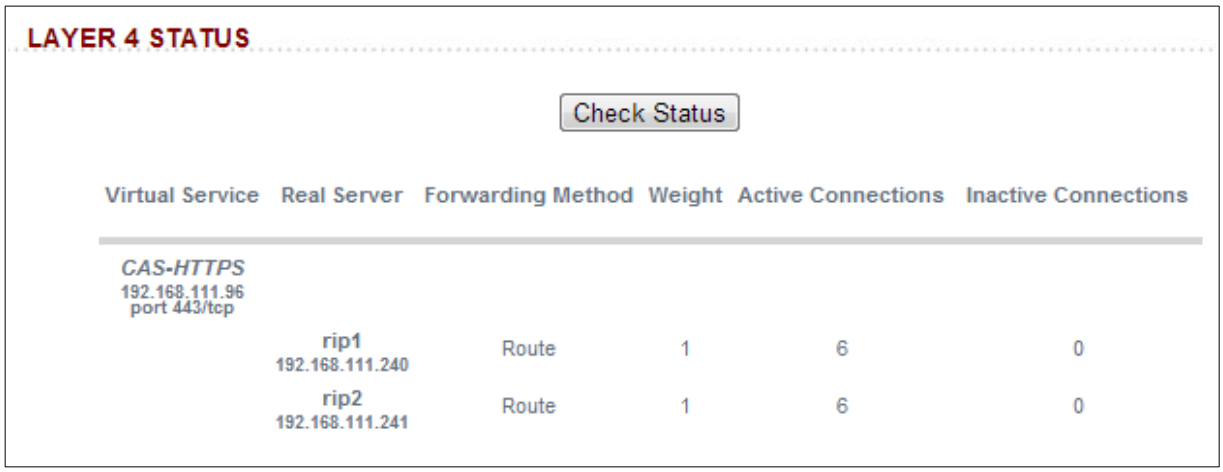

Layer 7 Statistics Report

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

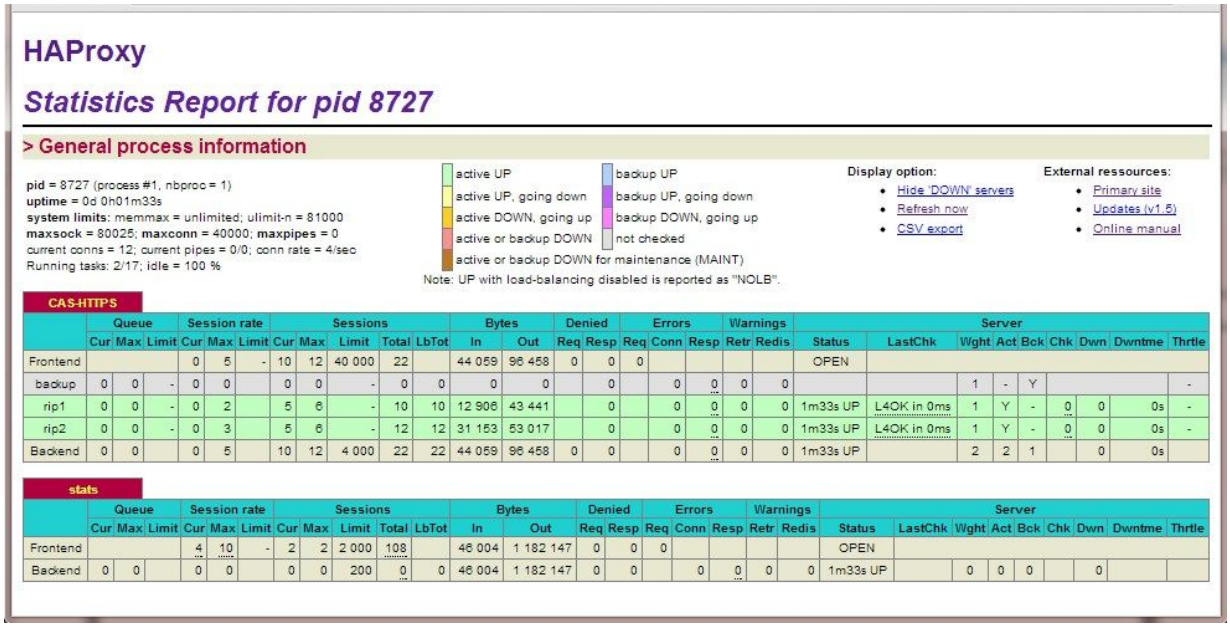

HAProxy

Statistics Report for pid 8727

> General process information

pid = 8727 (process #1, nbproc = 1)
uptime = 0d 0h01m33s
system limits: memmax = unlimited; ulimit-n = 81000
maxsock = 80025; maxconn = 40000; maxpipes = 0
current conns = 12; current pipes = 0/0; conn rate = 4/sec
Running tasks: 2/17; idle = 100 %

active UP | backup UP
active UP, going down | backup UP, going down
active DOWN, going up | backup DOWN, going up
active or backup DOWN | not checked
active or backup DOWN for maintenance (MAINT)
Note: UP with load-balancing disabled is reported as "NOLB".

Display option:
- Hide 'DOWN' servers
- Refresh now
- CSV export

External ressources:
- Primary site
- Updates (v1.5)
- Online manual

CAS-HTTPS

	Queue			Session rate			Sessions					Bytes		Denied		Errors			Warnings		Server							
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle
Frontend				0	5	-	10	12	40 000	22		44 059	96 458	0	0	0					OPEN							
backup	0	0	-	0	0		0	0	-	0	0	0	0		0		0	0	0	0			1	-	Y				-
rip1	0	0	-	0	2		5	6	-	10	10	12 906	43 441		0		0	0	0	0	1m33s UP	L4OK in 0ms	1	Y	-	0	0	0s	-
rip2	0	0	-	0	3		5	6	-	12	12	31 153	53 017		0		0	0	0	0	1m33s UP	L4OK in 0ms	1	Y	-	0	0	0s	-
Backend	0	0		0	5		10	12	4 000	22	22	44 059	96 458	0	0		0	0	0	0	1m33s UP		2	2	1		0	0s	

stats

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

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*, and then viewed using the 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

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

15. Conclusion

Loadbalancer.org appliances provide a very cost effective solution for highly available load balanced Exchange 2013 environments.

16. Appendix

1 – Enabling Layer 7 Transparency using TProxy

As mentioned previously, Layer 7 SNAT mode is non-transparent by default. If a fully transparent configuration is required, TProxy can be used. The main points to note are that two subnets must be used and the default gateway on the Exchange Servers must be set to be the load balancer.

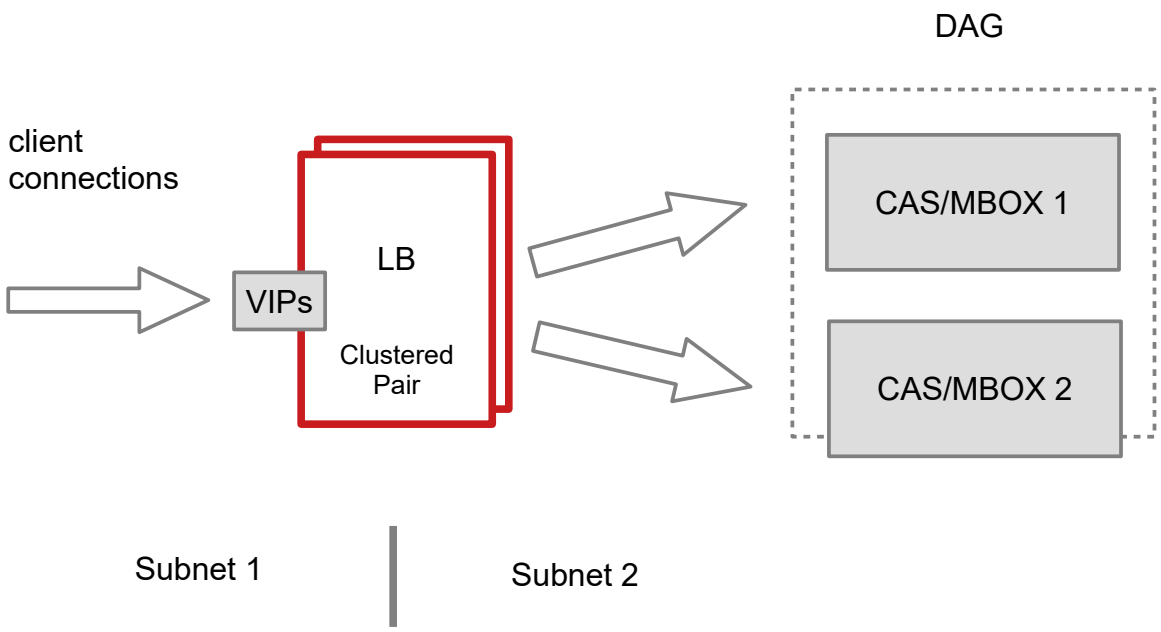

Key points to note:

- 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
- 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 allow failover to the slave
- TProxy must be enabled using the WebUI: Go to *Cluster Configuration* > *Layer 7 – Advanced Configuration* and set **Transparent Proxy** to 'On' and 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.

2 – Limiting inbound SMTP Connections using Firewall Rules

Since layer 7 is not transparent by default, it's not possible to filter inbound SMTP connections by IP address at the receive connector on the Hub Transport Server. One way to address this is to add firewall rules to the load balancer to limit which hosts can connect inbound on port 25.

Rules can be added using the WebUI option: *Maintenance > Firewall Script*. Simply copy/paste/edit the examples below into the firewall script then click **Update**.

> 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.

EXAMPLES:

1) to 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.

2) to 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.

> 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.

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 page 9.

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

Layer 4 DR Mode - Solving the ARP Problem:

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

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

Configuring the Real Servers in this way is referred to as '*Solving the ARP problem*". The steps required depend on the particular version of Windows being used. For detailed steps on solving the ARP problem for Windows 2012/2016 Please refer to page 41.

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**

The steps:

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**

5 – Clustered Pair Configuration – Adding a Slave Unit

If you initially configured just the master unit and now need to add a slave - our recommended procedure, please refer to the relevant section below for more details:

Note: A number of settings are not replicated as part of the master/slave pairing process and therefore must be manually configured on the slave appliance. These are listed below:

- Hostname & DNS settings
- Network settings including IP addresses, bonding configuration and VLANs
- Routing configuration including default gateways and static routes
- Date & time settings
- Physical – Advanced Configuration settings including Internet Proxy IP address & port, Firewall table size, SMTP relay and Syslog server
- SNMP settings
- Graphing settings
- Firewall Script & Firewall Lockdown Script settings
- Software updates

Version 7:

Please refer to **Chapter 8 – Appliance Clustering for HA** in the v7 Administration Manual.

Version 8:

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

- Deploy a second appliance that will be the slave and configure initial network settings
- Using the WebUI, navigate to: *Cluster Configuration > High-Availability Configuration*

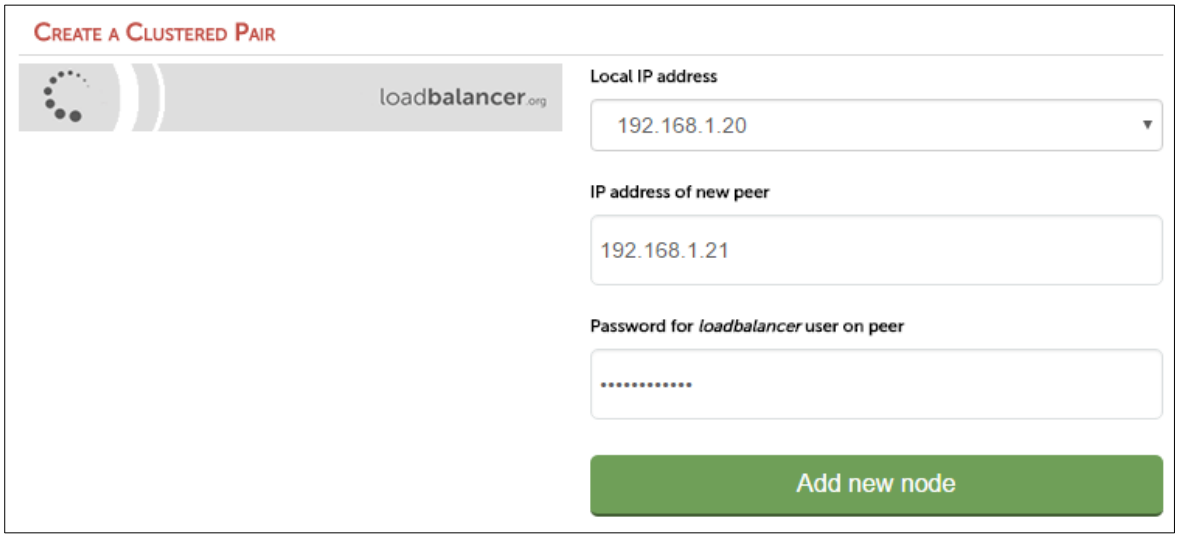

- Specify the IP address and the *loadbalancer* users password (the default is 'loadbalancer') for the slave (peer) appliance as shown above
- Click **Add new node**
- The pairing process now commences as shown below:

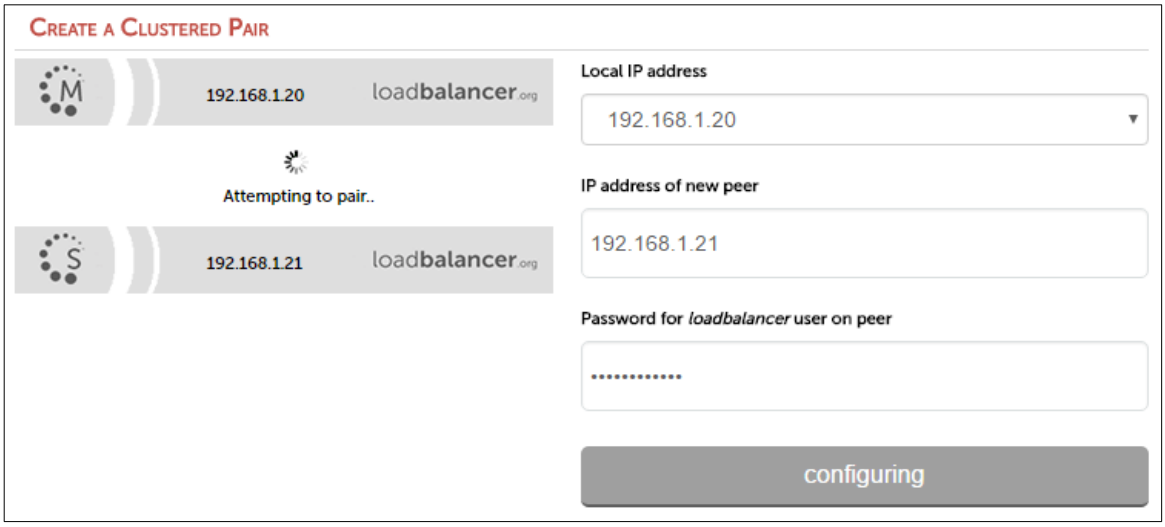

- Once complete, the following will be displayed:

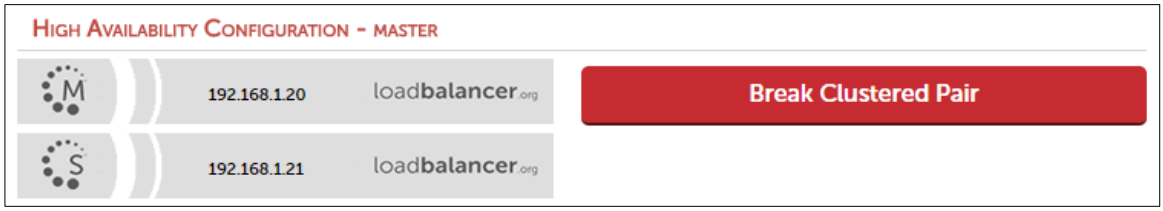

- To finalize the configuration, restart heartbeat and any other services as prompted in the blue message box at the top of the screen

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

Note: Please refer to chapter 9 – Appliance Clustering for HA in the Administration Manual for more detailed information on configuring HA with 2 appliances.

6 – Solving the ARP Problem

When using Layer 4 DR mode, the ARP problem must be solved. This involves configuring each Real Server to be able to receive traffic destined for the VIP, and ensuring that each Real Server does not respond to ARP requests for the VIP address – only the load balancer should do this.

The steps below are for Windows 2012 / 2016, for other versions of Windows please refer to chapter 6 in the Administration Manual.

Step 1: Install the Microsoft Loopback Adapter

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

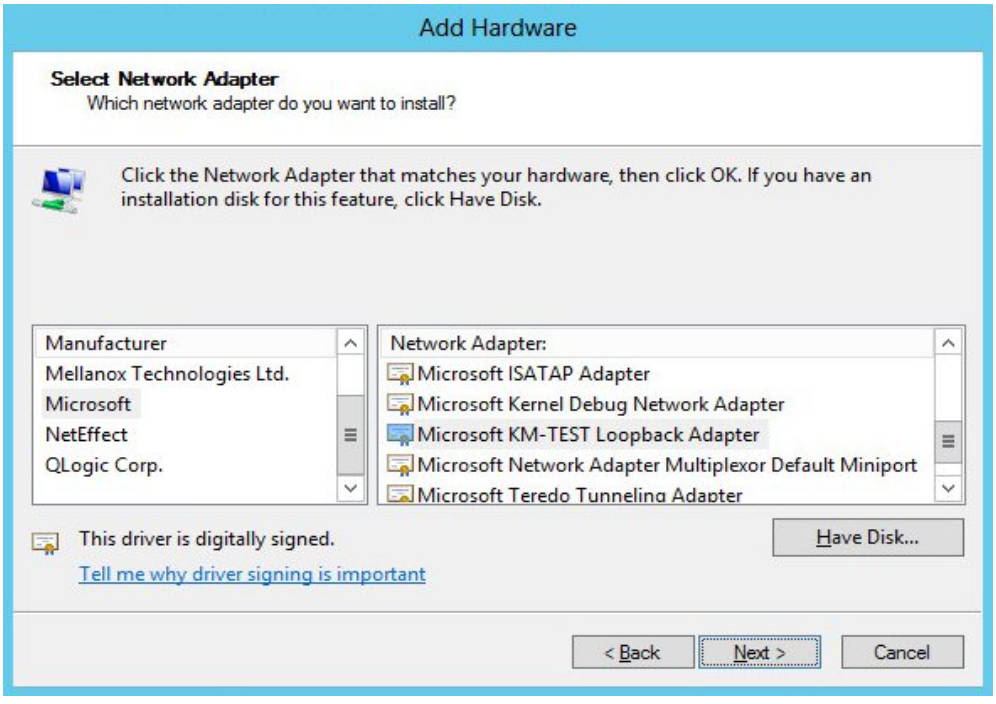

6. Click **Next** to start the installation, when complete click **Finish**

Step 2: Configure the Loopback Adapter

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

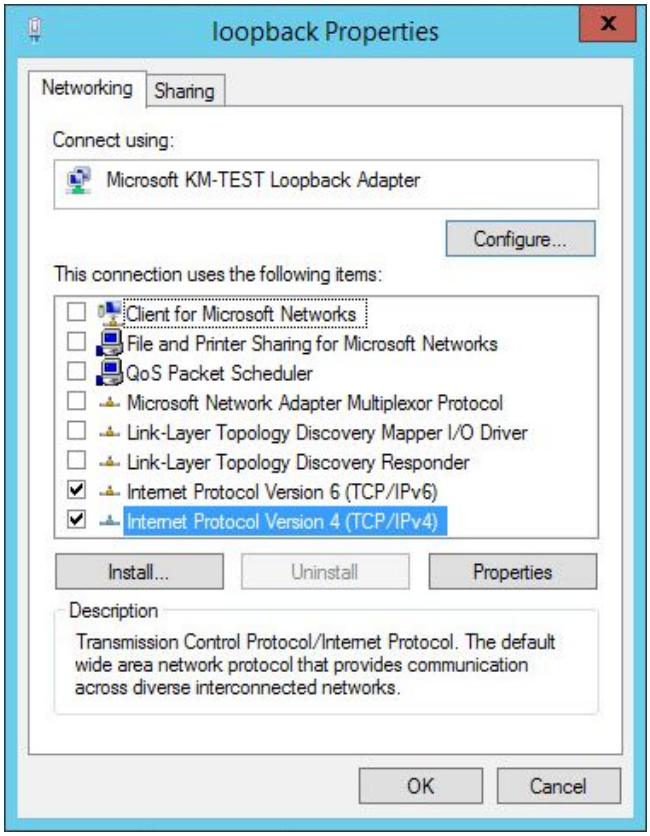

Note: Leaving both checked ensures that both IPv4 and IPv6 are supported. Select one If preferred.

5. If configuring IPv4 addresses, select **Internet Protocol Version (TCP/IPv4)**, click **Properties** and configure the IP address to be *the same as the address you've used for the Virtual Service (VIP)* with a subnet mask of 255.255.255.255 , e.g. 192.168.30.10/255.255.255.255 as shown below:

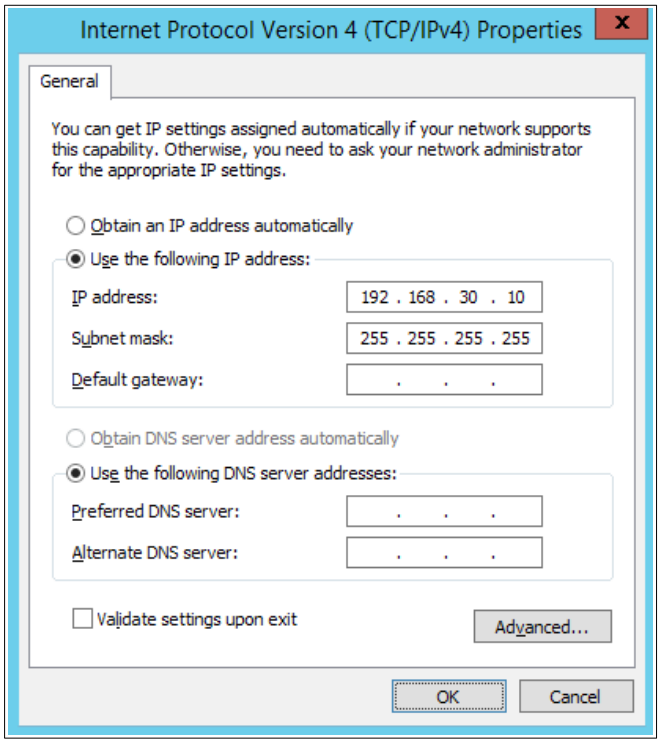

6. If configuring IPv6 addresses select **Internet Protocol Version (TCP/IPv6)**, click **Properties** and configure the IP address to be *the same as the address you've used for the Virtual Service (VIP)* and set the *Subnet Prefix Length* to be the same as your network setting , e.g. 2001:470:1f09:e72::15/64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

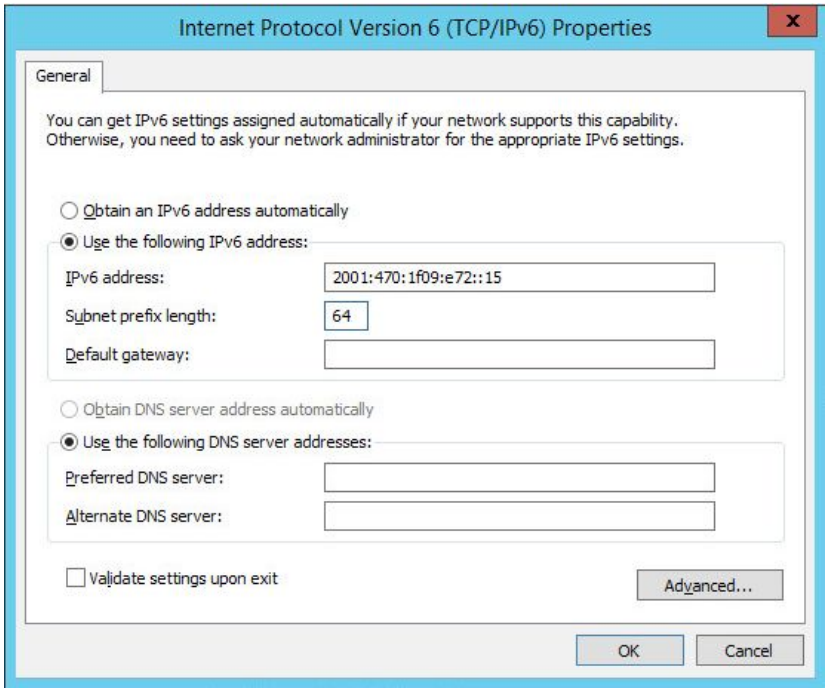

7. Click **OK** on TCP/IP Properties, then click **Close** on Ethernet Properties to save and apply the new settings
8. Now repeat the above process on the other Windows 2012/2016 Real Servers

Step 3: Configure the strong/weak host behavior

Windows Server 2000 and Windows Server 2003 use the weak host model for sending and receiving for all IPv4 interfaces and the strong host model for sending and receiving for all IPv6 interfaces. You cannot configure this behavior. The Next Generation TCP/IP stack in Windows 2008 and later supports strong host sends and receives for both IPv4 and IPv6 by default. To ensure that Windows 2012/2016 is running in the correct mode to be able to respond to the VIP, the following commands must be run on each Real Server:

For IPv4 addresses:

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

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

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

For IPv6 addresses:

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

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

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

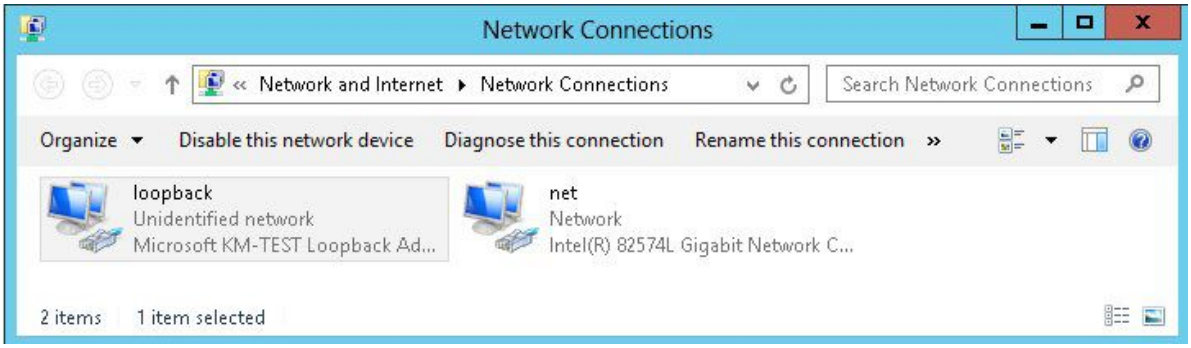

Note: The names for the NICs are case sensitive, so make sure that the name used for the interface and the name used in the commands match exactly.

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

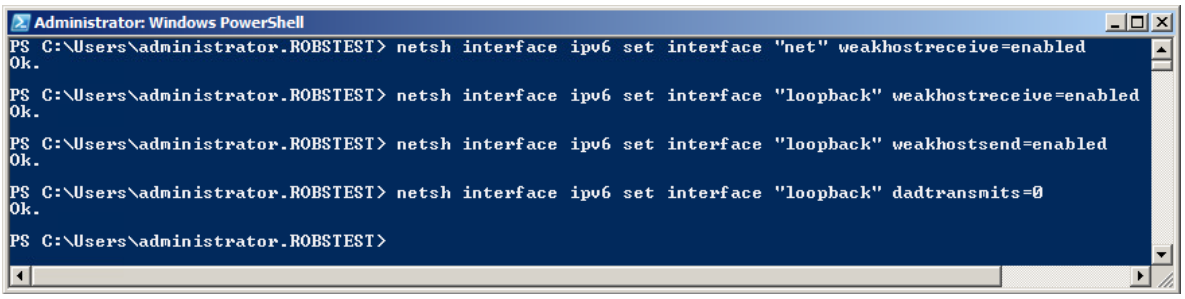

Note: This shows an IPv6 example, use the IPv4 commands if you're using IPv4 addresses.

Repeat steps 1 – 3 on all remaining Exchange Server(s).

17. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.5.0	6 August 2019	Styling and layout	General styling updates	RJC
1.5.1	17 January 2020	Added note explaining how to disable "Secure Mode" to unlock the firewall script page	Required update	RJC
1.5.2	3 June 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.5.3	25 June 2021	Minor updates	Required update	RJC

About Loadbalancer.org

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

United Kingdom

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

Canada

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

United States

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

Germany

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