Load Balancing Microsoft Exchange 2013 with FortiADC

Highly Available, High Performing, and Scalable Deployment with FortiADC D-Series Appliances

Exchange 2013 and Application Delivery

Microsoft® Exchange Server 2013 was released in late 2012 as the successor to Microsoft Exchange Server 2010. It introduced a number of new features as well as changes to existing features. Enhancements were added with Cumulative Update 1 (CU1) in 2013 and again, later in 2013 with the release of Cumulative Update 2 (CU2).

This guide was written for CU2. This version of the guide provides configuration information for Fortinet's FortiADC 3.1 software release.

Load Balancing Requirements for Exchange

Microsoft recognizes the need for load balancing client access in all but the smallest Exchange deployments. For Microsoft's overview of load balancing recommendations in Exchange 2013, please see:

http://technet.microsoft.com/en-us/library/jj898588%28v=exchg.150%29.aspx

As stated in the above document, Exchange 2013 relies upon Client Access servers to provide connections for users. Load balancing these Client Access servers can improve availability, by providing redundancy, as well as efficiency through intelligently balancing load across the servers.

> *"In Exchange Server 2010, client connections and processing were handled by the Client Access server role. This required that both external and internal Outlook connections, as well as mobile device and third-party client connections, be load balanced across the array of Client Access servers in a deployment to achieve fault tolerance and efficient utilization of servers."*

> *"In Exchange Server 2013, there are two primary types of servers—the Client Access server and the Mailbox server. The Client Access servers in Exchange 2013 serve as lightweight, stateless proxy servers, allowing clients to connect to Exchange 2013 Mailbox servers. Exchange 2013 Client Access servers provide a unified namespace and authentication."*

While software load balancers and reverse proxy solutions can be adequate for smaller deployments, larger deployments will benefit from the features and capacity provided by a hardware load balancer. Among other issues, Microsoft recognizes the following limitations with Windows Network Load Balancing, the most popular software based load balancing solution for Exchange:

Important Note:

This guide is written only for the **FortiADC D-series platform**. The instructions included within are not designed to be used with the FortiADC E-series platform application delivery controllers.

WNLB can't be used on Exchange servers where mailbox DAGs are also being used because WNLB is incompatible with Windows failover clustering. If you're using an Exchange 2013 DAG and you want to use WNLB, you need to have the Client Access server role and the Mailbox server role running on separate servers.

WNLB doesn't detect service outages. WNLB only detects server outages by IP address. This means that if a particular web service, such as Outlook Web App, fails, but the server is still functioning, WNLB won't detect the failure and will still route requests to that Client Access server. Manual intervention is required to remove the Client Access server experiencing the outage from the load balancing pool.

Using WNLB can result in port flooding, which can overwhelm networks.

Because WNLB only performs client affinity using the source IP address, it's not an effective solution when the source IP pool is small. This can occur when the source IP pool is from a remote network subnet or when your organization is using network address translation.

The FortiADC Difference

There are a number of hardware load balancing products available on the market with a wide range of features and capabilities. FortiADC differentiates itself by providing superior value; advanced acceleration features, high performance, reliability, and security from a market leader.

FortiADC not only load balances Internet service requests across multiple servers, but also accelerates application performance and provides application aware features that monitor server load and improve server response times – by as much as 25%. In addition to basic load balancing, FortiADC provides:

- Automatic server and application health monitoring
- Intelligent, application aware load balancing policies (least connections, fastest response time, static weight, and round robin)
- Redundant High Availability (HA) configurations

For more information on how FortiADC can make your applications work better, faster, and more economically, please visit www.fortinet.com.

Using FortiADC with Exchange 2013

For the purposes of this deployment guide, we assume a working Exchange deployment that will be augmented by the addition of FortiADC (or a pair of FortiADCs in a failover configuration). If you are setting up a new deployment of Exchange, we recommend that you first set up your Exchange configuration without FortiADC, verify each of your intended client access methods, and then follow this document to deploy FortiADC into that configuration.

Logically, FortiADC sits in between clients accessing Exchange and the Exchange servers, as shown in the following diagram:

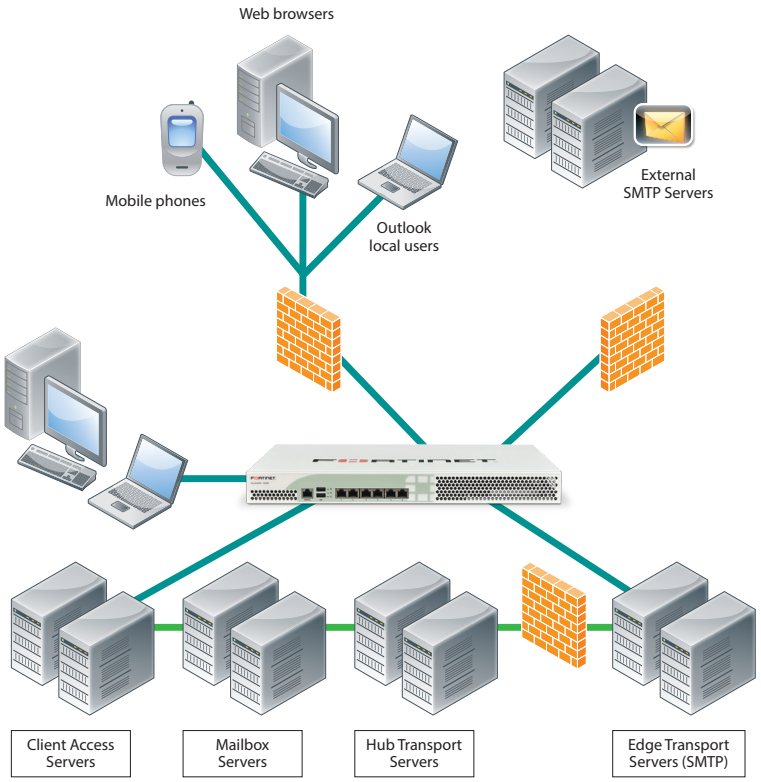

Clients can access Exchange via a number of applications and protocols (generally called Exchange services in the remainder of this guide):

Outlook Web App (OWA) – (known as Outlook Web Access in previous Exchange releases). Internal and external clients initiate OWA sessions over HTTP using a web browser, or Outlook Web App Light.

Outlook Anywhere (OA) – Outlook clients access Exchange by tunneling the Outlook MAPI (Messaging Application Programming Interface) protocol over an HTTP connection.

ActiveSync (AS) – Mobile clients can synchronize with Exchange services, which push data to the mobile device, over an HTTP connection.

POP3 and IMAP4 – External and internal third-party mail programs use these protocols (Post Office Protocol v3 and Internet Message Access protocol v4) to retrieve and send email.

Remote Procedure Call Client Access (RPC CA) – RPC CA services include the PortMapper, MAPI access to Outlook, and the AddressBook application. In Exchange 2013 all RPC services are tunneled via HTTPS.

SMTP – External mail servers forward mail to Exchange through Edge Servers or Hub Transport Servers using the Simple Mail Transfer Protocol (SMTP).

All of the services above are routed through FortiADC and load balanced to the appropriate Client Access Servers.

NOTE: Load balancing in Exchange 2013 has been changed in 3 main ways from load balancing in Exchange 2010.

1. SSL Offload support has been removed. This means that instead of layer 7 load balancing being recommended, now layer 4 load balancing of port 443 traffic is the only method supported.
2. Connections to Client Access servers is now stateless. This means that there is no longer a need for persistent connections between the clients and the Client Access servers.
3. Outlook anywhere, formerly known as RPC over HTTP is now the default method for communication via RPC in Exchange 2013. All Outlook connectivity takes place over outlook anywhere.

Hardware and Software Used in This Guide

To develop this deployment guide, the following hardware and software was used:

- FortiADC model 200D
- OS Version 3.1 build 0094
- Custom Server hardware running VMware ESX 4
- Several VM servers running Microsoft Server 2012
- Microsoft Exchange Server 2013 CU2
- Appropriately configured clients to test client access

Note that the hardware and software required for your configuration will vary from the above depending on your testing and production environment. Microsoft Hyper-V, for example, could be used in place of VMware.

If you do not have locally available clients of all types, Microsoft offers two alternatives that you can use in a testing environment to validate your configuration prior to putting it into production:

- The Exchange Load Generator 2010 can be installed on a local server and can be configured to generate Exchange traffic for the various protocols.
- The Exchange Remote Connectivity Analyzer is an online tool that you configure to test your Internet accessible Exchange configuration.

Click on the links above for more details on configuring and using these tools.

Server Health Checks

The ADC has an ICMP healthcheck pre-defined but in order to check not just the overall health of the server but rather the availability of the Exchange services, we are going to define an additional healthcheck for TCP port 443.

Load Balancing Policy

IIn previous versions of Microsoft Exchange, Microsoft recommended using the **least connections** load balancing policy, which routes requests to servers that are more lightly loaded in terms of number of open connections. Microsoft has since changed its recommendation, since it is possible that using a load balancing algorithm like **least connections** can lead to overloading a server when it is first brought online. Microsoft now recommends using a policy that does not depend on weighted criteria, such as **round robin** – which simply routes requests evenly across all available servers, regardless of performance. The result is that while **round robin** makes it less likely that any one server will be overloaded when it is brought online, it leads to an imbalance in the distribution of requests across all servers when servers are brought offline and online.

It should be noted that FortiADC's **least connections** setting is less prone to overloading a new server due to the warm-up feature which can be defined individually on each server within a server pool. A setting of at least 3 seconds should be used when using least connections with Exchange 2013, this can be

adjusted if needed.

Besides **round robin** and **least connections**, FortiADC also offers these load balancing policies:

static weight -- distributes requests among the servers depending on their assigned initial weights. A server with a higher initial weight gets a higher percentage of the incoming requests. Think of this method as a weighted round robin implementation

fastest response -- dispatches the highest percentage of requests to the server with the shortest response time. FortiADC does this carefully: if FortiADC sends too many requests to a server, the result can be an overloaded server with slower response time. The fastest response policy optimizes the cluster-wide response time. The fastest response policy also checks the number of active connections and server agent values (if configured); but both of these have less of an influence than they do under the adaptive load balancing policy. For example, if a server's active connection count and server agent values are high, FortiADC might not dispatch new requests to that server even if that server's response time is the fastest in the cluster.

Source Network Address Translation (spoof) Setting

In general, Microsoft recommends using SNAT for all configurations unless specific requirements prevent it. The FortiADC has multiple options relating to NAT for each TCP Virtual Server. Direct Routing performs no NAT at all, NAT changes only the destination address from the Virtual Server IP to a real server's IP, and Full NAT changes both the source and destination addresses. In this configuration Full NAT will be used, requiring an IP Pool, specifying the source addresses, to be created.

Note that Direct Server Return (DSR) configurations are supported by Exchange, but in general SNAT-enabled configurations are recommended by Microsoft to avoid the additional complexity and drawbacks of DSR. For example: to use DSR with any load balancer requires configuration of a special loopback adapter on each server and is supported only for Layer 4 services.

See the FortiADC *Installation and Administration Guide* for more information on FortiADC network configuration.

SSL Offload and Acceleration

Although in Exchange 2010 Microsoft supported SSL offloading of HTTPS traffic to a device with dedicated hardware, in Exchange 2013, as of CU2 they have currently decided not to support SSL offloading.

The FortiADC Administrative Interface

This guide assumes that you have already set up FortiADC on your organization's network. Full instructions are available in the printed startup guides and CD-ROM delivered with your FortiADC. Documentation is also available from our support site at:

https://support.fortinet.com

Once FortiADC has an IP address on the network, open the Administrative Interface by opening the following URL in your web browser:

`http://<FortiADC_IP_addr>`

Where *<FortiADC_IP_addr>* is FortiADC's management IP address. Log in to FortiADC using a login with administrator privileges. This opens the graphical user interface, as shown in the following figure:

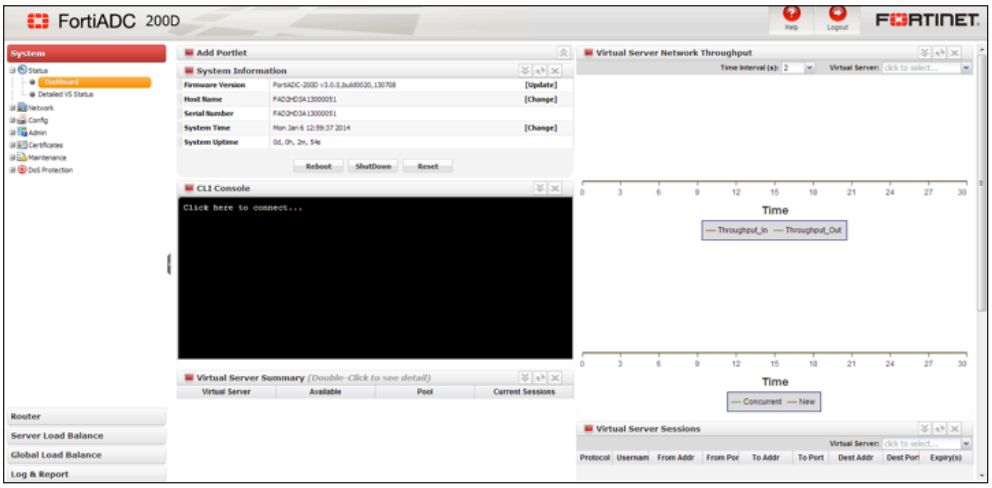

The clusters, servers, responders, and match rules you create for Exchange 2013 will be displayed in the left frame, while configuration details are displayed and modified in the right frame.

Virtual Server Configuration Summary

The following table summarizes the Virtual Servers that can be used for Exchange 2013 – A single Virtual Server can provide a virtual front-end for multiple Exchange services. In this case, for the services offered below, 3 Virtual Servers will be required. A fourth will be created to redirect client requests from port 80 to port 443 to help ensure that only secure connections will be accepted.

Exchange Service	Cluster Type	Cluster Port	Server Port	Affinity (Persistence)	Server Health
OWA	L4 TCP	443	443	None	TCP port 443
Outlook Anywhere	L4 TCP	443	443	None	TCP port 443
ActiveSync	L4 TCP	443	443	None	TCP port 443
HTTP redirect	HTTP	80	n/a	None	n/a
IMAP4	L4 TCP	993	993	None	TCP port 993
SMTP	L4 TCP	25	25	None	TCP port 25

Creating Resources for the Virtual Servers

In order to create Virtual Servers for use with Exchange 2013 we need to first configure the resources that they will require.

In this case we will need to create Health Checks for TCP ports 25, 443, and 993, a server Pool for each group of services offered via SMTP, HTTPS and IMAPS, and an IP Pool in order to enable the SNAT feature and ensure correct routing of return traffic.

These can all be configured within the expandable resources menu in the Server Load Balance section of the GUI.

Creating Health Checks

Server pools are groups of servers that can be assigned as a unit to an FortiADC virtual cluster: the IP address/port that clients access when requesting Exchange services.

In general, a server pool is required for each group of FortiADC servers that offer the same service, or same group of services.

For the sample Exchange implementation outlined above three different health checks will be needed, one each for TCP ports 25, 443, and 993, to be used to check SMTP, HTTPS, and IMAPS respectively.

To create SMTP, HTTPS, and IMAPS Health Checks for Exchange:

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI
2. Click on the **Resources** link within the Server Load Balance tab to expand the resources menu.
3. Click on the **Health Check** link and the Health Check configuration page will appear in the right hand pane.
4. Click on the **add** button at the top left of the Health Check configuration page to bring up the Add Health Check panel.
5. Choose a name for your Health Check and choose TCP as the protocol.
6. Set the port to 25.
7. Set the Interval to 5s, the Timeout to 3s, Down retry to 2, and Up retry to 3.
8. Click the green **Save** button.

Repeat steps 4-8 for each of the 3 healthchecks to be created for ports 25, 443, and 993.

Creating Server Pools

Server pools are groups of servers that can be assigned as a unit to a FortiADC Virtual Server. A server pool consists of member servers and a defined Health Check.

Although the IP of the Virtual Server is used by clients to connect to the Exchange services, the server pools contain members whose IP remains hidden. The port used to connect to the Virtual Server by clients does not need to be the same as the port used by members of the Server Pool associated with that Virtual Server.

In general, a server pool is required for each group of FortiADC servers that offer the same service, or same group of services and are to be load balanced.

For example, when using an Exchange implementation providing the services outlined in the Virtual Server configuration summary, and with 4 Client Access Servers, the following server pools would provide separation and redundancy.

Server Pool	FortiADC Servers
HTTPpool	Exch1, Exch2
Mailpool	Exch3 Exch4
SMTPpool	Exch3 Exch4

Only 3 server pools will be needed in this example, the organization and grouping of servers into server pools depends on the number and distribution of services across your Exchange servers.

To create a server pool for OWA and add members to it, do the following:

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI
2. Click on the **Resources** link within the Server Load Balance tab to expand the resources menu.
3. Click on the **Pool** link to open the Pool Configuration page in the right hand pane.
4. Click the **Add** button at the top right of the Pool Configuration pane to open the Add Pool panel.
5. Enter a name for the pool, in this case **HTTPSpool**.
6. Click the **Health Check** option to enable Health Checks for this Pool.
7. From the **Health Check** section that expands, drag the TCP port 443 Health Check that was created earlier from the right hand side to the left hand side to enable it.
8. Scroll down to the **Members** section of the **Add Pool** panel and click on the **Add** button.
9. Click **OK** on the pop up dialogue to save the Pool.
10. Within the **Members** sub panel, enter the IP address of the server to be

added to the pool as well as the port being used for the service, in this case 443

11. Click on the check box to Inherit Health Checks.
12. Click on the **OK** button.
13. Repeat steps 8-12 for each of the servers to be added to the pool, in this case Exch1 and Exch2.
14. Click the green **Save** button.

To create a server pool for SMTP and add members to it, do the following:

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI
2. Click on the **Resources** link within the Server Load Balance tab to expand the resources menu.
3. Click on the **Pool** link to open the Pool Configuration page in the right hand pane.
4. Click the **Add** button at the top right of the Pool Configuration pane to open the Add Pool panel.
5. Enter a name for the pool, in this case **SMTPpool**.
6. Click the **Health Check** option to enable Health Checks for this Pool.
7. From the **Health Check** section that expands, drag the TCP port 25 Health Check that was created earlier from the right hand side to the left hand side to enable it.
8. Scroll down to the **Members** section of the Add Pool panel and click on the Add button.
9. Click **OK** on the pop up dialogue to save the Pool.
10. Within the **Members** sub panel, enter the IP address of the server to be added to the pool as well as the port being used for the service, in this case 25.
11. Click on the check box to Inherit Health Checks.
12. Click on the **OK** button.
13. Repeat steps 8-12 for each of the servers to be added to the pool, in this case Exch3 and Exch4.
14. Click the green **Save** button.

To create a server pool for IMAPS and add members to it, do the following:

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI
2. Click on the **Resources** link within the Server Load Balance tab to expand the resources menu.
3. Click on the **Pool** link to open the Pool Configuration page in the right hand pane.

4. Click the **Add** button at the top right of the Pool Configuration pane to open the Add Pool panel.
5. Enter a name for the pool, in this case **Mailpool**.
6. Click the **Health Check** option to enable Health Checks for this Pool.
7. From the **Health Check** section that expands, drag the TCP port 993 Health Check that was created earlier from the right hand side to the left hand side to enable it.
8. Scroll down to the **Members** section of the Add Pool panel and click on the Add button.
9. Click **OK** on the pop up dialogue to save the Pool.
10. Within the **Members** sub panel, enter the IP address of the server to be added to the pool as well as the port being used for the service, in this case 993.
11. Click on the check box to Inherit Health Checks.
12. Click on the **OK** button.
13. Repeat steps 8-12 for each of the servers to be added to the pool, in this case Exch3 and Exch4.
14. Click the green **Save** button.

Creating an IP Pool

Since the Exchange servers will respond to clients by sending packets to the source address that they have seen, and since it is usually required that these packets traverse the ADC on their way back to the client, it is generally recommended to use SNAT to substitute addresses owned by the ADC for the actual client IP addresses. In this way, return traffic from the servers is guaranteed to pass via the ADC on its way back to the client.

An IP Pool is an address or group of addresses to be used in conjunction with the SNAT feature. The IP Pool defines the IP addresses used by SNAT when forwarding packets from the ADC to the Exchange servers and thus the addresses seen by the Exchange servers as the sources of client requests.

In this case we will be using the simplest IP Pool possible, consisting of only one address, that of the internal interface connecting to the Exchange servers.

To create an IP Pool.

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI
2. Click on the **Resources** link within the Server Load Balance tab to expand the resources menu.
3. Click on the **IP Pool** link to open the IP Pool Configuration page in the right hand pane.
4. Click on the **Add** button at the top left of the IP Pool Configuration pane.
5. Enter a name for the IP Pool.
6. From the drop down menu, choose the interface through which your Exchange servers will be reached.
7. Click on the green **Save** button.

Configuring OWA, Outlook Anywhere, and ActiveSync

In this Exchange configuration, Outlook Web App (OWA), Outlook Anywhere (OA), ActiveSync (AS), and the RPC MAPI services are all configured to run together on the same server or servers in the CAS array.

This means that FortiADC can be easily configured for OWA, OA, AS,using a single TCP Virtual Server that will provide access for all three services.

1. To create a TCP Virtual Server for port 443.
2. Click on the **Server Load Balance** tab on the left hand pane of the GUI.
3. Click on the **Virtual Servers** link to expand the Virtual Servers menu.
4. Click on the **Virtual Server** link to open the Virtual Server Configuration page in the right hand pane.
5. Click on the **Add** button at the top left of the Virtual Server Configuration panel.
6. Enter a name for the Virtual Server, in this case ExchangeHTTPS.
7. Check the **L4 load balance option** within the Type field.
8. Select **Full NAT** from the Packet Forwarding drop down menu.
9. In the **General section**, enter the IP address of the Virtual Server.
10. Set the Port as 443.
11. Select an Interface for the Virtual Server, usually the External interface.
12. Select **LB_PROF_TCP** from the Profile drop down menu.
13. Select **LB_METHOD_ROUND_ROBIN** from the Method drop down menu
14. Select **HTTPSpool** from the Pool drop down menu.
15. Select the IP Pool created in the IP Pool section from the Source Pool drop down menu.
16. Click the green **Save** button.

Required Name Service Changes

The FQDN used by clients to access the Exchange 2013 services must be changed in DNS (the Domain Name Service) and/or Active Directory to point to the FortiADC HTTPS Virtual Server IP address.

Creating a Virtual Server for IMAPS

To support mailbox access from IMAP clients, you need to start the IMAP service on the Exchange Client Access Servers. For clients to send email through Exchange, you'll also need to configure SMTP. It is highly recommended to use IMAPS for security by enabling it on the Exchange server.

IMAPS is load balanced using a Layer 4 TCP cluster, and persistence is not needed. SSL offloading is not performed by FortiADC for IMAPS.

To create a TCP Virtual Server for IMAPS.

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI.
2. Click on the **Virtual Servers** link to expand the Virtual Servers menu.
3. Click on the **Virtual Server** link to open the Virtual Server Configuration page in the right hand pane.
4. Click on the **Add** button at the top left of the Virtual Server Configuration panel.
5. Enter a name for the Virtual Server, in this case ExchangeIMAPS.
6. Check the **L4 load balance** option within the Type field.
7. Select **Full NAT** from the Packet Forwarding drop down menu.
8. In the **General section**, enter the IP address of the Virtual Server.
9. Set the Port as 993.
10. Select an Interface for the Virtual Server, usually the External interface.
11. Select **LB_PROF_TCP** from the Profile drop down menu.
12. Select **LB_METHOD_ROUND_ROBIN** from the Method drop down menu
13. Select **IMAPSpool** from the Pool drop down menu.
14. Select the IP Pool created in the IP Pool section from the Source Pool drop down menu.
15. Click the green **Save** button.

Required Name Service Changes

The FQDN or IP used by clients to access the IMAPS services must be changed in DNS (the Domain Name Service) and/or Active Directory to point to the FortiADC IMAPS Virtual Server IP address.

Configuring Virtual Server for SMTP

SMTP connections in Exchange 2013 may be configured using either Edge Transport Servers or Hub Transport Servers. Just use the appropriate IP addresses for your configuration when adding servers in Step 6, below. Persistence is not used with the SMTP protocol.

To create a TCP Virtual Server for SMTP.

1. Click on the **Server Load Balance** tab on the left hand pane of the GUI.
2. Click on the **Virtual Servers** link to expand the Virtual Servers menu.
3. Click on the **Virtual Server** link to open the Virtual Server Configuration page in the right hand pane.
4. Click on the **Add** button at the top left of the Virtual Server Configuration panel.
5. Enter a name for the Virtual Server, in this case ExchangeSMTP.
6. Check the **L4 load balance** option within the Type field.
7. Select **Full NAT** from the Packet Forwarding drop down menu.
8. In the **General section**, enter the IP address of the Virtual Server.
9. Set the Port as 25.
10. Select an Interface for the Virtual Server, usually the External interface.
11. Select **LB_PROF_TCP** from the Profile drop down menu.
12. Select **LB_METHOD_LEAST_CONNECTION** from the Method drop down menu
13. Select **SMTPpool** from the Pool drop down menu.
14. Select the IP Pool created in the IP Pool section from the Source Pool drop down menu.
15. Click the green **Save** button.

Required Name Service Changes

The FQDN used by clients and other mail servers to access the Exchange 2013 SMTP services must be changed in DNS (the Domain Name Service) and/or Active Directory to point to the FortiADC SMTP Virtual Server IP address.

Summary

FortiADC provides the load balancing, application acceleration, and high availability features demanded by medium to large Microsoft Exchange Server 2013 configurations. This document has presented a step-by-step guide to configuring FortiADC's features for an Exchange 2013 environment.

About Fortinet's ADC Solutions

From the leader in Network Security comes a new breed of Application Delivery Controller (ADC), FortiADC, built for your needs today and in the future. FortiADC solutions meet the challenge of delivering mission critical applications reliably, securely and at a value that others can't match.

We offer a broad selection of hardware and virtual appliances to cover your needs whether you're a small business looking to expand your website or an enterprise that has to span applications across data centers around the globe.

All FortiADCs offer global server load balancing (GSLB) at no extra cost. If you need to bridge your application across two or more data centers for disaster recovery or to improve response times, the built-in GSLB is easy to setup and manage. For even greater flexibility and more connectivity choices, Fortinet's FortiDirector GSLB provides a subscription-based GSLB service that can bridge traffic between multiple data centers, single servers or to host-based services. With the ability to route traffic based on server health, network status or even time of day, FortiDirector gives you even greater versatility to mnage your applications.

For more information on Fortinet's portfolio of ADC solutions, please visit www.fortinet.com or contact us directly at the numbers listed below for your region.

GLOBAL HEADQUARTERS
Fortinet Inc.
1090 Kifer Road
Sunnyvale, CA 94086
United States
Tel: +1.408.235.7700
Fax: +1.408.235.7737
www.fortinet.com/sales

EMEA SALES OFFICE
120 rue Albert Caquot
06560, Sophia Antipolis,
France
Tel: +33.4.8987.0510
Fax: +33.4.8987.0501

APAC SALES OFFICE
300 Beach Road 20-01
The Concourse
Singapore 199555
Tel: +65.6513.3730
Fax: +65.6223.6784

LATIN AMERICA SALES OFFICE
Prol. Paseo de la Reforma 115 Int. 702
Col. Lomas de Santa Fe,
C.P. 01219
Del. Alvaro Obregón
México D.F.
Tel: 011-52-(55) 5524-8480