Monitoring Exchange 2007 and 2010 Environments

Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Trademarks

Microsoft Windows, Windows NT, Windows 2003, and Windows 2000 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

Table of Contents

Table of Figures

1

Introduction

Exchange 2007/2010 provides a reliable messaging system, with built-in protection against spam and viruses. Using Exchange 2007/2010, users throughout an organization can access e-mail, voice mail, calendars, and contacts from a wide variety of devices and from any location.

As a messaging system that is widely used in both large corporations and small businesses, Exchange Server has always been scalable in both directions. However, new demands on messaging – such as compliance, security, and disaster recovery – have created new challenges for delivering a messaging system that works great in small businesses and large enterprises alike. To rise to these new challenges, the architecture of Exchange server 2007/2010 has been updated to take advantage of 64-bit hardware, simplified administration and routing, and to enable an Exchange server to host one or more server roles.

Exchange server 2007/2010 allows you to assign predefined roles to specific servers. These roles allow organizations to control mail flow, increase security, and distribute services, as shown in the following illustration.

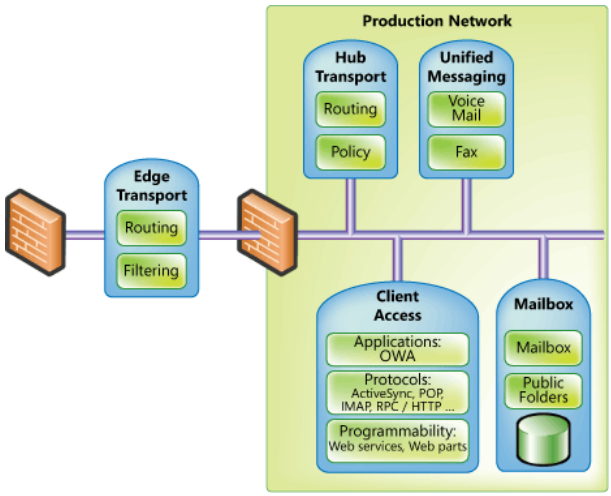

Figure 1.1: Exchange server 2007/2010 server roles

In Exchange server 2007/2010, roles are predefined and chosen during installation. The role selected during installation ensures that only the necessary services and components are installed. Not only does this simplify deployment, but it also enables more efficient management and hardware utilization over time.

a. **Client Access role**. Similar to the front-end server in earlier versions of Exchange, this server proxies Internet client traffic to the correct mailbox server.

b. **Mailbox role**. This role hosts user mailboxes stored in databases that can be replicated or clustered.

c. **Hub Transport role**. This role provides internal routing of all messages – from Edge servers, Unified Messaging (UM) servers, or between two users on the same mailbox database. The Hub Transport role is also where messaging policy is enforced for messages moving within and outside the organization.

d. **Unified Messaging role**. This role enables PBX integration to allow voice mail and fax messages delivered to Exchange mailboxes, and provides voice dial-in capabilities to Exchange Server. This role and its services are explained in more detail later in this paper.

e. **Edge Transport role**. This server resides outside your internal network and provides on-premise e-mail security, antivirus, and anti-spam services for Exchange.

The error-free functioning of each of the aforesaid roles is essential to ensure that Exchange 2007/2010 discharges its duties efficiently and effectively. Performance issues experienced by any of these components can delay or even halt message delivery. To avoid such anomalies, each of the above-mentioned server roles require 24 x 7 monitoring. For this purpose, eG Enterprise prescribes specialized monitoring models that provide top-down monitoring of each of the server roles discussed above. This way, administrators can focus on the performance of every role individually, and can perform indepth analysis of the issues experienced by each role.

In addition to these role-based models, eG Enterprise also provides an *Exchange 2007* model and an *Exchange 2010* model that provide integrated monitoring across server roles. Though the 2007 and 2010 versions of Exchange are the same in terms of functionality and architecture, eG Enterprise offers two separate models for these versions because the Windows services that execute on both these versions are different and need to be tracked separately. These models provide administrators an overview of the health of the entire Exchange 2007 and Exchange 2010 (as the case may be) environments, and enables them to accurately pinpoint that server role which is serving as a road-block to the timely delivery of Exchange services.

1.1 How does eG Enterprise Monitor the Exchange 2007 and Exchange 2010 Servers and their Server Roles

eG Enterprise adopts an agent-based approach to monitoring the *Exchange 2007* and *Exchange 2010* models and also those models that correspond to each of the Exchange 2007/2010 server roles. The agent-based approach requires that you install and configure the eG agent on the Exchange 2007/2010 host (if one of the 'integrated' *Exchange 2007* or *Exchange 2010* models is being used) or on the host on which the server role to be monitored exists.

This internal agent, once started, periodically runs a wide variety of tests on the Exchange 2007/2010 server/server role to extract useful performance data. Some of these tests , namely – the ExchangeMailboxStatus test, the ExchangeStorageGroup test, and the ExchangeQueueStats test – require **Exchange Administrator** privileges to execute. Therefore, prior to monitoring an Exchange 2007/2010 server/server role using eG Enterprise, make sure that you configure the eG agent to run with the privileges of an **Exchange Administrator**.

This document discusses each of the monitoring models that eG Enterprise offers.

2

Monitoring the Client Access Servers

The Client Access server role is required in every Exchange server 2007/2010 organization.

The Client Access server role supports the Microsoft Outlook Web Access and Microsoft Exchange ActiveSync client applications and the Post Office Protocol version 3 (POP3) and Internet Message Access Protocol version 4rev1 (IMAP4) protocols. The Client Access server role also supports services, such as the Autodiscover service and Web services.

The Client Access server role accepts connections to your Exchange 2007/2010 server from a variety of different clients. Software clients such as Microsoft Outlook Express and Eudora use POP3 or IMAP4 connections to communicate with the Exchange server. Hardware clients, such as mobile devices, use ActiveSync, POP3, or IMAP4 to communicate with the Exchange server.

Failure of any of these protocols or issues in network connection between the clients and the server can disrupt client-server communication and cause critical emails to go undelivered. In order to avoid such adversities, the operations of the Client Access server should be continuously monitored, and administrators proactively alerted to abnormalities.

eG Enterprise prescribes a specialized *Microsoft Exchange CAS* model for monitoring Client Access servers.

Figure 2.1: The layer model of the Client Access server

Each layer of Figure 2.1 reports critical performance metrics extracted from the Client access server, that enable administrators to answer the following questions easily and effectively.

- How efficient is the ActiveSync engine? Is it taking too long to process requests? How many requests to ActiveSync are still in queue?
- Were requests to the Availability service processed quickly?
- Were all connection requests to mailboxes serviced by the cache, or were any requests missed?
- What is the current session load on Outlook web access (OWA)?
- Does OWA respond swiftly to user requests?
- Does OWA take too long to complete search requests?
- Did any requests for web access fail?
- Is the Exchange web server responding promptly to requests?

Since the bottom 6 layers of Figure 2.1 have already been discussed in the *Monitoring Unix and Windows Servers* document, the sections to come discuss the **Client Access Services** layer only.

2.1 The Client Access Services Layer

The tests mapped to this layer monitor the critical services offered by the Client Access server, which include:

- Exchange ActiveSync service
- Exchange Availability service
- Outlook Web Access service
- Exchange Web service

In addition, the layer monitors the availability of the Exchange 2007/2010 server to send and receive mails.

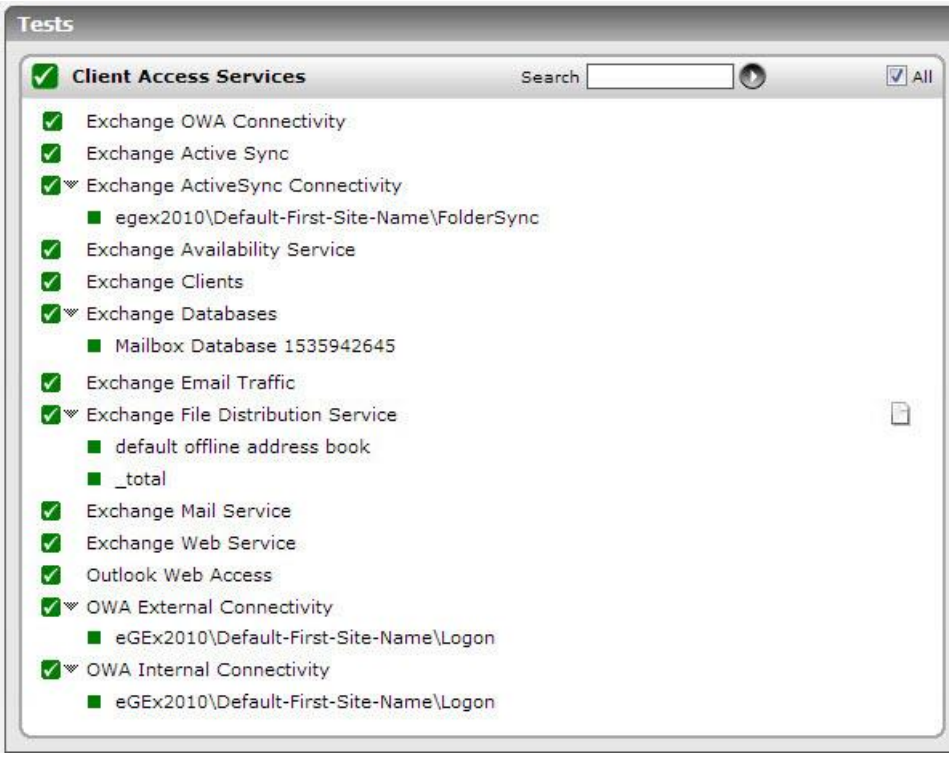

Figure 2.2: The tests associated with the Client Access Services layer

2.1.1 Exchange File Distribution Service Test

An Offline Address Book (OAB) is a container that stores a collection of Offline Address Lists. Typically, users download these address lists to obtain information about other users in their organization.

Exchange server 2007/2010 introduces a new mechanism for distributing Offline Address Books that does not involve Public Folders; it instead uses HTTP(S) and the Background Intelligent Transfer Service (BITS). This new web-based OAB distribution process depends on several components working together:

- Exchange System Attendant Service – this service runs on the mailbox server to create the OAB.
- Exchange File Distribution Service – this service runs on CAS (Client Access) servers and is responsible for obtaining the OAB content from the OABGen server.
- OAB Virtual Directory – This is an IIS virtual directory on a CAS server where the OAB is downloaded from.
- Autodiscover – Autodiscover runs on a CAS server and returns the correct OAB URL for a given client connection.

The OAB is typically generated on a mailbox server by the Exchange System Attendant Service. At configured intervals (default: every 8 hours), the Exchange Fle Distribution Service (FDS) on the CAS server polls the mailbox server for new OAB files. The first poll happens when the Exchange File Distribution Service starts; so, the exact time a server polls will be different on each CAS. If polling reveals new files, the Exchange Fle Distribution Service downloads the files from the mailbox server. The copied files are stored in a web distribution folder on the CAS server. The user then connects to the

AutoDiscover service via Outlook to get the closest OAB distribution URL. Autodiscover returns the URL to the CAS server. Outlook then connects with BITS to the URL provided, and downloads the OAB.

From this, we can infer that the location of the CAS server, the quality of the network link between the CAS and the mailbox server, and the polling interval are key factors that influence the speed, frequency, and overall efficiency of the OAB download performed by the Exchange File Distribution Service. Carelessly made changes to the polling interval and issues with network connectivity can therefore significantly impact the OAB distribution process, thereby delaying users access to the latest information pertaining to other users.

Using the **Exchange File Distribution Service** test, you can periodically monitor the OAB downloads performed by the FDS service on the CAS server and promptly capture slowdowns (if any) in the downloading process and changes in polling interval (if any).

Purpose	Periodically monitors the OAB downloads performed by the FDS service on the CAS server and promptly captures slowdowns (if any) in the downloading process and changes in polling interval (if any)		
Target of the test	A server configured with the Client Access server role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** - The port number of the client access server. By default, this is 110.		
Outputs of the test	One set of results for the Client Access server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td rowspan="2">test</td>
<td>Download task queued:
Indicates whether any download tasks have been queued or not.</td>
<td></td>
<td>If one/more download tasks are in queue, this measure will report the value Yes. If no download tasks are in queue, then this measure will report the value No.

The numeric values that correspond to these measure values are as follows:

<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Yes</td><td>1</td></tr>
<tr><td>No</td><td>0</td></tr>
</table>

The value Yes for this measure could indicate that the CAS server is taking too much time to download the OAB files, thereby causing subsequent download tasks to be queued. The prolonged downloads could be due to a poor network link between the CAS server and the OAB Generation server. Slowdowns can also occur if a large number of OAB files are downloaded, or if the size of the OAB files is huge.

Note:

By default, this measure will report the Measure Values in the table above to indicate if any download task is pending. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</td>
</tr>
<tr>
<td>Download task completed:
Indicates the number of OAB downloads that have been completed since the last measurement period.</td>
<td>Number</td>
<td>Ideally, the value of this measure should be high. A very low value could indicate downloading bottlenecks that might require further investigation.

Another reason for a change in the value of this measure is a change in the polling interval. While an increase in the polling frequency, can increase the value of this measure, a decrease in the polling frequency can cause a less number of download tasks to be completed and can hence, reduce the value of this measure.

A reduction in the polling interval can cause critical updates to OAB files to be available to end users only after a long time.</td>
</tr>
</table>

2.1.2 Exchange Mail Service Test

This test monitors the availability and performance of a Microsoft Exchange 2007/2010 mail server from an external perspective. The test mimics a mail client activity by using the Exchange Web Service for sending and receiving mails.

For this test to execute smoothly, the external agent executing the test should be in the same domain as the Exchange 2007/2010 server.

Purpose	Monitors the availability and performance of a Microsoft Exchange 2007/2010 mail server from an external perspective		
Target of the test	A server configured with the Client Access server role		
Agent deploying the test	An external agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** - The port number of the client access server. By default, this is 110. 4. **FROM USER NAME** - Provide the name of the **Exchange Administrator**, using whose credentials the eG agent executing this test is running. Alternatively, you can provide the email ID of the **Exchange Administrator** as well. 5. **FROM USER PASSWORD** – Specify the password of the **Exchange Administrator**. 6. **CONFIRM PASSWORD** – Confirm the password by retyping it here. 7. **EXCHANGE DOMAIN NAME** - Provide a valid domain in which the target server is running. 8. **WEB SERVICE URL** – To enable the test to connect to Exchange Web Services, you need to provide the **External Web Service URL** here. To know what URL to provide, run the following command from the Exchange server's powershell command prompt: **Get-WebServicesVirtualDirectory -server <servername> | select name, *url* | fl** For instance, if your Exchange server's name is **Exchange**, then your command will be: **Get-WebServicesVirtualDirectory -server Exchange | select name, *url* | fl** Upon successful execution of the command, a list of URLs will be displayed. Note down the URL displayed against the label **ExternalUrl** , and enter it against **WEB SERVICE URL**.		
Outputs of the test	One set of results for the Client Access server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Send mail availability:** Indicates the availability of the mail server for receiving the mails sent by the test.	Percent	A value of 0 indicates that the test was not successful in sending a mail. Possible reasons for this could include the mail server being down, the network connection to the server not being available, or the test configuration information being incorrect.
	Sent messages: Indicates the number of messages sent to the mail server.	Number	A value of –1 indicates that the mail server may be down or the configuration information may be incorrect.
	Avg time to send messages: Indicates time taken to send a mail to the mail server.	Secs	A high value of this measure could indicate high network traffic or that the mail server is busy.
	Receive mail availability: Indicates the availability of the exchange server for sending mails to the mail client.	Percent	The value of 0 indicates that the test was not successful in receiving a mail message from the Exchange server. Possible reasons could be incorrect configuration information.
	Received messages: Indicates the number of messages received by the mail client from the mail server.	Number	The value of 0 indicates that the test was not successful in receiving mail messages from the Exchange server. The possible reasons could be: The sent messages could be in the message queue of the mail server but not routed to the mail box Configuration information may be incorrect Network failure The mail service may not be running in the user account
	Mail received time: Indicates the time taken by the mail client to receive a mail from the mail server.	Secs	A high value in this measure indicates that the mail server is busy or the network traffic is high.

	Avg roundtrip time: The average of the round trip time (the time lapse between transmission and reception of a message by the server) of all the messages received by the mail server during the last measurement period.	Mins	This is a key measure of quality of the mail service. An increase in roundtrip time may be indicative of a problem with the mail service. Possible reasons could include queuing failures, disk space being full, etc.
	Max roundtrip time: The high water mark of the round trip time (the time lapse between transmission and reception of a message by the server) of all messages received by the mail server during the last measurement period.	Mins	If the value of the **Received messages** measure is 1, then the value of this measure will be the same as the **Avg roundtrip time** measure.

2.1.3 Exchange Active Sync Test

By default, when you install the Client Access server role on a computer that is running Microsoft Exchange server 2007/2010, you enable Microsoft Exchange ActiveSync. Exchange ActiveSync lets you synchronize a mobile device with your Exchange 2007/2010 mailbox.

Exchange ActiveSync is an Microsoft Exchange synchronization protocol (HTTP and XML) that is optimized to work together with high-latency and low-bandwidth networks. Exchange ActiveSync enables mobile device users to access their e-mail, calendar, contacts, and tasks and to continue to be able to access this information while they are working offline.

The performance of Microsoft Exchange ActiveSync is affected by many factors. These include the number of users who are synchronizing with Exchange ActiveSync, the types of mobile devices that are synchronizing with it, and how much data each user synchronizes between the Microsoft Exchange server and the mobile device. By using monitoring, you can understand the factors that affect the performance of Exchange ActiveSync.

This test measures the health of the ActiveSync engine.

Purpose	Measures the health of the ActiveSync engine
Target of the test	A server configured with the Client Access server role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** - The port number of the client access server. By default, this is 110.

Outputs of the test	One set of results for the Client Access server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	ActiveSync request processing time: Indicates the average time elapsed waiting for a request to complete.	Secs	This measure includes Ping Request Time, which can increase the general response time. Adding ping counters helps clarify where performance is being impacted.
	Ping commands pending on the server: Indicates the number of ping commands that are currently pending on the server.	Number	
	Ping commands dropped: Indicates the number of Ping commands per second whose connection to the client was dropped before a response could be issued.	Dropped/sec	
	ActiveSync requests to the server: Indicates the number of HTTP requests that are received from the client via ASP.NET per second.	Reqs/Sec	
	ActiveSync requests queued for processing: Indicates the number of HTTP requests that are currently waiting to be assigned to a thread.	Number	A steady increase in this value over time is a cause for concern, as it is indicative of a processing bottleneck.
	Sync commands processed: Indicates the number of sync commands that are currently processed by the server.	Number	
	Worker threads busy: Indicates the number of worker threads that are presently busy processing requests.	Number	

	Worker threads idle: Indicates the number of worker threads that are currently idle.	Number	Ideally, this value should be low.

2.1.4 Exchange Availability Service Test

The Microsoft Exchange server 2007/2010 Availability service improves information workers' calendaring and meeting scheduling experience by providing secure, consistent, and up-to-date free and busy information to computers running Microsoft Office Outlook 2007. Outlook 2007 uses the Autodiscover service to obtain the URL of the Availability service. The Autodiscover service is similar to the Domain Name System (DNS) Web service for Exchange 2007/2010 Web services. Essentially, the Autodiscover service helps Outlook 2007 locate various Web services, such as the Unified Messaging (UM), Offline Address Book (OAB), and Availability services.

The following figure illustrates the process flow for the Availability service.

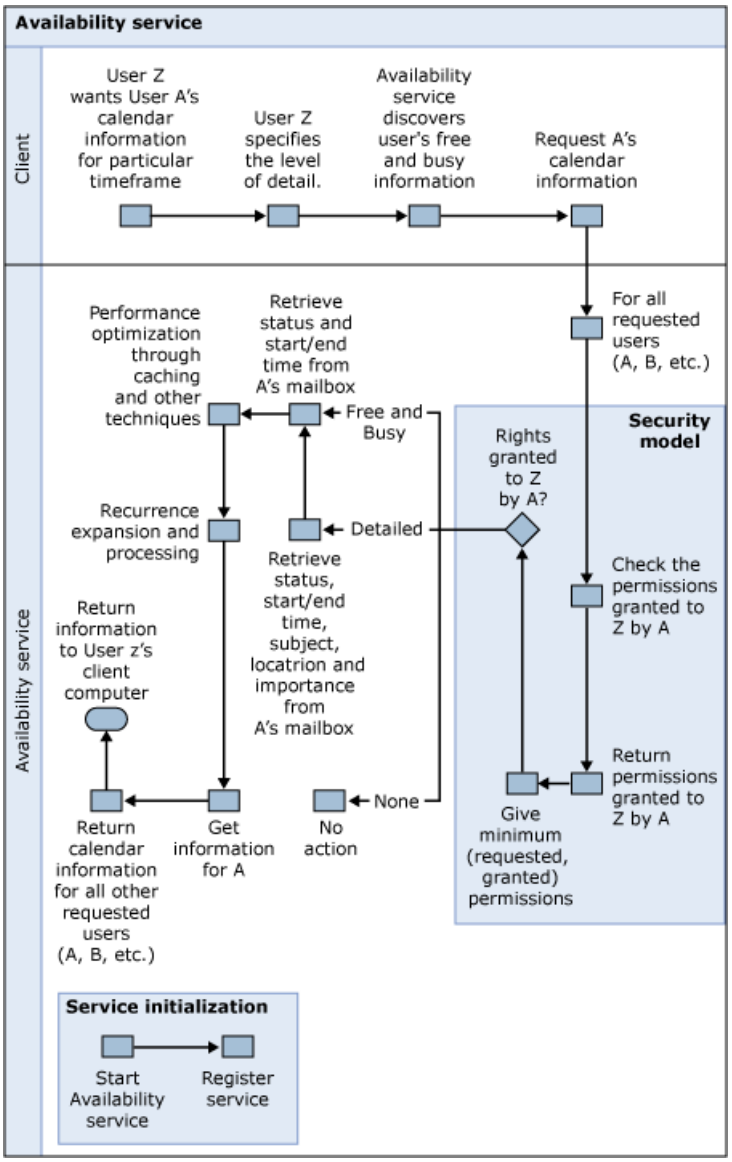

Figure 2.3: The Availability Service

This test reports statistics indicating how healthy the Availability Service is.

Purpose	Reports statistics indicating how healthy the Availability Service is
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** - The port number of the client access server. By default, this is 110.

Outputs of the test	One set of results for the Mailbox server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Availability requests to the web service: Indicates the number of requests serviced per second.	Reqs/Sec	The request can be only for free busy or include suggestions. One request may contain multiple mailboxes. A very low request servicing rate could indicate a processing bottleneck.
	Avg. mailboxes processed per request: Indicates the average number of mailboxes processed per request.	Mailboxes/Req	
	Mailbox connection hits: Indicates the number of mailboxes opened per second without creating a new connection.	Hits/Sec	Ideally, this value should be high. A low cache hit rate could increase processing overheads.
	Mailbox connection misses: Indicates the number of mailboxes opened per second, by creating a new connection, because there is no available connection in the cache.	Misses/Sec	Ideally, this measure should be low. A very high cache miss rate could indicate insufficient connections in the cache to service requests. You might want to consider resizing the cache.

2.1.5 Outlook Web Access Test

Outlook Web Access (OWA) is a HyperText Transfer Protocol (HTTP) virtual server that enables users to access their Microsoft Exchange inbox using a Web browser. In the event that users are unable to access their mailbox, you can use the metrics reported by this test to determine whether the problem in HTTP access is local to the client access server or not.

Purpose	Helps determine whether the problem in HTTP access is local to the client access server or not
Target of the test	A server configured with the Client Access role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** - The port number of the client access server. By default, this is 110.		
Outputs of the test	One set of results for the Client Access server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Outlook web access sessions: Indicates the number of Outlook Web Access user sessions that are created per second.	Login/Sec	
	Current unique web access users: Indicates the number of unique users currently logged on to Outlook Web Access.	Number	This value monitors the number of unique active user sessions, so that users are only removed from this count after they log off or their session times out.
	Response time for web access - average: Indicates the average time in seconds that elapsed between the beginning and end of an OEH or ASPX request.	Secs	This is a good measure of the latency that a client is experiencing. Higher values may indicate high user load or higher than normal CPU time.
	Search time during web access - average: Indicates the average time that elapsed while waiting for a search to complete.	Secs	Ideally, this value should be low at all times.
	Request rate for web access: Indicates the number of requests handled by Outlook Web Access per second.	Reqs/Sec	
	Failed requests for web access: Indicates the number of Outlook Web Access requests that failed, per second.	Reqs/Sec	A zero value is typically desired. If the measure reports a non-zero value, the reason for the same should be investigated.

	Store logon failures for web access: Indicates the percentage of Outlook Web Access user logons to Microsoft Exchange Mailbox servers that have failed currently.	Percent	Ideally, this value should be low.

2.1.6 Exchange Web Service Test

This test monitors requests from Exchange clients and reveals how quickly the Exchange 2007/2010 server responds to these requests.

Purpose	Monitors requests from Exchange clients and reveals how quickly the Exchange 2007/2010 server responds to these requests		
Target of the test	An Exchange 2007/2010 server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110.		
Outputs of the test	One set of results for the Exchange 2007/2010 server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Request rate to Exchange web server: Indicates the number of requests from clients that are processed each second.	Reqs/Sec	
	Avg. response time for web service: Indicates the average time (in milliseconds) that has elapsed between the beginning and end of requests.	Msecs	A high value for this measure is indicative of a slowdown in the responsiveness of the server.

2.1.7 Exchange RPC HTTP Test

This test assists you in assessing the load and issues with the RPC/HTTP Proxy component.

Purpose	Assesses the load and issues with the RPC/HTTP Proxy component		
Target of the test	An Exchange 2007/2010 server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110.		
Outputs of the test	One set of results for the Exchange 2007/2010 server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Current incoming RPCs: Indicates the number of front-end HTTP connections.	Number	This measure serves as a good indicator of the load imposed by the user.
	Current unique users: Indicates the number of unique users currently connected to a back-end server via RPC/HTTP.	Number	This measure serves as a good indicator of level of user load.
	RPC/HTTP requests: Indicates the rate of RPC/HTTP request send to the back-end server	Reqs/Sec	This measure indicates the current Outlook Anywhere load.
	Failed backend connections: Indicates the rate at which the RPC proxy attempts are occurring but fail to establish a connection to a back-end server.	Conns/Sec	Ideally, this value should be 0.

2.1.8 Exchange OWA Connectivity Test

This test verifies whether the Microsoft Office Outlook web app is running as expected. This test can be used to test Outlook Web App connectivity for all Microsoft Exchange Server 2010 virtual directories on a specified Client Access server for all mailboxes on servers running Exchange that are in the same Active Directory site.

This test is also used to test the connectivity for an individual Exchange Outlook Web App URL. To execute this test, you need to setup a test account in the exchange forest and you need to run the script which is available in the following location *\scripts\ new-TestCasConnectivityUser.ps1* .

Purpose	Verifies whether the Microsoft Office Outlook web app is running as expected		
Target of the test	An Exchange 2007/2010 server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **CLIENT ACCESS SERVER** – Specify the complete qualified hostname of the client access server in the **CLIENT ACCESS SERVER** text box. 6. **OUTLOOK WEB APP URL** – Specify the website URL in the **OUTLOOK WEB APP URL**. By Default *none* will be provided.		
So thOutputs of the test	One set of results for the Exchange 2007/2010 server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Logon latency: Indicates the delay time in logging in to Outlook Web App.	Msecs	
	OWA connectivity status: Indicates the status of the OWA connection.	MB	If the value of this measure reports to 100, it implies that the OWA connection is successfull, otherwise this measure reports to 0.

2.1.9 Exchange ActiveSync Connectivity Test

Exchange ActiveSync lets you synchronize a mobile device with your Exchange 2010 mailbox, so that you can check your emails from your mobile phone itself! Whenever a mobile phone user complaints that he/she is unable to check or is experiencing slowness when checking emails on his/her mobile phone, Exchange administrators need to quickly determine what is causing the non-sync – is it because ActiveSync is unable to synchronize with the user's mailbox? Or is it because ActiveSync is taking too long to perform the synchronization? At which stage of the synchronization did the failure/delay occur? This test helps answer all these questions. The test periodically checks ActiveSync connectivity at every stage (a.k.a scenario) of the synchronization – eg., the Logon stage, the FolderSync stage, the Options stage, etc. - reports issues and latencies (if any) in connectivity, and leads you to the exact stage at which the failure/slowdown occurred.

Purpose	Periodically checks ActiveSync connectivity at every stage of the synchronization – eg., the Logon stage, the FolderSync stage, the Options stage, etc. - reports issues and latencies (if any) in connectivity, and leads you to the exact stage at which the failure/slowdown occurred		
Target of the test	An Exchange 2010 server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **CLIENT ACCESS SERVER** – Specify the fully-qualified domain name of the Client Access server.		
Outputs of the test	One set of results for each *<ClientAccessServer>/<LocalSiteNameofClientAccessServer>/<SynchronizationStage/Scenario tested>* combination		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr><td>test</td><td>ActiveSync connectivity status:
Indicates whether the ActiveSync connectivity check was successful or not at this stage/scenario of the synchronization.</td><td></td><td>If the value of this measure is Success, it indicates that the ActiveSync connectivity check was successful at this stage. If the value of this measure is Failure, it indicates that mailbox synchronization using ActiveSync failed at this stage. The numeric values that correspond to these measure values are as follows:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Success</td><td>1</td></tr><tr><td>Failure</td><td>0</td></tr></table>
Note:
Typically, this measure reports the Measure Values listed in the table above to indicate the ActiveSync connectivity status. However, in the graph of this measure, the Numeric values are used to represent the connectivity status.</td></tr>
<tr><td></td><td>ActiveSync latency:
Indicates the time taken by ActiveSync to successfully complete this stage/scenario of the synchronization.</td><td>Secs</td><td>A low value is desired for this measure. A high value indicates that this stage/scenario of the synchronization is taking too long to complete.
Compare the value of this measure across stages/scenarios to know where the maximum delay occurred. This will greatly aid troubleshooting.</td></tr>
</table>

2.1.10 OWA Internal Connectivity Test

Outlook Web App (OWA) is a browser-based email client accessible from the web. It allows you to check your email from computers that do not have an email client (such as Outlook 2010) installed. If an internal user (i.e., intranet user) complains that he/she is unable to check emails using OWA, you can run this test to figure out whether/not OWA is accessible, and if so, how long it takes to connect to OWA over the intranet. This test attempts to connect to the OWA URL from the intranet, and for every stage (a.k.a scenario) of the connection process, reports whether/not that stage completed successfully or not and the time taken for completion. This way, the test not only reports an OWA connectivity failure/slowdown, it also points you to the exact stage at which the failure/slowdown may have occurred. This brings connectivity issues in the internal network and their probable causes to light.

Purpose	Attempts to connect to the OWA URL from the intranet, and for every stage (a.k.a scenario) of the connection process, reports whether/not that stage completed successfully or not and the time taken for completion
Target of the test	An Exchange 2010 server

<table>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Client Access server.
3. PORT – The port number of the client access server. By default, this is 110.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each <ClientAccessServer>/<LocalSiteNameofClientAccessServer>/<SynchronizationStage/Scenario tested> combination</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>OWA internal connectivity status:
Indicates whether the OWA connectivity check was successful or not at this stage/scenario of the connection.</td><td>Percent</td><td>The value 0 for this measure indicates that the connectivity check failed at this stage of the interaction. The value 100 on the other hand indicates that this stage of the interaction was cleared successfully.
Use the detailed diagnosis of this measure to know the OWA URL that the test tried to connect to.</td></tr>
<tr><td>Logon internal latency:
Indicates the time taken for the successful completion of this stage/scenario.</td><td>Secs</td><td>A low value is desired for this measure. A high value indicates a connection bottleneck. Compare the value of this measure across descriptors to isolate the exact stage/scenario that took the maximum time to complete, and investigate further to determine why.</td></tr>
</table>

2.1.11 OWA External Connectivity Test

Outlook Web App (OWA) is a browser-based email client accessible from the web. It allows you to check your email from computers that do not have an email client (such as Outlook 2010) installed. If an external user (i.e., internet user) complains that he/she is unable to check emails using OWA, you can run this test to figure out whether/not that OWA is accessible over the internet, and if so, how long that connection takes. This test attempts to connect to the OWA URL from the internet, and for every stage (a.k.a scenario) of the connection process, reports whether/not that stage completed successfully or not and the time taken for completion. This way, the test not only reports an OWA connectivity failure/slowdown, it also points you to the exact stage at which the failure/slowdown may have occurred. This brings connectivity issues in the internet and their probable causes to light.

Purpose	Attempts to connect to the OWA URL from the internet, and for every stage (a.k.a scenario) of the connection process, reports whether/not that stage completed successfully or not and the time taken for completion
Target of the test	An Exchange 2010 server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.
Outputs of the test	One set of results for each *<ClientAccessServer>/<LocalSiteNameofClientAccessServer>/<SynchronizationStage/Scenario tested>* combination

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	OWA external connectivity status: Indicates whether the OWA connectivity check was successful or not at this stage/scenario of the connection.	Percent	The value 0 for this measure indicates that the connectivity check failed at this stage of the interaction. The value 100 on the other hand indicates that this stage of the interaction was cleared successfully. Use the detailed diagnosis of this measure to know the OWA URL that the test tried to connect to.
	Logon external latency: Indicates the time taken for the successful completion of this stage/scenario.	Secs	A low value is desired for this measure. A high value indicates a connection bottleneck. Compare the value of this measure across descriptors to isolate the exact stage/scenario that took the maximum time to complete, and investigate further to determine why.

Use the detailed diagnosis of the *OWA internal connectivity status* measure to know the OWA URL that the test tried to connect to. Sometimes, an incorrect URL may also report incorrect results. To avoid this, its best to check the URL of the OWA using the detailed diagnosis of the *OWA internal connectivity status* measure.

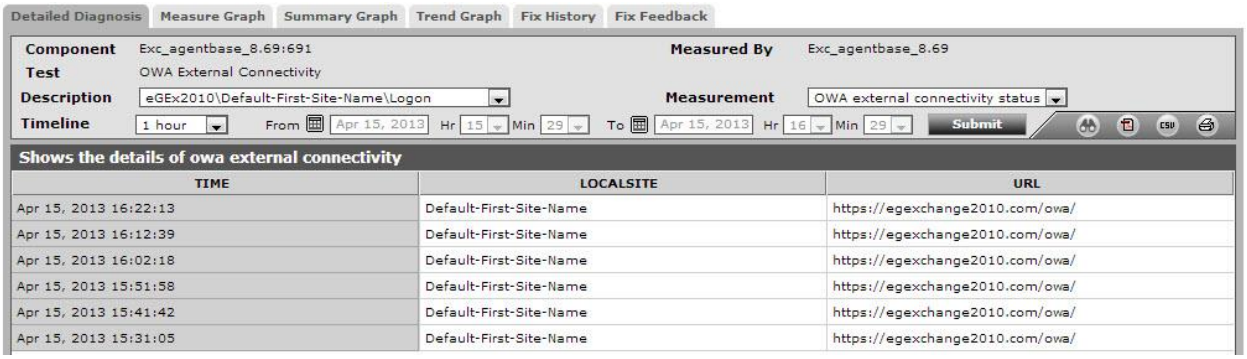

Figure 2.4: The detailed diagnosis of the OWA external connectivity status test

2.1.12 ActiveSync Device Status

Administrators must constantly track the devices connecting to ActiveSync, so that they can proactively identify devices that are unable to sync with user mailboxes, the users using these devices, and the probable reason for the non-sync, much before device users even notice that something is wrong! Likewise, administrators should also be able to zero-in on devices that are connected to ActiveSync, but have been inactive for long time periods, so that they can take efforts to clear out such devices en masse. To isolate such devices, administrators can use the **ActiveSync Device Status** test. This test reports the count of devices that are using ActiveSync without a glitch, those that are having problems using Activesync, and the stale (inactive) devices. Using the detailed diagnosis of this measure, the devices that are operating well, those that are not, and those that are stale can be clearly isolated.

Purpose	Reports the count of devices that are using ActiveSync without a glitch, those that are having problems using Activesync, and the stale (inactive) devices		
Target of the test	A server configured with the Client Access Server role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **INACTIVE DEVICE AGE IN DAYS** – Specify the minimum duration (in days) for which a device should not have synchronized with its mailbox for it to be counted as a stale/inactive device. 6. **SHOW DD FOR OK STATUS DEVICE** – By default, this flag is set to **No**, indicating that detailed metrics will not be available by default for the *Device with OK status* measure reported by this test. To ensure that this test collects detailed metrics for this measure, set this flag to **Yes**. 7. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for the Client access server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Device with OK status:** Indicates the count of devices that are currently able to synchronize with their mailboxes via ActiveSync.	Number	**Detailed diagnostics will be available for this measure only if the SHOW DD FOR OK STATUS DEVICE flag is set to 'Yes'.** If available, then, you can use the detailed diagnosis of this measure to know which devices are able to connect to ActiveSync and synchronize with their mailboxes, and which users are using these devices.
	Device with not OK status: Indicates the number of devices that are currently unable to synchronize with their Exchange mailboxes via ActiveSync.	Number	Ideally, the value of this measure should be 0. A non-zero value for this measure implies that one/more devices are unable to synchronize with their Exchange mailboxes. To know these devices and their users, use the detailed diagnosis of this measure.
	Stale devices: Indicates the current number of stale devices.	Number	Use the detailed diagnosis of this measure to know which devices are inactive, which users are using such devices, and how long these devices have remained inactive.

2.1.13 Exchange ActiveSync Servers Test

Where Exchange ActiveSync is used to synchronize mobile devices with Exchange server mailboxes, Exchange administrators may want to know which devices are connecting to the server at any given point in time, so that accesses by unauthorized devices can be instantly detected and blocked. Administrators may also want to track the usage of mailboxes by mobile devices over time and identify the most and the least effective users, so that access policies can be accordingly drawn. Moreover, when a device user complains of a slowdown when accessing his/her mailbox, administrators may want to take a look at the network traffic generated by every device that is connecting to the server at the time of the slowdown, so that devices that are choking the bandwidth and causing the slowness can be accurately isolated. The **Exchange ActiveSync Servers** test performs all these checks periodically and provides Exchange administrators with actionable information that will enable them to take well-informed

and intelligent performance/policy decisions.

This test auto-discovers the devices that are synchronizing with the Exchange 2010 mailboxes via ActiveSync, and for each device, reports the number of hits/accesses made by that device and the amount of data transmitted and received by that device. In the process, the test points administrators to the following:

- Devices that are currently connected to the Exchange server; unauthorized devices can thus be quickly captured;
- Devices that are accessing the Exchange server mailboxes frequently and those that seldom use the mailboxes; sizing and policy decisions can be taken based on this observation
- Devices that are consuming excessive bandwidth resources and could hence be contributing to the sluggish quality of the network;

<table>
<tr><td>Purpose</td><td colspan="3">Auto-discovers the devices that are synchronizing with the Exchange 2010 mailboxes via ActiveSync, and for each device, reports the number of hits/accesses made by that device and the amount of data transmitted or received by that device</td></tr>
<tr><td>Target of the test</td><td colspan="3">A server configured with the Mailbox role</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Client Access server.
3. PORT – The port number of the client access server. By default, this is 110.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell v2, which enables you to administer every part of Microsoft Exchange Server 2007/2010. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. To enable the test to load the Exchange management shell snap-in (exshell.psc1) for script execution, you need to specify the full path to the Exchange management shell in the XCHGEXTENSIONSHELLPATH text box. For instance, your specification can be, c:|progra~1|micros~1|exchan~1|v14|bin|exshell.psc1.
5. LOGFILE NAME – The Client Access Server is an IIS web server that hosts Exchange-related web pages. This is why, like any other IIS web server, the client access server creates a daily log of its activities – including Exchange ActiveSync-related activities - in the C:|inetpub|logs|logfiles|W3SVC1| directory by default. To report metrics on ActiveSync, this test parses the client access server's log file, reads the ActiveSync-related errors/warnings/general information messages that were recently logged (i.e., during the last 5 minutes) from the file, and writes them to a ActiveSynchLog.log file it creates in the <EG_AGENT_INSTALL_DIR>\agent\logs directory. Then, the test reads the metrics of interest from this log file and reports them to the eG manager. To enable the test to do the above, you need to specify the exact path to the directory that contains the client access server's logs in the LOGFILENAME text box.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each device or IP address that is currently accessing the mailboxes on the Exchange server</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

	Total hits: Indicates the current number of hits/accesses to the Exchange mailbox server from this device.	Number	Comparing the value of this measure across devices will help you to identify the device that is constantly accessing the Exchange mailbox server and that which is not using the server as frequently. Based on these usage metrics, administrators can define access policies. Also, this measure serves as a good indicator of the level of device activity on the Exchange server; based on this knowledge, administrators can right-size their Exchange infrastructure – i.e., decide on how much CPU, memory, bandwidth, and disk resources the Exchange server has to be allocated so that it can handle the ActiveSync load.
	Data sent: Indicates the amount of data this device is currently sending to the Exchange mail server.	KB	Compare the value of these measures across the devices to identify the device that is currently generating the maximum amount of network traffic when interacting with its mailbox on the Exchange server. In the event of a slowdown, this comparative analysis will point administrators to that device which is engaged in bandwidth-intensive conversations with the Exchange server, thus causing accesses to slow down. During normal operations on the other hand, administrators can analyze these measures over time to gauge the average network throughput of ActiveSync activities; this can help them decide whether/not more network resources need to be allocated to handle ActiveSync load efficiently.
	Data received: Indicates the amount of data currently received by this device from the Exchange mail server.	KB	
	Average unique devices: Indicates the number of unique devices currently accessing the ActiveSync server.		

2.1.14 Exchange ActiveSync Requests Status Test

When a mobile device attempts to synchronize with a mailbox on the Exchange server, the server returns an HTTP status code to the device indicating the status of the synchronization attempt. Some of the most critical HTTP status codes for ActiveSync and their interpretations are as follows:

HTTP status code	Description
HTTP_200	Indicates that the device successfully connected to the Exchange server and synchronized with the mailbox on the server.

HTTP_401	Indicates one or all of the following: • The credentials provided to access the server are incorrect; • The user is not enabled for synchronization
HTTP_404	Indicates that an issue exists with the user account
HTTP_404	Indicates that the file requested is not found on the server
HTTP_449	Indicates that the synchronization attempt should be retried
HTTP_500	Indicates one or all of the following: • The Internet Information Service is unavailable. • Windows Integrated Authentication is not enabled on the Exchange Server virtual directory of the server where the mailbox of the user resides. • Synchronization is tried when the mailbox is being moved.
HTTP_502	Indicates an error in the proxy server used to connect to the ActiveSync Server
HTTP_503	Indicates that the ActiveSync service is unavailable

Periodic review of these status codes and the synchronization attempts that resulted in these codes is imperative to understand how error-prone ActiveSync on the Exchange server is, identify the errors that occur frequently, investigate why these errors occur, and easily troubleshoot them. This is where the **Exchange ActiveSync Requests Status** test helps!

This test automatically discovers the HTTP status codes returned by the Exchange server for ActiveSync accesses. For each status code so discovered, the test reports the number and percentage of accesses that returned that status code. This way, the test points administrators to status codes that were returned most often, thus shedding light on ActiveSync errors that occurred frequently.

Purpose	Automatically discovers the HTTP status codes returned by the Exchange server for ActiveSync accesses. For each status code so discovered, the test reports the number and percentage of accesses that returned that status code
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Client Access server.
3. PORT – The port number of the client access server. By default, this is 110.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell v2, which enables you to administer every part of Microsoft Exchange Server 2007/2010. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. To enable the test to load the Exchange management shell snap-in (exshell.psc1) for script execution, you need to specify the full path to the Exchange management shell in the XCHGEXTENSIONSHELLPATH text box. For instance, your specification can be, c:|progra~1|micros~1|exchan~1|v14|bin|exshell.psc1.
5. LOGFILE NAME – The Client Access Server is an IIS web server that hosts Exchange-related web pages. This is why, like any other IIS web server, the client access server creates a daily log of its activities – including Exchange ActiveSync-related activities - in the C:|inetpub|logs|logfiles|W3SVC1| directory by default. To report metrics on ActiveSync, this test parses the client access server's log file, reads the ActiveSync-related errors/warnings/general information messages that were recently logged (i.e., during the last 5 minutes) from the file, and writes them to a ActiveSynchLog.log file it creates in the <EG_AGENT_INSTALL_DIR>\agent\logs directory. Then, the test reads the metrics of interest from this log file and reports them to the eG manager. To enable the test to do the above, you need to specify the exact path to the directory that contains the client access server's logs in the LOGFILENAME text box.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each status code returned by the ActiveSync server</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total hits:
Indicates the current number of hits to the Exchange mailbox server that returned this status code.</td><td>Number</td><td rowspan="2">Compare the value of these measures across status codes to identify the status code that is returned frequently. High values for the 4xx or 5xx class of status codes is a cause for concern, as they indicate client and server errors respectively. If such status codes are returned often, administrators will have to look up the Microsoft documentation to understand what error condition each code represents and how to resolve it.</td></tr>
<tr><td>Hit ratio:
Indicates the percentage of hits to the Exchange mailbox server that returned this status code.</td><td>Percent</td></tr>
</table>

2.1.15 Exchange ActiveSync Devices Test

In environments where ActiveSync is enabled, it is normal for users wielding different types of devices to synchronize their mailbox with their device. In such environments, administrators should pay close attention to the device types that are connected to the Exchange server mailboxes at any given point in time, so that unsupported device types

can be detected and the users using such types of devices identified and advised accordingly. It is also essential that administrators study how frequently each of these device types are accessing the Exchange server and monitor the level of activity generated by these device types on the server and on the network. If a device users complains of delays in accessing his/her mailbox, then this visibility will enable administrators to identify those device types to which the slowdown can be attributed. In addition, administrators will also need to know from time-to-time how much load ActiveSync imposes on the Exchange server and the network, across all device types! This aggregated measure will enable administrators to figure out whether/not the Exchange server is sized right to handle the load. To receive such in-depth insights into ActiveSync performance – both at the per-device type level and across all device types – administrators can use the **Exchange ActiveSync Devices** test.

This test auto-discovers the device types currently synchronizing with the Exchange server. For each device type, the test reports the number of ActiveSync accesses made by that device type and the number and size of items transmitted and received by that device type. This way, the test leads administrators to those device types that are utilizing the available network and server resources excessively, thus degrading the experience of some or all device users. Detailed metrics provided by the test also help administrators identify all the users who are using devices of a particular type and pinpoint the exact user who is engaged in a resource-intensive interaction with the Exchange server mailbox. Additionally, the test reports metrics across all device types, thus enabling administrators to measure the current load on the server and the network and assess the ability of the server to handle that load.

Purpose	Auto-discovers the device types currently synchronizing with the Exchange server. For each device type, the test reports the number of ActiveSync accesses made by that device type and the number and size of items transmitted and received by that device type
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Client Access server.
3. PORT – The port number of the client access server. By default, this is 110.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell v2, which enables you to administer every part of Microsoft Exchange Server 2007/2010. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. To enable the test to load the Exchange management shell snap-in (exshell.psc1) for script execution, you need to specify the full path to the Exchange management shell in the XCHGEXTENSIONSHELLPATH text box. For instance, your specification can be, c:|progra~1|micros~1|exchan~1|v14|bin|exshell.psc1.
5. LOGFILE NAME – The Client Access Server is an IIS web server that hosts Exchange-related web pages. This is why, like any other IIS web server, the client access server creates a daily log of its activities – including Exchange ActiveSync-related activities - in the C:|inetpub|logs|logfiles|W3SVC1| directory by default. To report metrics on ActiveSync, this test parses the client access server's log file, reads the ActiveSync-related errors/warnings/general information messages that were recently logged (i.e., during the last 5 minutes) from the file, and writes them to a ActiveSynchLog.log file it creates in the <EG_AGENT_INSTALL_DIR>\agent\logs directory. Then, the test reads the metrics of interest from this log file and reports them to the eG manager. To enable the test to do the above, you need to specify the exact path to the directory that contains the client access server's logs in the LOGFILENAME text box.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each device type that is accessing the Exchange server; an additional All descriptor is also supported, which reports a set of aggregated metrics across all device types</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total hits:
Indicates the number of hits/accesses made by this device type to the Exchange server mailbox.</td><td>Number</td><td>Comparing the value of this measure across device types will help administrators identify that device type which is very actively synchronizing with the Exchange mailbox.
Using the detailed diagnosis of this measure, administrators can also identify the precise user who is making the maximum number of accesses, which device that user is using, and the details of that device.
Based on this information, access policies can be defined.
Also, by observing the variations in the value of this measure for the All descriptor, administrators can effectively gauge the typical level of activity on the Exchange server and figure out if the server is sized right to handle this load.</td></tr>
</table>

	Total items sent: Indicates the number of items currently sent from this device type to the Exchange server.	Number	These measures indicate how much network traffic and I/O load is generated by each of the device types. By comparing the value of these measures across device types, administrators can easily and accurately identify that device type that is engaged in resource-intensive communication with the Exchange server. In the event of a slowdown, the results of this comparative analysis will lead administrators to that device type that could be contributing to the slowdown. Once the device type is identified, you can use the detailed diagnosis of the *Total hits* measure to know which user of that device type is actually choking the network/server and what device he/she is currently using.
	Total items received: Indicates the number of items currently received by this device type from the Exchange server.	Number	
	Data sent: Indicates the amount of data currently sent from this device type to the Exchange server.	KB	
	Data received: Indicates the amount of data currently received by this device type from the Exchange server.	KB	

2.1.16 Exchange ActiveSync Policy Compliance Test

Exchange ActiveSync mailbox policies let you apply a common set of policy or security settings to a user or group of users. With the help of these policies, Exchange administrators can indicate what specific devices – thus users – connecting to ActiveSyc, can do.

EAS policies are applied to users; each user can have zero policies or one EAS policy at any given time. If you don't explicitly assign a policy to a user, the default policy is applied instead. During the initial sync of a new device (that is, one that has not been synchronized to the server before), the device and server exchange what EAS calls a policy key. Think of the policy key as a GUID or MAC address; it's a unique key that indicates one specific policy. If the device and server keys do not match, the device is required to request the most recent policy and then apply it. The process of applying a policy to the device is known as provisioning. On most devices, the user will see a dialog box indicating that the server is applying a policy and asking whether to accept it. If the user declines the policy, the server might or might not allow the device to continue to sync to it; the exact behavior depends on whether the default policy on the server allows non-provisioned devices.

Not every device that connects to ActiveSync will implement every setting defined in a policy; some devices may even lie about the policy settings that they implement. Hence, the onus of determining the number of devices that comply with the policy settings and to what extent is the compliance, lies with the administrator. To determine this, administrators can use the **Exchange ActiveSync Policy Compliance** test. This test reports the count and percentage of devices connecting to ActiveSync that are fully compliant, partially compliant, and completely non-compliant with their mailbox policies. This way, the test reveals the degree of compliance to configured policies.

Purpose	Reports the count and percentage of devices connecting to ActiveSync that are fully compliant, partially compliant, and completely non-compliant with their mailbox policies. This way, the test reveals the degree of compliance to configured policies.
Target of the	A server configured with the Client Access Server role

<table>
<tr><td>test</td><td colspan="3"></td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Client Access server.
3. PORT – The port number of the client access server. By default, this is 110.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell v2, which enables you to administer every part of Microsoft Exchange Server 2007/2010. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. To enable the test to load the Exchange management shell snap-in (exshell.psc1) for script execution, you need to specify the full path to the Exchange management shell in the XCHGEXTENSIONSHELLPATH text box. For instance, your specification can be, c:|progra~1|micros~1|exchan~1|v14|bin|exshell.psc1.
5. LOGFILE NAME – The Client Access Server is an IIS web server that hosts Exchange-related web pages. This is why, like any other IIS web server, the client access server creates a daily log of its activities – including Exchange ActiveSync-related activities - in the C:|inetpub|logs|logfiles|W3SVC1| directory by default. To report metrics on ActiveSync, this test parses the client access server's log file, reads the ActiveSync-related errors/warnings/general information messages that were recently logged (i.e., during the last 5 minutes) from the file, and writes them to a ActiveSynchLog.log file it creates in the <EG_AGENT_INSTALL_DIR>\agent\logs directory. Then, the test reads the metrics of interest from this log file and reports them to the eG manager. To enable the test to do the above, you need to specify the exact path to the directory that contains the client access server's logs in the LOGFILENAME text box.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of compliance – Compliant, Partially compliant, Not compliant, Unknown</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total hits:
Indicates the number of devices currently accessing ActiveSync that are of this compliance type .</td><td>Number</td><td>Compare the value of this measure across compliance types to know how compliant the maximum number of devices are - fully compliant? partially compliant? non-compliant? or unknown? (i.e., the compliance level cannot be determined)</td></tr>
<tr><td>Hits ratio:
Indicates the percentage of devices currently accessing ActiveSync that are of this compliance type.</td><td>Percent</td><td>Compare the value of this measure across compliance types to know the degree of compliance of devices accessing ActiveSync - fully compliant? partially compliant? or non-compliant? or unknown? (i.e., the compliance level cannot be determined)</td></tr>
</table>

2.1.17 Exchange ActiveSync User Agent Test

Devices communicating with Exchange via ActiveSync identify themselves to Exchange using a 'User Agent' string and a 'User Agent Type' string. For instance, an iPhone may identify itself to Exchange using the user agent string 'Apple-iPhone/xxx.xxx' and the user agent type 'iPhone'. While the user agent string is unique for every device, multiple devices can be of the same user agent type. By tracking the types of user agents that are accessing Exchange via ActiveSync, administrators can determine which type of devices are attempting to synchronize with the Exchange mailboxes. In times of an overload, this information may point administrators to the exact type of devices that could be contributing to the heavy load. To obtain this useful information, administrators can use the **Exchange ActiveSync User Agent** test. For every user agent type, this test reports the total number of user agents of that type that are accessing ActiveSync at any given point in time. In addition, it also reports the number of unique devices of each type synchronizing with Exchange. This not only indicates the current synchronization load on Exchange, but also helps identify the user agent types (i.e., device types) that could be contributing to the workload.

Purpose	For every user agent type, this test reports the total number of user agents of that type that are accessing ActiveSync at any given point in time. In addition, it also reports the number of unique devices of each type synchronizing with Exchange.
Target of the test	A server configured with the Client Access Server role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Client Access server. 3. **PORT** – The port number of the client access server. By default, this is 110. 4. **XCHGEXTENSIONSHELLPATH** - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell v2, which enables you to administer every part of Microsoft Exchange Server 2007/2010. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. To enable the test to load the Exchange management shell snap-in (exshell.psc1) for script execution, you need to specify the full path to the Exchange management shell in the XCHGEXTENSIONSHELLPATH text box. For instance, your specification can be, *c:\progra~1\micros~1\exchan~1\v14\bin\exshell.psc1*. 5. **LOGFILE NAME** – The Client Access Server is an IIS web server that hosts Exchange-related web pages. This is why, like any other IIS web server, the client access server creates a daily log of its activities – including Exchange ActiveSync-related activities - in the *C:\inetpub\logs\logfiles\W3SVC1* directory by default. To report metrics on ActiveSync, this test parses the client access server's log file, reads the ActiveSync-related errors/warnings/general information messages that were recently logged (i.e., during the last 5 minutes) from the file, and writes them to a **ActiveSynchLog.log** file it creates in the **<EG_AGENT_INSTALL_DIR>\agent\logs** directory. Then, the test reads the metrics of interest from this log file and reports them to the eG manager. To enable the test to do the above, you need to specify the exact path to the directory that contains the client access server's logs in the **LOGFILENAME** text box.
Outputs of the test	One set of results for each user agent type

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Total hits: Indicates the current number of synchronization requests from user agents of this type.	Number	This is a good indicator of the current synchronization load on the Exchange server. You can compare the value of this measure across user agents to know which type of user agents are actually overloading the server.
	Unique devices: Indicates the number of unique devices of this user agent type that are currently accessing ActiveSync.	Number	Compare the value of this measure across user agent types to identify the device type that is significantly impacting the server workload.

2.1.18 Exchange ActiveSync Device Errors Test

In order to enable administrators to quickly troubleshoot current issues with ActiveSync, the eG Exchange Monitor intelligently reads ActiveSync-related errors/warnings/general information or status messages related to ActiveSync commands captured recently (i.e., in the last 5 minutes) from the client access server's log file and writes them to the **ActiveSynchLog.log** file it creates in the **<EG_AGENT_INSTALL_DIR>\agent\logs** directory. At specified intervals, this test scans the **ActiveSynchLog.log** file for configured patterns of errors and reports the number and nature of such errors (if found).

At least one of the following tests should be running and reporting metrics for this test to work:

- Exchange ActiveSync Servers Test
- Exchange ActiveSync Status Test
- Exchange ActiveSync Users Test
- Exchange ActiveSync Policy Compliance Test
- Exchange ActiveSync User Agents Test

Purpose	Scans the **ActiveSynchLog.log** file for configured patterns of errors and reports the number and nature of such errors (if found)
Target of the test	An Exchange Client Access Server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured. 3. **PORT** – The port at which the server listens 4. **ALERTFILE –** By default, the full path to the **ActiveSynchLog.log** file is set here. 5. **SEARCHPATTERN** - Enter the specific patterns of alerts to be monitored. The pattern should be in the following format: *<PatternName>:<Pattern>*, where *<PatternName>* is the pattern name that will be displayed in the monitor interface and *<Pattern>* is an expression of the form - *expr* or expr or *expr or expr*, etc. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. For example, say you specify *ItemNotFound:Item** in the **SEARCHPATTERN** text box. This indicates that "ItemNotFound" is the pattern name to be displayed in the monitor interface. "Item*" indicates that the test will monitor only those lines in the alert log which start with the term "Item". Similarly, if your pattern specification reads: *UserDisabledForSync:*Sync*, then it means that the pattern name is *UserDisabledForSync* and that the test will monitor those lines in the alert log which end with the term *Sync*. A single pattern may also be of the form e1+e2, where + signifies an OR condition. That is, the <*PatternName*> is matched if either e1 is true or e2 is true. Multiple search patterns can be specified as a comma-separated list. For example: *ItemNotFound:Item*, UserDisabledForSync:*Sync* If you want all the messages in a log file to be monitored, then your specification would be: *<PatternName>:**. 6. **LINES** - Specify two numbers in the format x:y. This means that when a line in the alert file matches a particular pattern, then x lines before the matched line and y lines after the matched line will be reported in the detail diagnosis output (in addition to the matched line). The default value here is 0:0. Multiple entries can be provided as a comma-separated list. If you give 1:1 as the value for **LINES**, then this value will be applied to all the patterns specified in the **SEARCHPATTERN** field. If you give 0:0,1:1 as the value for **LINES** and if the corresponding value in the **SEARCHPATTERN** filed is like *ItemNotFound:Item*, UserDisabledForSync:*Sync*, then: 0:0 will be applied to *ItemNotFound* pattern 1:1 will be applied to *UserDisabledForSync* pattern

7. **EXCLUDEPATTERN** - Provide a comma-separated list of patterns to be excluded from monitoring in the **EXCLUDEPATTERN** text box. For example **critical*, *exception**. By default, this parameter is set to 'none'.

8. **UNIQUEMATCH** - By default, the **UNIQUEMATCH** parameter is set to **FALSE**, indicating that, by default, the test checks every line in the log file for the existence of each of the configured **SEARCHPATTERNS**. By setting this parameter to **TRUE**, you can instruct the test to ignore a line and move to the next as soon as a match for one of the configured patterns is found in that line. For example, assume that *Pattern1:*fatal*,Pattern2:*error** is the **SEARCHPATTERN** that has been configured. If **UNIQUEMATCH** is set to **FALSE**, then the test will read every line in the log file completely to check for the existence of messages embedding the strings 'fatal' and 'error'. If both the patterns are detected in the same line, then the number of matches will be incremented by 2. On the other hand, if **UNIQUEMATCH** is set to **TRUE**, then the test will read a line only until a match for one of the configured patterns is found and not both. This means that even if the strings 'fatal' and 'error' follow one another in the same line, the test will consider only the first match and not the next. The match count in this case will therefore be incremented by only 1.

9. **ROTATINGFILE** - This flag governs the display of descriptors for this test in the eG monitoring console.

   If this flag is set to **true** and the **ALERTFILE** text box contains the full path to a specific (log/text) file, then, the descriptors of this test will be displayed in the following format: *Directory_containing_monitored_file:<SearchPattern>*. For instance, if the **ALERTFILE** parameter is set to *c:|eGurkha|logs|syslog.txt*, and **ROTATINGFILE** is set to **true**, then, your descriptor will be of the following format: *c:|eGurkha|logs:<SearchPattern>*. On the other hand, if the **ROTATINGFILE** flag had been set to **false**, then the descriptors will be of the following format: *<FileName>:<SearchPattern>* - i.e., *syslog.txt:<SearchPattern>* in the case of the example above.

   If this flag is set to **true** and the **ALERTFILE** parameter is set to the directory containing log files, then, the descriptors of this test will be displayed in the format: *Configured_directory_path:<SearchPattern>.* For instance, if the **ALERTFILE** parameter is set to *c:|eGurkha|logs*, and **ROTATINGFILE** is set to **true**, then, your descriptor will be: *c:|eGurkha|logs:<SearchPattern>*. On the other hand, if the **ROTATINGFILE** parameter had been set to **false**, then the descriptors will be of the following format: *Configured_directory:<SearchPattern>* - i.e., *logs:<SearchPattern>* in the case of the example above.

   If this flag is set to true and the **ALERTFILE** parameter is set to a specific file pattern, then, the descriptors of this test will be of the following format: *<FilePattern>:<SearchPattern>*. For instance, if the **ALERTFILE** parameter is set to *c:|eGurkha|logs|*sys**, and **ROTATINGFILE** is set to **true**, then, your descriptor will be: **sys*:<SearchPattern>*. In this case, the descriptor format will not change even if the **ROTATINGFILE** flag status is changed .

10. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.

11. **CASESENSITIVE** - This flag is set to **No** by default. This indicates that the test functions in a 'case-insensitive' manner by default. This implies that, by default, the test ignores the case of your **ALERTFILE** and **SEARCHPATTERN** specifications. If this flag is set to **Yes** on the other hand, then the test will function in a 'case-sensitive' manner. In this case therefore, for the test to work, even the case of your **ALERTFILE** and **SEARCHPATTERN** specifications should match with the actuals.

12. **ROLLOVERFILE** - By default, this flag is set to **false**. Set this flag to **true** if you want the test to support the 'roll over' capability of the specified **ALERTFILE**. A roll over typically occurs when the timestamp of a file changes or when the log file size crosses a pre-determined threshold. When a log file rolls over, the errors/warnings that pre-exist in that file will be automatically copied to a new file, and all errors/warnings that are captured subsequently will be logged in the original/old file. For instance, say, errors and warnings were originally logged to a file named *error_log*. When a roll over occurs, the content of the file *error_log* will be copied to a file named *error_log.1*, and all new errors/warnings will be logged in *error_log*. In such a scenario, since the **ROLLOVERFILE** flag is set to **false** by default, the test by default scans only *error_log.1* for new log entries and ignores *error_log*. On the other hand, if the flag is set to **true**, then the test will scan both *error_log* and *error_log.1* for new entries.

    If you want this test to support the 'roll over' capability described above, the following conditions need to be fulfilled:

    - The **ALERTFILE** parameter has to be configured only with the name and/or path of one/more alert files. File patterns or directory specifications should not be specified in the **ALERTFILE** text box.
    - The roll over file name should be of the format: "<**ALERTFILE**>.1", and this file must be in the same directory as the **ALERTFILE**.

13. **OVERWRITTENFILE** - By default, this flag is set to **false**. Set this flag to **true** if log files do not 'roll over' in your environment, but get overwritten instead. In such environments typically, new error/warning messages that are captured will be written into the log file that pre-exists and will replace the original contents of that log file; unlike when 'roll over' is enabled, no new log files are created for new entries in this case. If the **OVERWRITTENFILE** flag is set to **true**, then the test will scan the new entries in the log file for matching patterns. However, if the flag is set to **false**, then the test will ignore the new entries.

14. **USEUTF8** – Set this flag to **Yes**, if the test needs to use the UTF-8 encoding format for reading from the specified alert file.

15. **USEUTF16** - Set this flag to **Yes**, if the test needs to use the UTF-16 encoding format for reading from the specified alert file.

16. **ENCODEFORMAT** – By default, this is set to *none*, indicating that no encoding format applies by default. However, if the test has to use a specific encoding format for reading from the specified **ALERTFILE** , then you will have to provide a valid encoding format here - eg., *UTF-8, UTF-16*, etc. Where multiple log files are being monitored, you will have to provide a comma-separated list of encoding formats – one each for every log file monitored. Make sure that your encoding format specification follows the same sequence as your **ALERTFILE** specification. In other words, the first encoding format should apply to the first alert file, and so on. For instance, say that your alertfile specification is as follows: *D:|logs|report.log,E:|logs|error.log, C:|logs|warn_log*. Assume that while *UTF-8* needs to be used for reading from *report.log* , *UTF-16* is to be used for reading from *warn_log* . No encoding format need be applied to *error*.log. In this case, your **ENCODEFORMAT** specification will be: *UTF-8,none,UTF-16*.

If your **ALERTFILE** specification consists of file patterns that include wildcard characters (eg., */tmp/db/*dblogs*,/tmp/app/*applogs**), then such configurations will only be supported in the ANSI format, and not the UTF format.

17. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.

18. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

    The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

    - The eG manager license should allow the detailed diagnosis capability
    - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Outputs of the test	One set of results for every **SEARCHPATTERN** configured		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Recent errors: Indicates the number of errors that were added to the **ActiveSynchLog.log** file when the test was last executed.	Number	The value of this measure is a clear indicator of the number of "new" errors detected in ActiveSync. The detailed diagnosis of this measure, if enabled, provides the detailed descriptions of the errors of the configured patterns.

2.1.19 Exchange ActiveSync Device Commands Test

In order to enable administrators to quickly troubleshoot current issues with ActiveSync, the eG Exchange Monitor intelligently reads ActiveSync-related errors/warnings/general information or status messages related to ActiveSync commands captured recently (i.e., in the last 5 minutes) from the client access server's log file and writes them to the **ActiveSynchLog.log** file it creates in the **<EG_AGENT_INSTALL_DIR>\agent\logs** directory. At specified intervals, this test scans the **ActiveSynchLog.log** file for configured patterns of command-related messages and reports the number and nature of messages (if found) matching the configured patterns.

At least one of the following tests should be running and reporting metrics for this test to work:

- Exchange ActiveSync Servers Test
- Exchange ActiveSync Status Test
- Exchange ActiveSync Users Test
- Exchange ActiveSync Policy Compliance Test
- Exchange ActiveSync User Agents Test

Purpose	Scans the **ActiveSynchLog.log** file for configured patterns of command-related messages and reports the number and nature of messages (if found) matching the configured patterns.
Target of the test	An Exchange Client Access Server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured. 3. **PORT** – The port at which the server listens 4. **ALERTFILE –** By default, the full path to the **ActiveSynchLog.log** file is set here. 5. **SEARCHPATTERN** - Enter the specific patterns of alerts to be monitored. The pattern should be in the following format: *<PatternName>:<Pattern>*, where *<PatternName>* is the pattern name that will be displayed in the monitor interface and *<Pattern>* is an expression of the form - *expr* or expr or *expr or expr*, etc. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. For example, say you specify *Sync:Sync** in the **SEARCHPATTERN** text box. This indicates that "Sync" is the pattern name to be displayed in the monitor interface. "Sync*" indicates that the test will monitor only those lines in the alert log which start with the term "ORA-". Similarly, if your pattern specification reads: *SendMail:*SendMail*, then it means that the pattern name is *SendMail* and that the test will monitor those lines in the alert log which end with the term *SendMail*. A single pattern may also be of the form e1+e2, where + signifies an OR condition. That is, the <*PatternName*> is matched if either e1 is true or e2 is true. Multiple search patterns can be specified as a comma-separated list. For example: *Sync:Sync*, SendMail:*SendMail* If you want all the messages in a log file to be monitored, then your specification would be: *<PatternName>:**. 6. **LINES** - Specify two numbers in the format x:y. This means that when a line in the alert file matches a particular pattern, then x lines before the matched line and y lines after the matched line will be reported in the detail diagnosis output (in addition to the matched line). The default value here is 0:0. Multiple entries can be provided as a comma-separated list. If you give 1:1 as the value for **LINES**, then this value will be applied to all the patterns specified in the **SEARCHPATTERN** field. If you give 0:0,1:1 as the value for **LINES** and if the corresponding value in the **SEARCHPATTERN** filed is like *Sync:Sync*, SendMail:*SendMail*, then: 0:0 will be applied to *Sync* pattern 1:1 will be applied to *SendMail* pattern

7. **EXCLUDEPATTERN** - Provide a comma-separated list of patterns to be excluded from monitoring in the **EXCLUDEPATTERN** text box. For example **critical*, *exception**. By default, this parameter is set to 'none'.

8. **UNIQUEMATCH** - By default, the **UNIQUEMATCH** parameter is set to **FALSE**, indicating that, by default, the test checks every line in the log file for the existence of each of the configured **SEARCHPATTERNS**. By setting this parameter to **TRUE**, you can instruct the test to ignore a line and move to the next as soon as a match for one of the configured patterns is found in that line. For example, assume that *Pattern1:*fatal*,Pattern2:*error** is the **SEARCHPATTERN** that has been configured. If **UNIQUEMATCH** is set to **FALSE**, then the test will read every line in the log file completely to check for the existence of messages embedding the strings 'fatal' and 'error'. If both the patterns are detected in the same line, then the number of matches will be incremented by 2. On the other hand, if **UNIQUEMATCH** is set to **TRUE**, then the test will read a line only until a match for one of the configured patterns is found and not both. This means that even if the strings 'fatal' and 'error' follow one another in the same line, the test will consider only the first match and not the next. The match count in this case will therefore be incremented by only 1.

9. **ROTATINGFILE** - This flag governs the display of descriptors for this test in the eG monitoring console.

   If this flag is set to **true** and the **ALERTFILE** text box contains the full path to a specific (log/text) file, then, the descriptors of this test will be displayed in the following format: *Directory_containing_monitored_file:<SearchPattern>*. For instance, if the **ALERTFILE** parameter is set to *c:|eGurkha|logs|syslog.txt*, and **ROTATINGFILE** is set to **true**, then, your descriptor will be of the following format: *c:|eGurkha|logs:<SearchPattern>*. On the other hand, if the **ROTATINGFILE** flag had been set to **false**, then the descriptors will be of the following format: *<FileName>:<SearchPattern>* - i.e., *syslog.txt:<SearchPattern>* in the case of the example above.

   If this flag is set to **true** and the **ALERTFILE** parameter is set to the directory containing log files, then, the descriptors of this test will be displayed in the format: *Configured_directory_path:<SearchPattern>.* For instance, if the **ALERTFILE** parameter is set to *c:|eGurkha|logs*, and **ROTATINGFILE** is set to **true**, then, your descriptor will be: *c:|eGurkha|logs:<SearchPattern>*. On the other hand, if the **ROTATINGFILE** parameter had been set to **false**, then the descriptors will be of the following format: *Configured_directory:<SearchPattern>* - i.e., *logs:<SearchPattern>* in the case of the example above.

   If this flag is set to true and the **ALERTFILE** parameter is set to a specific file pattern, then, the descriptors of this test will be of the following format: *<FilePattern>:<SearchPattern>*. For instance, if the **ALERTFILE** parameter is set to *c:|eGurkha|logs|*sys**, and **ROTATINGFILE** is set to **true**, then, your descriptor will be: **sys*:<SearchPattern>*. In this case, the descriptor format will not change even if the **ROTATINGFILE** flag status is changed .**DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.

10. **CASESENSITIVE** - This flag is set to **No** by default. This indicates that the test functions in a 'case-insensitive' manner by default. This implies that, by default, the test ignores the case of your **ALERTFILE** and **SEARCHPATTERN** specifications. If this flag is set to **Yes** on the other hand, then the test will function in a 'case-sensitive' manner. In this case therefore, for the test to work, even the case of your **ALERTFILE** and **SEARCHPATTERN** specifications should match with the actuals.

11. **ROLLOVERFILE** - By default, this flag is set to **false**. Set this flag to **true** if you want the test to support the 'roll over' capability of the specified **ALERTFILE**. A roll over typically occurs when the timestamp of a file changes or when the log file size crosses a pre-determined threshold. When a log file rolls over, the errors/warnings that pre-exist in that file will be automatically copied to a new file, and all errors/warnings that are captured subsequently will be logged in the original/old file. For instance, say, errors and warnings were originally logged to a file named *error_log*. When a roll over occurs, the content of the file *error_log* will be copied to a file named *error_log.1*, and all new errors/warnings will be logged in *error_log*. In such a scenario, since the **ROLLOVERFILE** flag is set to **false** by default, the test by default scans only *error_log.1* for new log entries and ignores *error_log*. On the other hand, if the flag is set to **true**, then the test will scan both *error_log* and *error_log.1* for new entries.

    If you want this test to support the 'roll over' capability described above, the following conditions need to be fulfilled:

    - The **ALERTFILE** parameter has to be configured only with the name and/or path of one/more alert files. File patterns or directory specifications should not be specified in the **ALERTFILE** text box.
    - The roll over file name should be of the format: "<**ALERTFILE**>.1", and this file must be in the same directory as the **ALERTFILE**.

12. **OVERWRITTENFILE** - By default, this flag is set to **false**. Set this flag to **true** if log files do not 'roll over' in your environment, but get overwritten instead. In such environments typically, new error/warning messages that are captured will be written into the log file that pre-exists and will replace the original contents of that log file; unlike when 'roll over' is enabled, no new log files are created for new entries in this case. If the **OVERWRITTENFILE** flag is set to **true**, then the test will scan the new entries in the log file for matching patterns. However, if the flag is set to **false**, then the test will ignore the new entries.

13. **USEUTF8** – Set this flag to **Yes**, if the test needs to use the UTF-8 encoding format for reading from the specified alert file.

14. **USEUTF16** - Set this flag to **Yes**, if the test needs to use the UTF-16 encoding format for reading from the specified alert file.

15. **ENCODEFORMAT** – By default, this is set to *none*, indicating that no encoding format applies by default. However, if the test has to use a specific encoding format for reading from the specified **ALERTFILE** , then you will have to provide a valid encoding format here - eg., *UTF-8, UTF-16*, etc. Where multiple log files are being monitored, you will have to provide a comma-separated list of encoding formats – one each for every log file monitored. Make sure that your encoding format specification follows the same sequence as your **ALERTFILE** specification. In other words, the first encoding format should apply to the first alert file, and so on. For instance, say that your alertfile specification is as follows: *D:|logs|report.log,E:|logs|error.log, C:|logs|warn_log*. Assume that while *UTF-8* needs to be used for reading from *report.log* , *UTF-16* is to be used for reading from *warn_log* . No encoding format need be applied to *error*.log. In this case, your **ENCODEFORMAT** specification will be: *UTF-8,none,UTF-16*.

If your **ALERTFILE** specification consists of file patterns that include wildcard characters (eg., */tmp/db/*dblogs*,/tmp/app/*applogs**), then such configurations will only be supported in the ANSI format, and not the UTF format.

16. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.

17. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

    The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

    - The eG manager license should allow the detailed diagnosis capability
    - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Outputs of the test	One set of results for every **SEARCHPATTERN** configured		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	New messages: Indicates the number of messages that were added to the **ActiveSynchLog.log** file when the test was last executed.	Number	The detailed diagnosis of this measure, if enabled, provides the detailed descriptions of the messages of the configured patterns.

3

Monitoring the Mailbox Servers

The Mailbox server role hosts mailbox databases, which contain user's mailboxes. If you plan to host user mailboxes, public folders, or both, the Mailbox server role is required.

In Exchange server 2007/2010, the Mailbox server role integrates with the Active Directory directory service better than the mailbocx features and functionality in earlier versions of Exchange. This improved integration makes deployment and operation tasks much easier. The Mailbox server role also improves the information worker experience by providing richer calendaring functionality, resource management, and offline address book downloads.

The Mailbox server must interact directly with the following:

- Active Directory directory service server
- Hub Transport server
- Client Access server
- Unified Messaging (UM) server
- Microsoft Outlook clients

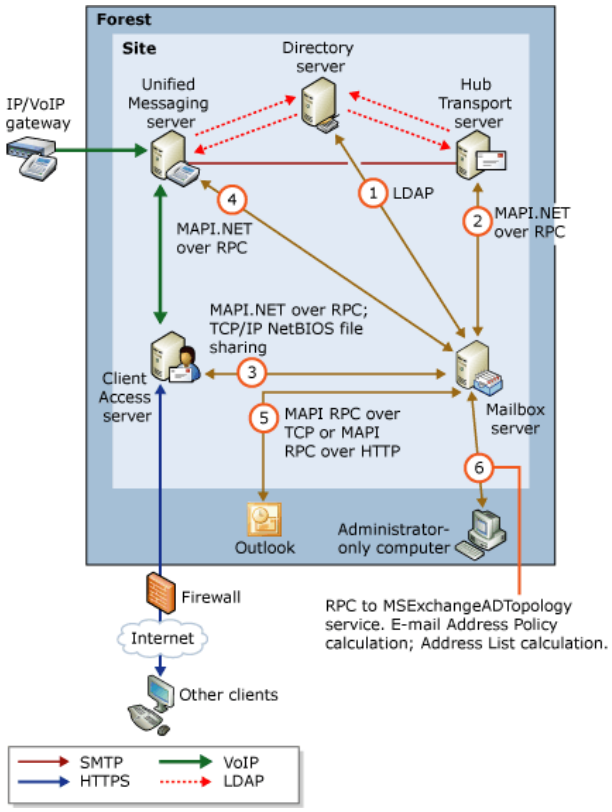

Figure 3.1: The relationship between the Mailbox server and the other server roles, clients, and the Active Directory server

Figure 3.1 shows what protocol the Mailbox server uses to communicate with each of these roles or computers. Each numbered interaction in Figure 3.1 corresponds to the following list, describing what types of information is shared between these roles and computers.

1. The Mailbox server accesses recipient, server, and organization configuration information from Active Directory.
2. The Store driver on the Hub Transport server places messages from the transport pipeline into the appropriate mailbox. The Store driver on the Hub Transport server also adds messages from the Outbox of a sender on the Mailbox server to the transport pipeline.
3. The Client Access server sends requests from clients to the Mailbox server, and returns data from the Mailbox server to the clients. The Client Access server also accesses offline address book files on the Mailbox server through NetBIOS file sharing. The types of data that the Client Access server sends between the client and the Mailbox server are messages, free/busy data, client profile settings, and offline address book data.
4. The Unified Messaging server retrieves e-mail and voice mail messages and calendar information from the Mailbox server for Outlook Voice Access. The Unified Messaging server also retrieves storage quota information from the Mailbox server.
5. Outlook clients that are inside your firewall can access a Mailbox server directly to send and retrieve messages. Outlook clients outside the firewall can access a Mailbox server using remote procedure call (RPC) over Hypertext Transfer Protocol (HTTP).
6. The administrator-only computer retrieves Active Directory topology information from the Microsoft Exchange Active Directory Topology service. It also retrieves e-mail address policy information and address list information.

We can thus conclude that the Mailbox server plays a crucial role in ensuring the uninterrupted flow of mails in an

Exchange organization. Therefore, it is quiet evident that performance issues that the Mailbox server experiences can stall the delivery of emails indefinitely. If such adversities are to be avoided, the Mailbox server has to be monitored 24 x 7 for anomalies, and issues reported should be resolved quickly.

eG Enterprise provides a specialized *Microsoft Exchange Mailbox* model to monitor the Mailbox server inside-out and proactively report potential issues in its performance.

Figure 3.2: Layer model of the Microsoft Exchange Mailbox Server

Every layer of Figure 3.2 above is mapped to a wide variety of tests that run periodic checks on the health of the key components and services offered by the Mailbox server. Using the results of these health-checks, administrators can find answers to the following performance queries:

- Is the Active Directory cache adequately sized to handle requests from the mailbox server?
- Did any search requests to any domain controller fail owing to a bad network link or the non-availability of the domain controller?
- Were any LDAP fatal errors experienced while communicating with a domain controller?
- Did too many bind calls to any domain controller fail?
- Is any domain controller responding too slowly to read and search requests?
- How effectively was the database buffer pool used in serving database requests?
- Is the Exchange store experiencing a processing bottleneck? Which component of the Exchange store is responsible for this – the mailbox store or the public folder store? Are the send and receive queues of the Exchange store too long?
- How soon does the Exchange store deliver mails to local recipients?
- Is the Exchange store Virtual Memory used optimally?
- Has the mailbox of any user exhausted or is about to exhaust its storage quota? Are there too many deleted mails in these mailboxes? Would clearing the deleted mails free space in these mailboxes?
- How efficient is the Exchange Search engine? Have too many mailboxes been left to crawl on the database?
- Does Exchange Search take too long to connect to the Exchange store? Does it index documents

quickly?

- Are too many events awaiting processing by the mailbox assistants?

The sections to come discuss the top 3 layers of the Mailbox server role, as all other layers have already been dealt with in the *Monitoring Unix and Windows Servers* document.

3.1 The Exchange Directory Access Layer

The Mailbox server talks to Active Directory for recipient, server, and organization configuration information. The tests mapped to the Exchange Directory Access layer monitor these interactions to bring problems to light.

Figure 3.3: The tests mapped to the Exchange Directory Access test

3.1.1 ActiveDirectory Access Cache Test

This test reveals whether the AD cache is being utilized effectively or not.

Purpose	Reveals whether the AD cache is being utilized effectively or not		
Target of the test	A server configured with the Mailbox server role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Mailbox server 3. **PORT** - The port number of the Mailbox server. By default, this is 691.		
Outputs of the test	One set of results for the every AD cache used by the Mailbox server role being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	AD access cache hits: Indicates the rate at which requested objects were found in the cache.	Hits/Sec	

	AD access cache misses: Indicates the rate at which requested objects were not found in the cache.	Misses/Sec	A high value of this measure is indicative of the ineffectiveness of the AD cache.
	LDAP search rate from cache: Indicates the number of LDAP search requests issued per second.	Searches/sec	
	Outstanding LDAP async notifies: Indicates the number of LDAP notification requests that are currently outstanding.	Number	
	Outstanding LDAP async reads: Indicates the number of LDAP read requests that are currently outstanding.	Number	
	Outstanding LDAP async searches: Indicates the number of LDAP search requests that are currently outstanding.	Number	

3.1.2 Active Directory Accesses Test

Exchange 2007/2010 uses the Active Directory directory service site topology to determine how messages are transported in the organization.

Exchange 2007/2010 is a site-aware application. Site-aware applications can determine their own Active Directory site membership and the Active Directory site membership of other servers by querying Active Directory. In Exchange 2007/2010, the Microsoft Exchange Active Directory Topology service is responsible for updating the site attribute of the Exchange server object. When an Exchange server role has to determine the Active Directory site membership of another Exchange server role, it can query Active Directory to retrieve the site name.

The Mailbox server role uses Active Directory site membership information to determine which Hub Transport servers are located in the same Active Directory site as the Mailbox servers. The Mailbox server submits messages for routing and transport to a Hub Transport server that has the same Active Directory site membership as the Mailbox server. The Hub Transport server performs recipient resolution and queries Active Directory to match an e-mail address to a recipient account. The recipient account information includes the fully qualified domain name (FQDN) of the user's Mailbox server. The FQDN is used to determine the Active Directory site of the user's Mailbox server. The Hub Transport server delivers the message to Mailbox server within its same Active Directory site, or it relays the message to another Hub Transport server for delivery to a Mailbox server that is outside the Active Directory site. If there are no Hub Transport servers in the same Active Directory site as a Mailbox server, mail cannot flow to that Mailbox

server.

For processing all the Active Directory queries that are required for the aforesaid transactions, the Mailbox server role once again uses site membership to determine which domain controllers and global catalog servers to use. The Mailbox server role then binds to the identified directory servers whenever it needs to read from or write to Active Directory.

Any slowdown therefore, in the communication between the **Mailbox** server role and the marked global catalog servers / domain controllers can significantly delay the identification of the Hub Transport server that the Mailbox server needs to interact with; this in turn can cause delays in message delivery/processing. This test periodically monitors the network connection between the mailbox server role and each identified domain controller, so that communication bottlenecks are swiftly identified and resolved.

Purpose	Periodically monitors the network connection between the mailbox server role and each identified domain controller, so that communication bottlenecks are swiftly identified and resolved		
Target of the test	A server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Mailbox server 3. **PORT** - The port number of the Mailbox server. By default, this is 691		
Outputs of the test	One set of results for every domain controller used by the Mailbox server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	LDAP read calls: Indicates the number of Depth 0 read calls per second that were made by the mailbox server role to this domain controller.	Calls/Sec	
	LDAP search calls: Indicates the number of LDAP Depth 1 or 2 search calls per second that were made by the mailbox server role to this domain controller.	Calls/Sec	

	LDAP searches timed out: Indicates the number of LDAP searches that timed out during the last minute on this domain controller.	Timeouts/min	A high value could indicate any of the following: • Loss of the network connection between the Mailbox server role and the Active Directory directory service domain controller • Non-availability of the domain controller • Critical issues with one/more Active Directory resources To resolve this error, do one or more of the following: • Verify network connectivity between the Mailbox server and the domain controllers it uses. • Ensure that the domain controllers the Mailbox server uses are up and running. • Make sure that none of the Active Directory resource are experiencing performance issues
	LDAP fatal errors: Indicates the number of LDAP errors that caused the Exchange Active Directory Provider to close the LDAP connection without marking the domain controller down during the last minute.	Errors/Min	Ideally, this value should be 0.
	LDAP disconnects: Indicates the number of LDAP errors that caused Exchange Active Directory Provider to mark the domain controller down during the last minute.	Disconnects/Min	

	User search operations failed: Indicates the number of Exchange Active Directory Provider client's searches that failed on this domain controller during the last minute.	Failures/Min	
	Bind failures: Indicates the number of LDAP bind calls that failed during the last minute	Failures/Min	A large number of bind call failures is a cause for concern, as it can disrupt the execution of Active Directory queries.
	Long running LDAP operations: Indicates the number of LDAP operations that the mailbox server performed on this domain controller that took longer than the specified threshold per minute. (Default threshold is 15 minutes.)	Operations/Min	A high value generally indicates performance problems on the said domain controller(s) or network congestion. To resolve this, do one or more of the following: • Ensure that the quality of the network link between the Mailbox server and the domain controllers is good. • Ensure that the domain controller is not experiencing issues in internal operations. You can investigate CPU usage, as well as disk and memory bottlenecks, on your Active Directory directory service servers. • Consider using a dedicated Exchange server and a global catalog server for the expansion of dynamic distribution groups and large distribution groups.
	LDAP pages retrieved: Indicates the number of additional pages retrieved from this domain controller per second.	Pages/sec	
	Outstanding requests to Active Directory: Indicates the number of currently pending LDAP operations to this domain controller.	Number	A high value of this measure or a steady increase in this value is indicative of the poor query processing capability of the domain controller, and would warrant further investigation.

	LDAP read time: Indicates the average time (in ms) taken to send an LDAP read request to the specified domain controller and receive a response.	Msecs	A low value is desired for this measure. A high value or a value that increases consistently is indicative of a gradual slowdown in the domain controller.
	LDAP search time: Indicates the average time (in ms) to send an LDAP search request and receive a response.	Msecs	High LDAP search latencies can be caused by high remote procedure call (RPC) latencies and by increasing queues. High LDAP search latencies generally indicate one of the following problems: • Performance problem with the network connection to the domain controller. • Performance problems with the domain controller itself. To reduce the time it takes for LDAP searches, do one or more of the following: • Ensure that the network performance between the Mailbox server and the domain controllers it uses is not the bottleneck. • Monitor the Searches/sec performance counter to see if there is an unexpected surge in the number of searches the Mailbox server is requesting from the Active Directory directory service. • Ensure that this domain controller is not experiencing performance problems. You can investigate CPU usage, as well as disk and memory bottlenecks, on your Active Directory servers.

3.1.3 Exchange AD Processes Test

The XchgADProcesses test reports whether there is a slow-down in communicating with the global catalogs.

Purpose	Reports whether there is a slow-down in communicating with the global catalogs
Target of the	An Exchange server 2000/2003/2007

test			
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The IP address of the machine configured with the Mailbox server role. 3. **PORT** – The port number through which the Mailbox server communicates. The default is 691. 4. **ISPASSIVE** – If the value chosen is **YES**, then the Exchange server under consideration is a passive server in an Exchange cluster. No alerts will be generated if the server is not running. Measures will be reported as "Not applicable' by the agent if the server is not up. **This parameter can be ignored while configuring this test for a managed "Exchange Mailbox" server**.		
Outputs of the test	One set of results for every Exchange server process		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Ldap read time: The average time that an LDAP read request from the Exchange server takes to be fulfilled.	Secs	The average value should be less than 50 milliseconds. Spikes (Maximum) should not be higher than 100milliseconds.
	Ldap search time: The average time that an LDAP search request takes to be fulfilled.	Secs	The Average value should be less than 50 milliseconds. Spikes (Maximum) should not be higher than 100 milliseconds.

3.1.4 Exchange Clients Test

Monitoring the RPC activity to a Mailbox server and the responsiveness of the server to RPC requests can provide an indication of user satisfaction levels with the performance of the Mailbox server. Foreground RPCs happen during interactions of Outlook clients with the Mailbox server, and any slow down or failure of these RPCs will be directly visible to users of the Mailbox server. Background RPCs are caused by "behind-the-scene" activities internal to the Mailbox server.

This test monitors the performance of RPC activities on the Mailbox server.

Purpose	Monitors the performance of RPC activities on the Mailbox server
Target of the test	A server configured with the Mailbox role
Agent deploying the	An internal agent or remote agent

<table>
<tr><td>test</td><td colspan="3"></td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST –The host for which the test is to be configured.
3. PORT – The port number through which the Mailbox server communicates. The default is 691.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Exchange server being monitored</td></tr>
<tr><td rowspan="6">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>RPC attempts:
The rate at which RPC calls were attempted to the Exchange server during the last measurement period.</td><td>Attempts/sec</td><td>This metric provides an indicator of the RPC workload on the server.</td></tr>
<tr><td>RPC failures:
This metric is the rate of failed RPCs to the Exchange server during the last measurement period.</td><td>Failures/Sec</td><td>Typically, this value should be low</td></tr>
<tr><td>RPC successes:
The rate of successful RPC requests handled by the Exchange server during the last measurement period.</td><td>Successes/Sec</td><td></td></tr>
<tr><td>Foreground RPC failures:
This metric is the client-reported rate of failed foreground RPCs during the last measurement period.</td><td>Failures/sec</td><td>Typically, this value should be low.</td></tr>
<tr><td>Foreground RPC successes:
Shows the client-reported rate of successful foreground RPCs handled by the Exchange server during the last measurement period.</td><td>Successes/Sec</td><td></td></tr>
</table>

	RPC success ratio: This metric is the ratio of the foreground RPC successes to the total number of foreground RPCs attempted during the last measurement period, expressed as a percentage.	%	This metric is one measure of client satisfaction levels with the Exchange server. The closer this value is to 100, the better the client satisfaction level.
	RPCs with latency > 2secs: Shows the client-reported rate of successful RPCs during the last measurement period with latencies > 2 seconds.	Rpcs/sec	
	RPCs with latency > 5secs: Shows the client-reported rate of successful RPCs during the last measurement period with latencies > 5 seconds.	Rpcs/sec	
	RPCs with latency > 10secs: Shows the client-reported rate of successful RPCs during the last measurement period with latencies > 10 seconds	Rpcs/sec	
	Fast RPC ratio: This metric indicates whether client RPCs are happening fast or not.	%	This metric is another key measure of client performance. This metric is computed as the ratio of successful client RPCs with latency less than 2 seconds to the total number of successful RPCs, expressed as a percentage. Hence, the value of this metric indicates the percentage of client RPCs that are taking more than 2 seconds. A value closer to 100 indicates better RPC performance.
	RPC Failed server too busy error rate: Indicates the rate at which client RPCs failed (since the store was started) due to the server too busy ROC error.	Errors/Sec	The value of these measures should be 0 at all times. Non-zero values may indicate that RPC threads are exhausted or client throttling is occurring for clients running versions of Outlook earlier than Microsoft Office Outlook.

	RPC Failed server too busy errors: Indicates the number of client RPCs that failed owing to the server too busy ROC error.	Number	
	Query processor threads: Indicates the number of query processor threads that are currently running queries that are not optimized.	Number	Ideally, the value of this measure should be 0. A non-zero value is indicative of the existence of threads that are executing very slowly because they are running unoptimized/inefficient queries.
	Search threads: Indicates the number of search threads that are currently running queries that are not optimized.	Number	Ideally, the value of this measure should be 0. A non-zero value is indicative of the existence of threads that are executing very slowly because they are running unoptimized/inefficient queries.
	RPC client backoff: Indicates the rate at which the server notifies the client to backoff.	Backoffs/Sec	The Exchange server allows administrators to manage end-user performance by preventing client applications, such as Outlook for example, from sending too many Remote Procedure Call [RPC] requests per second to Exchange, causing the server to suffer in terms of performance. When Exchange determines that a client is having a negative effect on the server, it will send a "back-off" request to the client telling it to delay sending any additional requests for a specified time in order to reduce the performance effect on the server. If these back-off requests are sent by the server frequently, it could indicate that the Exchange server is being bombarded with RPC requests; this is a sign of a server overload condition.

3.2 The Exchange Store Layer

The Exchange store is a storage platform that provides a single repository for managing multiple types of information in one infrastructure.

The Exchange store has several logical components that interact with each other. These components can reside on a single server, or they can be distributed across multiple servers.

- Storage groups (including recovery storage groups), which are logical containers for Exchange databases and their associated system and transaction log files.

- Mailbox databases, which contain the data, data definitions, indexes, checksums, flags, and other information that comprise mailboxes
- Public folder databases, which contain the data, data definitions, indexes, checksums, flags, and other information that comprise any public folders

Using the tests mapped to this layer, you can determine the overall health of the Exchange store and each of its key components.

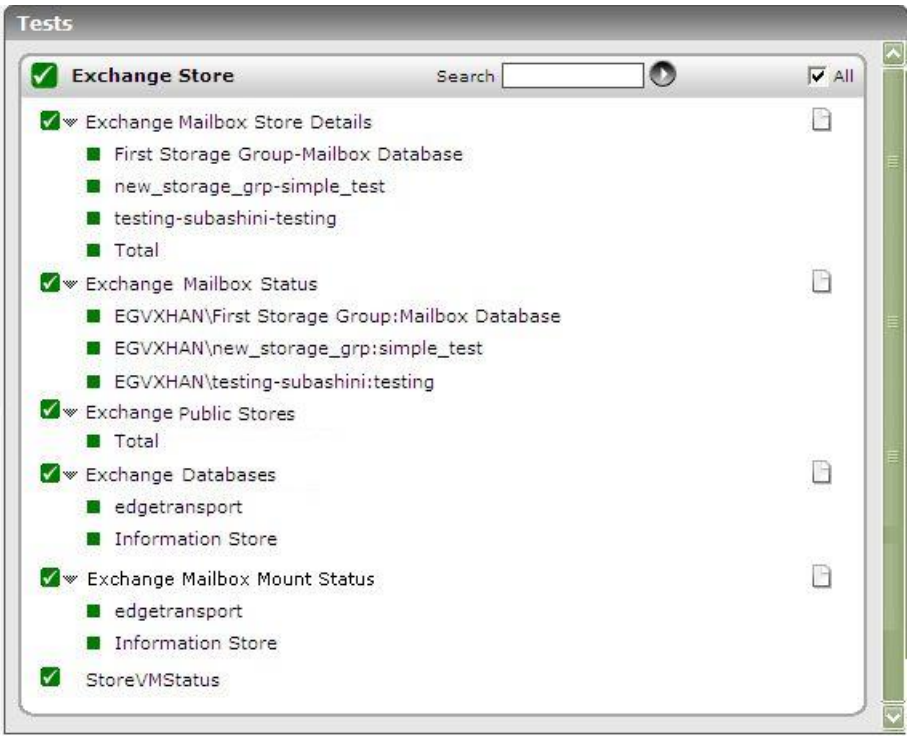

Figure 3.4: The tests associated with the Exchange Store layer

3.2.1 Exchange Database Test

This test measures the performance of the Exchange server database. The Exchange server database comprises of files with ".edb" and ".stm" extensions. A database is a collection of mailboxes. A pair of ".edb" and ".stm" files makes a mailbox.

When an Internet mail message enters into the Exchange server, the body of the message is saved in the ".stm" file, and the header information (From, To, Cc, Time Sent, and so on) is converted to Rich Text Format (RTF), and then stored in the ".edb" file. The transaction log file maintains the state and integrity of ".edb" and ".stm" files.

Purpose	This test monitors the performance of Exchange server database.
Target of the test	An Exchange server 2000/2003/2007
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The IP address of the machine configured with the Mailbox server role.
3. PORT – The port number through which the Mailbox server communicates. The default is 691.
4. ISPASSIVE – If the value chosen is YES, then the Exchange server under consideration is a passive server in an Exchange cluster. No alerts will be generated if the server is not running. Measures will be reported as "Not applicable' by the agent if the server is not up. This parameter can be ignored while configuring this test for a managed "Exchange Mailbox" server.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Mailbox server being monitored</td></tr>
<tr><td rowspan="4">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Database cache hit ratio:
This measure shows the percentage of database requests that were fulfilled by the database buffer pool without incurring disk input/output activity.</td><td>Percent</td><td>A significantly low value indicates that the Exchange server is not having enough free memory. Increasing the memory available to the server may solve this problem.</td></tr>
<tr><td>Database tables cache hit ratio:
This measure shows the percentage of database tables opened using the cached schema information.</td><td>Percent</td><td>A significantly low value indicates that the Exchange server is not having enough free memory. Increasing the memory available to the server may solve this problem.</td></tr>
<tr><td>Log record waits:
This measure shows the number of log records that cannot be added to the log buffers because the log buffers are full.</td><td>Records/Sec</td><td>This measure should be as close to zero as possible.
Abnormal values of this metric indicate that the size of the log buffer is insufficient.
The average value should be below 10 per second. Spikes (maximum values) should not be higher than 100 per second.</td></tr>
<tr><td>Log thread waits:
This measure shows the number of threads waiting for their data to be written to the log buffer so that the update of the database can be completed.</td><td>Number</td><td>This measure should be as low as possible.
A high value for this measure may indicate that the transaction log buffer might be a bottleneck.</td></tr>
</table>

3.2.2 Exchange Mailbox Status Test

Mounting a database puts it online, thereby making its data available to users. If a mailbox database is not mounted,

then users will be denied access to the mailbox data. It is therefore important that the mount status of the mailbox databases is monitored periodically.

The **ExchangeMailboxStatus** test reports the mount status of every mailbox database in an Exchange 2007 mailbox role-enabled server.

Purpose	Reports the mount status of every mailbox database in an Exchange 2007 mailbox role-enabled server
Target of the test	An Exchange server 2007 configured with the Mailbox role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The IP address of the machine configured with the Mailbox server role. 3. **PORT** – The port number through which the Mailbox server communicates. The default is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default.
Outputs of the test	One set of results for every Mailbox database being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Mount status of mailbox: Indicates the mount status of this mailbox database.	Percent	If the value of this measure is *100*, it indicates that the database is mounted. The value *0*, on the other hand, implies that the database is not mounted. An unmounted database can render critical data inaccessible to users. Commonly, mounting issues may occur owing to one/more of the following reasons: ▪ To mount a database, typically, the user should belong to the local Administrators group for the target server and should be assigned the Exchange Server Administrator role. If the user account used for mounting does not have these privileges, then the database will not mount. ▪ You can mount a database only if the Microsoft Exchange Information Store service is running. If this service is not running, then you would be unable to mount the database. ▪ An Exchange mailbox database might not be able to mount if it reaches the 16-GB limit ▪ If a file-level antivirus software deletes or modifies the transaction log files, then the database might not mount. ▪ Hardware issues can prevent a database from mounting. ▪ If Exchange runs out of hard drive space, then the databases might not mount. ▪ If hard disk NTFS file system permissions have been modified, then the databases might not mount.

3.2.3 Exchange Mailbox Mounts Test

This test reports the availability and size of every mailbox database in an Exchange 2010 mailbox role-enabled server.

Purpose	Reports the availability and size of every mailbox database in an Exchange 2010 mailbox role-enabled server
Target of the test	An Exchange server 2010 configured with the Mailbox role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The IP address of the machine configured with the Mailbox server role. 3. **PORT** – The port number through which the Mailbox server communicates. The default is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **SERVERNAME** - Provide the name of the Exchange 2010 server being monitored.
Outputs of the test	One set of results for every Mailbox database being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Mount status: Indicates the availability of this mailbox database.	Boolean	The value 1 for this measure indicates that the database is available, and the value 0 indicates that it is not.
	Database size: Indicates the size of this mailboxdatabase	MB	

3.2.4 Exchange Mailbox Stores Test

The Exchange store is responsible for data storage and management. It is the interface between the clients and the server running Exchange Server. There are three components of the Exchange Store, namely:

- Storage Groups
- Mailboxes
- Public Folders

This test reports statistics pertaining to the Mailbox component of the Exchange store.

Purpose	Measures the performance of the mailbox component of the Exchange Store		
Target of the test	A server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691.		
Outputs of the test	One set of results for every Information store being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active client logons: The number of clients that performed any active logons in the last ten minute interval.	Number	

	Average delivery time: This measure indicates the average time between the submission of a message to the mailbox store and the delivery to all local recipients (recipients on the same server) for the last 10 messages.	Secs	A non-zero value for this measure indicates a change in user workload. An abnormally high value for this measure indicates inability to deliver to one or more destinations. One of the possible reasons for this can be a network failure.
	Categorization count: Indicates the number of categorizations that currently exist in the mailbox store.	Number	Categorizations are created when a user creates a filtered view or performs a search. When the information store must maintain an excessive number of categorizations, performance can be affected.
	Client logons: The number of clients currently logged on.	Number	This measure is a good indicator of the current user activity in the Exchange Server. This information can be used by the administrator for planning the capacity of the mail server.
	Local delivery rate: Indicates the rate at which mails are delivered locally.	Mails/Sec	
	Messages queued for submission: The current number of submitted messages which are not yet processed by transport.	Number	
	Message delivery rate: Indicates the rate at which messages are delivered to all recipients.	Mails/Sec	
	Messages opened: This measure indicates the rate at which the requests to open the messages are being submitted to the information store.	Msgs/Sec	This measure shows the overall picture of user activity. An abnormally high value for this measure may indicate that the Exchange 2000 Server is overloaded.
	Messages sent: This measure indicates the rate at which messages are sent to transport by the information store.	Msgs/Sec	

	Messages submitted: Indicates the rate at which mails are submitted by clients.	Msgs/Sec	
	Receive queue size: This measure shows the number of messages that are currently in the receiving queue of the information store.	Number	This measure is usually zero under normal conditions. A non-zero value for this measure indicates that the SMTP service is choking up memory.
	Search task rate: Indicates the number of search tasks created per second.	Tasks/Sec	
	Slow FindRow Rate: Indicates the rate at which the slower FindRow needs to be used in the mailbox store.	Finds/Sec	Higher values indicate applications are crawling or searching mailboxes, which is affecting server performance. These include desktop search engines, customer relationship management (CRM), or other third-party applications.
	RPC average latency: Indicates the RPC latency averaged across all operations in the last 1024 packets.	Msec	This measure provides the best indication of whether counters with high database latencies are actually impacting Exchange health and client experience. Often, high RPC averaged latencies are associated with a high number of RPC requests, which should be less than 70 at all times. The value of this measure should not be higher than 25 msecs on an average.

3.2.5 Exchange Public Stores Test

This test reports statistics pertaining to the public folders of the Information store.

Purpose	Measures the performance of the public store component of the Exchange Information Store
Target of the test	An Exchange 2007 server configured with the Mailbox role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691.

Outputs of the test	One set of results for every Information store being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active client logons: The number of clients that performed any active in the last ten minute interval.	Number	
	Average delivery time: This measure indicates the average time between the submission of a message to the public folder store and the delivery to all local recipients (recipients on the same server) for the last 10 messages.	Secs	A non-zero value for this measure indicates a change in user workload. An abnormally high value for this measure indicates inability to deliver to one or more destinations. One of the possible reasons for this can be a network failure.
	Categorization count: Indicates the number of categorizations that currently exist in the public folder store.	Number	Categorizations are created when a user creates a filtered view or performs a search. When the information store must maintain an excessive number of categorizations, performance can be affected.
	Client logons: The number of clients currently logged on.	Number	This measure is a good indicator of the current user activity in the Exchange Server. This information can be used by the administrator for planning the capacity of the mail server.
	Messages queued for submission: The current number of submitted messages which are not yet processed by transport.	Number	
	Messages opened: This measure indicates the rate at which the requests to open the messages are being submitted to the information store.	Msgs/Sec	This measure shows the overall picture of user activity. An abnormally high value for this measure may indicate that the Exchange 2000 Server is overloaded.
	Messages delivered: Indicates the current number of messages delivered to all recipients.	Number	

	Messages sent: This measure indicates the rate at which messages are sent to transport by the information store.	Msgs/Sec	
	Messages submitted: Indicates the rate at which mails are submitted by clients.	Msgs/Sec	
	Receive queue size: This measure shows the number of messages that are currently in the receiving queue of the information.	Number	This measure is usually zero under normal conditions. A non-zero value for this measure indicates that the SMTP service is choking up memory.
	Search task rate: Indicates the number of search tasks created per second.	Tasks/Sec	
	Slow FindRow Rate: Indicates the rate at which the slower FindRow needs to be used in the mailbox store.	Finds/Sec	Higher values indicate applications are crawling or searching mailboxes, which is affecting server performance. These include desktop search engines, customer relationship management (CRM), or other third-party applications.
	Replication receive queue: This measure indicates the number of replication messages waiting to be processed currently.	Number	A consistent increase in this value could indicate a bottleneck in the processing of replication messages.

3.2.6 Store VM Status Test

Each store.exe process of a server has limited amount of memory called the Store Virtual memory that it can address. This test reports statistics related to the usage of the Exchange store's virtual memory.

Purpose	This test monitors the performance of Exchange server database.
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST – Indicates the IP address of the Mailbox server.
3. PORT – The port number through which the Mailbox server communicates. By default, this is 691.
4. ISPASSIVE – If the value chosen is YES, then the Exchange server under consideration is a passive server in an Exchange cluster. No alerts will be generated if the server is not running. Measures will be reported as “Not applicable” by the agent if the server is not up.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Exchange server being monitored</td></tr>
<tr><td rowspan="5">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Largest block size:
It is the largest free block of virtual memory.</td><td>MB</td><td>At no point should this value go below 32 MB. As you scale a server to accommodate more users and more usage, the server may run low on virtual memory. When a server is low on virtual memory, overall performance degrades as the low situation forces the store.exe process to use the page file and begin paging rapidly.</td></tr>
<tr><td>16MB free blocks in virtual memory:
The total number of free virtual memory blocks that are greater than or equal to 16MB.</td><td>Number</td><td>At no point should this value go below 1.</td></tr>
<tr><td>Free blocks in virtual memory:
The total number of free virtual memory blocks regardless of size.</td><td>Number</td><td>At no point should this value go below 1.</td></tr>
<tr><td>Large memory blocks in virtual memory:
The sum of all the free virtual memory blocks that are greater than or equal to 16MB</td><td>MB</td><td>At no point should this value go below 50 MB.</td></tr>
</table>

3.3 The Mailbox Services Layer

You can use the tests associated with this layer to monitor the core services offered by a Mailbox server, such as the Exchange Search and Mailbox Assistants service, and also monitor the RPC mechanisms between the Mailbox server and clients.

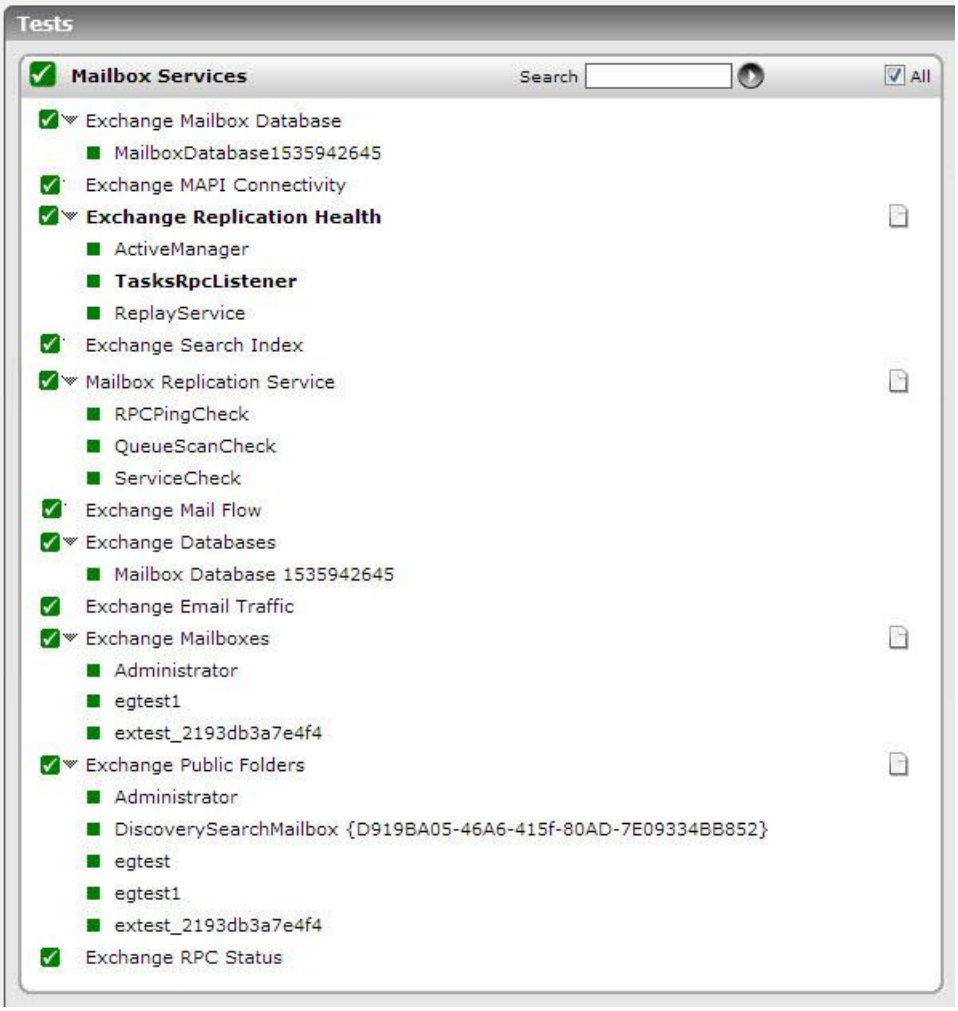

Figure 3.5: The tests mapped to the Mailbox Services layer

3.3.1 Exchange Public Folders Test

This test retrieves information about the public folders in the mailboxes on a server, including the number and size of items in the folder, the folder name and ID, and other information.

Purpose	Retrieves information about the public folders in the mailboxes on a server, including the number and size of items in the folder, the folder name and ID, and other information		
Target of the test	A server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **SERVERNAME** - Specify the name of the target Exchange server. 6. **FOLDER SCOPE** - Indicate whether the monitoring scope of this test should be restricted to important folders alone. If so, then pick the **Important** option from this list. In this case, critical folders such as inbox, outbox, deleted items, sent items and calendar details, will alone be monitored. On the other hand, to monitor all folders in a mailbox, pick the **All** option. 7. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**. 8. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for each monitored folder		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Items in folder:** Indicates the number of items in this folder.	Number	
	Folder size: Indicates the size of this folder.	KBytes	
	Items in folders and subfolder: Indicates the number of items available in this folder and its subfolders.	Number	
	Folder and subfolder size: Indicates the total size of the folder and its subfolders.	KBytes	

3.3.2 Exchange Mail Flow Test

This test checks whether mail can be successfully sent from and delivered to the system mailbox on a computer that has the Mailbox server role installed. This test is also used to verify whether the e-mail is sent between Mailbox servers within a defined latency threshold.

Purpose	Checks whether mail can be successfully sent from and delivered to the system mailbox on a computer that has the Mailbox server role installed. This test is also used to verify whether the e-mail is sent between Mailbox servers within a defined latency threshold
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **TARGET EMAIL ADDRESS** - This test sends mails to a configured mail ID to check whether the mails are successfully delivered to that ID. Specify the target email address here.		
Outputs of the test	One set of results for the server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Mail flow latency time: Indicates the time taken to deliver a message.	Secs	A very high value for this measure could indicate delivery bottlenecks.
	Mail flow result: Indicates whether the mail was successfully delivered or not.	Percent	The value 100 for this measure indicates that the mail flow is successful, and the value 0 for this measure indicates failure.

3.3.3 Exchange MAPI Connectivity Test

This test is used to verify exchange server functionality by logging on to the specified mailbox.

Purpose	Used to verify exchange server functionality by logging on to the specified mailbox
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **TARGETMAILADDRESS** - This test emulates a MAPI connection to a configured mailbox on the Exchange server to verify the availability of the MAPI connection and the time taken to establish the same. For this emulation, the email address of a mailbox on the Exchange server is required. Provide this email address against **TARGETMAILADDRESS**.		
Outputs of the test	One set of results for the server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Latency: Indicates the time taken to establish the MAPI connection.	MSecs	A very high value for this measure could indicate connection bottlenecks.
	MAPI connectivity status: Indicates whether the MAPI connection was successfully established or not.	Percent	This measure will report 100, if connection to a mailbox is successfull; otherwise, this measure will report 0.

3.3.4 Exchange Search Index Test

This test reports whether the Exchange search is currently enabled and is indexing new e-mail messages in a timely manner.

Purpose	Used to verify exchange server functionality by logging on to the specified mailbox
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **TARGETMAILADDRESS** - Provide the email address of a mailbox on the Exchange server.		
Outputs of the test	One set of results for the server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Search time: Indicates the time taken to index the messages.	Secs	A very high value for this measure could indicate issues in indexing messages.
	Is exchange search enabled?: Indicates whether the Exchange search is enabled or not.	Boolean	The value 1 for this measure indicates that the exchange search is enabled and value 0 indicates it is not enabled.

3.3.5 Exchange Mailbox Database Test

This test reports the size and moun status of each mailbox database on the Exchange 2007 serer, and also reveals the number of mailboxes in each database.

Purpose	Reports the size and moun status of each mailbox database on the Exchange 2007 serer, and also reveals the number of mailboxes in each database
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default.		
Outputs of the test	One set of results for each mailbox database on the server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Size of mailbox database: Indicates the size of this mailbox database.	GB	
	No of mailboxes in database: Indicates the number of mailboxes in this database.	Number	
	Mount status of mailbox database: Indicates whether this mailbox database has been mounted or not.	Boolean	The value '1' indicates that the mailbox is currently available - i.e., has been mounted, and '0' indicates that the mailbox is not available or has not been mounted yet.

3.3.6 Exchange Mailboxes Test

This test reports critical statistics pertaining to each mailbox on an Exchange server, such as the mailbox size, the mailbox quota, the size and number of deleted items, etc., and also reveals how effectively the mailbox has been utilized.

Purpose	This test reports critical statistics pertaining to each mailbox on an Exchange server, such as the mailbox size, the mailbox quota, the size and number of deleted items, etc., and also reveals how effectively the mailbox has been utilized.
Target of the test	A server configured with the Mailbox role
Agent	An internal agent

<table>
<tr><td>deploying the test</td><td colspan="3"></td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - Indicates how often this test needs to be executed.
2. HOST – Indicates the IP address of the Mailbox server.
3. PORT – The port number through which the Mailbox server communicates. By default, this is 691.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. SERVERNAME - Specify the name of the target Exchange server.
6. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
7. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each mailbox on the server monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Deleted item count:
Indicates the number of deleted messages in this mailbox.</td><td>Number</td><td></td></tr>
<tr><td>Total deleted item size:
Indicates the total size of all deleted items in this mailbox.</td><td>MB</td><td></td></tr>
</table>

	Total items: Indicates the total number of messages in this mailbox.	Number	
	Size: Indicates the total size of this mailbox.	MB	
	Quota: Indicates the allocated size of this mailbox.	MB	If the mailbox **Size** is dangerously close to the **Quota**, you may want to either reset the quota or delete unnecessary messages from the mailbox to free up space. This needs to be done so that the mailbox does not violate its quota, causing mails to bounce.
	Mailbox usage: Indicates the mailbox usage.	Percent	

3.3.7 Exchange Search Test

Exchange Search is a feature that enables fast searching of text in messages through the use of pre-built indexes. Exchange Search uses the Microsoft Search indexing engine, and creates the initial index by "crawling" all messages in mailboxes moved into an Exchange 2007/2010 database, and updates this index based on notifications from the information store as new messages arrive. This test summarizes the performance statistics of the Microsoft Search indexing engine for every Exchange 2007/2010 database.

Purpose	Summarizes the performance statistics of the Microsoft Search indexing engine		
Target of the test	A server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691.		
Outputs of the test	One set of results for Mailbox server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Mailboxes left to crawl:** Indicates the number of mailboxes that are currently left to be crawled on this database.	Number	Ideally, the value of this measure should be low. A high value is indicative of an inefficient search indexing engine.
	Documents to be indexed: Indicates the number of documents that are currently awaiting indexing.	Number	
	Documents indexed: Indicates the number of documents that are being actively indexed in this database.	Number	
	Documents successfully indexed: Indicates the number of documents that were currently indexed successfully.	Number	
	Documents failed during indexing: Indicates the number of documents have not been indexed in this database since the last measurement period.	Number	A low value is desired for this measure. An unusually high value might call for further investigation.
	Average latency of RPCs to the information store: Indicates the average latency of RPCs (in milliseconds) to the Exchange Information Store service.	Msecs	Typically, Exchange Search makes these RPC calls for crawling purposes for the given database. High RPC latencies are quiet obviously a cause for concern, as they can slow down crawling, and ultimately the search process.
	Average document indexing time: Indicates the average (in milliseconds) of how long it takes to index documents.	Msecs	A high value of this measure could indicate an indexing bottleneck.

	Is full crawl mode in use for indexing?: Indicates whether this database is going through a full crawl or not, currently	Status	When the database is still being crawled, it has a value of **1**. When the crawl is complete, the value is **0**.

3.3.8 Mailbox Assistants Test

The Microsoft Exchange Mailbox Assistant service provides functionality for Calendar Attendant, Resource Booking Attendant, Out of Office Assistant, and Managed Folder Mailbox Assistant.

The Exchange Assistants can be either event-based Assistants or time-based Assistants. The event-based Assistants start to process mailboxes on the occurrence of an event, such as on a change of Out-of-Office (OOF) information in one or more mailboxes. The time-based Assistants process the mailboxes periodically. Each time-based Assistant deploys an Assistants Driver that periodically checks whether the current time is within in a specified time window. When the current time reaches the specified time window, the Assistants Driver invokes the corresponding time-based Assistant. The time-based Assistant then obtains a list of mailboxes from the database and starts to process them.

To determine whether there are any mailbox assistant-related issues, use the MailboxAssistants test.

Purpose	To determine whether there are any mailbox assistant-related issues		
Target of the test	A server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691.		
Outputs of the test	One set of results for the Mailbox server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Events waiting to be processed by assistants: Indicates the current number of events in the in-memory queue waiting to be processed by the assistants.	Number	Ideally, the value of this measure should be low at all times. High values may indicate a performance bottleneck.

	Events processed rate by Exchange assistants: Indicates the number of events processed per second.	Events/Sec	While sporadic dips in this value can be ignored, a consistent decrease could be a cause for concern, and may warrant a thorough investigation.
	Events polled by Exchange assistants: Indicates the number of events polled per second.	Events/Sec	
	Delay between polling and event creation: Indicates the current latency between when the most recent MAPI event was polled and when the event was created.	Secs	
	Event queueing time - average: Indicates the average time (in seconds) that the event lives in the dispatcher queue.	Secs	Ideally, the value of this measure should be low. A very high value indicates that there are many events in queue with long waiting times. This in turn indicates a processing bottleneck.
	Event processing time - average: Indicates the average time (in seconds) that the assistants took for processing events.	Secs	Ideally, the value of this measure should be low. A very high value indicates that there are many events in queue with long waiting times. This in turn indicates a processing bottleneck.
	Queue size of event dispatchers - average: Indicates the average queue size of event dispatchers.	Number	A gradual but steady increase in the value of this measure over time, could indicate a problem with the dispatcher queue.
	Mailbox processing time - average: Indicates the average processing time of mailboxes for time-based assistants.	Secs	A low value is typically desired.

	Failed event dispatchers: Indicates the percentage of Event Dispatchers that are in failure mode, currently.	Percent	
	Mailboxes processed: Indicates the rate at which time-based assistants processed mailboxes.	Operations/Sec	
	Threads in use from the CLR thread pool: Indicates the current number of Threads used from the CLR thread pool.	Number	

3.3.9 Exchange PC Status Test

When using Outlook clients in MAPI mode, clients' actions in Outlook translate to remote procedure calls (RPCs) between the clients and the server. If the client is running in online mode, these RPC calls occur synchronously. Any delay by the server in fulfilling these synchronous requests directly affects user experience and the responsiveness of Outlook. Conversely, if the client is running in cached mode, the majority of these requests will be handled asynchronously. Asynchronous processing means that the performance of the RPC mechanism does not affect the overall user experience.

This test monitors the performance of RPC mechanisms between the clients and the Exchange server.

Purpose	Monitors the performance of RPC mechanisms between the clients and the Exchange server
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **ISPASSIVE** – If the value chosen is **YES**, then the Exchange server under consideration is a passive server in an Exchange cluster. No alerts will be generated if the server is not running. Measures will be reported as "Not applicable' by the agent if the server is not up.
Outputs of the test	One set of results for every Mailbox server being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	RPC operations: Indicates the rate of RPC operations handled by the Exchange information store during the last measurement period.	Operations/Sec	Generally, spikes in RPC requests that do not increase RPC operations/sec indicate that there are bottlenecks preventing the store from fulfilling the requests in a timely manner. It is relatively simple to identify where the bottlenecks are occurring with regards to RPC requests and RPC operations/sec. If the client experiences delays, but the RPC requests are zero and the RPC operations/sec are low, the performance problem is happening before Exchange processes the requests (that is, before the Microsoft Exchange Information Store service actually gets the incoming requests). All other combinations point to a problem either while Exchange processes the requests or after Exchange processes those requests.
	Current RPC requests: Indicates the number of MAPI RPC requests presently being serviced by the Microsoft Exchange Information Store service.	Number	The Exchange server is configured with a pre-set maximum number of RPC requests that can be handled simultaneously (default is 100). If this value is exceeded, client requests to the server will be rejected. This measure should be below 30 most of the time.
	RPC traffic: Indicates the number of MAPI RPC packets being handled by the Exchange Information Store during the last measurement period.	Packets/Sec	
	RPC latency: Indicates the RPC latency in milliseconds, averaged for the past 1024 packets.	Secs	This value should be below 50ms at all times. A slowdown in RPC packet processing can adversely impact the user experience.

3.3.10 Exchange Replication Test

Cluster continuous replication (CCR) is a high availability feature of Microsoft Exchange server 2007/2010 that combines the asynchronous log shipping and replay technology built into Exchange

2007/2010 with the failover and management features provided by the Cluster service.

CCR uses the database failure recovery functionality in Exchange 2007/2010 to enable the continuous and asynchronous updating of a second copy of a database with the changes that have been made to the active copy of the database. During installation of the passive node in a CCR environment, each storage group and its database is copied from the active node to the passive node. This operation is called *seeding*, and it provides a baseline of the database for replication. After the initial seeding is performed, log copying and replay are performed continuously.

In a CCR environment, the replication capabilities are integrated with the Cluster service to deliver a high availability solution

Local continuous replication (LCR) is a single-server solution that uses built-in asynchronous log shipping and log replay technology to create and maintain a copy of a storage group on a second set of disks that are connected to the same server as the production storage group. The production storage group is referred to as the *active* copy, and the copy of the storage group maintained on the separate set of disks is referred to as the *passive* copy.

The following figure illustrates a basic deployment of LCR:

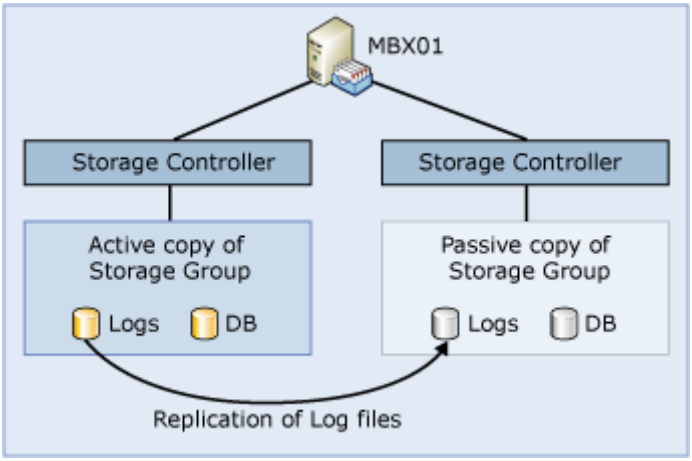

Figure 3.6: A basic LCR deployment

This test monitors the health of the replication in both LCR and CCR.

Purpose	Monitors the health of the replication in both LCR and CCR
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691.
Outputs of the test	One set of results for the Mailbox server being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Replay queue length: Indicates the number of log generations currently waiting to be replayed.	Number	
	Copy queue length: Indicates the number of log files currently waiting to be copied and inspected.	Number	A high value of this measure could indicate that Exchange database copy failed. Under such circumstances, do the following: • Remove passive DB copy • Delete DB and log files from disk • Force AD replication between AD sites • Dismount AD database • Delete all EO*.log files • Mount database • Create DB copy

3.3.11 Exchange Storage Groups Test

An Exchange *storage group* is a logical container for Exchange databases and their associated system and transaction log files.

This test monitors the mailboxes in a storage group to report key statistics such as the number of mailboxes that have exceeded their storage limit, the number of mails in the mailboxes, etc.

Purpose	Monitors the mailboxes in a storage group
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST – Indicates the IP address of the Mailbox server.
3. PORT – The port number through which the Mailbox server communicates. By default, this is 691.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
6. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
➢ The eG manager license should allow the detailed diagnosis capability
➢ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results every storage group auto-discovered from the Mailbox being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Disconnected users:
Indicates the number of users who are currently disconnected from their mailboxes in a storage group.</td><td>Number</td><td></td></tr>
</table>

	Users with storage limit warning status: Indicates the number of users whose mailboxes are currently reaching their storage quota and are therefore receiving the storage quota warning message.	Number	A *storage quota* is a storage size limit for a mailbox or a public folder. You can use the Exchange Management Console or the Exchange Management Shell to view or set the storage quotas for all of the mailboxes or public folders in a database. You can also use the Exchange Management Console or the Exchange Management Shell to set storage quotas on a per-mailbox basis, thereby overriding the storage quotas that are set at the database level. However, storage quotas for individual public folders can be viewed or set only in the Exchange Management Shell. If a mailbox size reaches or exceeds the **Issue warning** value that has been configured for a mailbox/database, then Exchange will automatically send a warning message to that mailbox. A non-zero value for this measure therefore, indicates that one/more mailboxes in the storage group are in danger of exceeding their storage quota. You can use the detailed diagnosis capability of this measure to identify the mailboxes that might be fast exhausting their storage quota. Based on the detailed diagnosis, you can decide whether to revise the storage quota or clear unnecessary mails from the listed mailboxes to make more space therein.
	Users with storage limit prohibit status: Indicates the number of users whose mailboxes have reached their storage quota currently and hence, will not be able to send any emails.	Number	If a mailbox size reaches or exceeds the limit specified against **Prohibit send at**, Exchange will prevent that mailbox user from sending new messages and will display a descriptive error message. A non-zero value for this measure is a cause for concern, as it indicates that one/more users might not be able to send out critical mails from their mailboxes. Use the detailed diagnosis of this test to identify the affected users, and then decide whether to increase the **Prohibit sent at** limit or simply remove unnecessary mails from the users' mailboxes to make space available.
	Total number of mails: Indicates the current number of e-mails in all mailboxes homed in a storage group.	Number	

	Total mail size: Indicates the total size of all mailboxes homed in this storage group.	MB	
	Total deleted mails: Indicates the number of deleted e-mails currently in all mailboxes homed in a storage group.	Number	
	Total deleted mail size: Indicates the current size of all deleted e-mails in all mailboxes homed in a storage group.	MB	

3.3.12 Exchange Email Traffic Test

Periodic workload monitoring is imperative to evaluate the processing ability of the Exchange 2010 server and to proactively detect potential overload conditions. By continuously monitoring the email traffic to and from the Exchange server, this test turns a spotlight on the workload of the Exchange server, helps detect overload conditions, and also points you to the source of the overload – mails sent/received by users in the intranet? Or mail traffic over the internet?

Purpose	Turns a spotlight on the workload of the Exchange server, helps detect overload conditions, and also points you to the source of the overload – mails sent/received by users in the intranet? Or mail traffic over the internet?
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST – Indicates the IP address of the Mailbox server.
3. PORT – The port number through which the Mailbox server communicates. By default, this is 691.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
6. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
7. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Mailbox server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Internal mails received:
Indicates the number of mails received by the Exchange server from the intranet.</td><td>Number</td><td>In the event of an overload, you can compare the value of this measure with the value of the Internal mails sent, External mails received, and External mails sent measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet?
To know the email IDs that received the emails, the number of emails that each ID received, and the total size of the emails to an ID, use the detailed diagnosis of this test.</td></tr>
</table>

	Internal mails received size: Indicates the total size of the mails received by the Exchange server from the intranet.	KB	
	Internal mails sent: Indicates the number of mails sent by the Exchange server to the intranet.	Number	In the event of an overload, you can compare the value of this measure with the value of the *Internal mails received, External mails received*, and *External mails sent* measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet? To know the email IDs that sent the emails, the number of emails that each ID sent, and the total size of the emails from an ID, use the detailed diagnosis of this test.
	Internal mail sent size: Indicates the total size of the mails sent by the Exchange server to the intranet.	KB	
	External mails received: Indicates the number of mails received by the Exchange server from the internet.	Number	In the event of an overload, you can compare the value of this measure with the value of the *Internal mails received*, *Internal mails sent*, and *External mails sent* measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet? To know the email IDs that received the emails, the number of emails that each ID received, and the total size of the emails to an ID, use the detailed diagnosis of this test.
	External mail received size: Indicates the total size of the mails received by the Exchange server from the internet.	KB	

	External mails sent: Indicates the number of mails sent by the Exchange server to the internet.	Number	In the event of an overload, you can compare the value of this measure with the value of the *Internal mails sent*, *Internal mails sent*, and *External mails received* measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet? To know the email IDs that sent the emails, the number of emails that each ID sent, and the total size of the emails from an ID, use the detailed diagnosis of this test.
	External mail sent size: Indicates the total size of mails sent by the Exchange server to the internet.	KB	

Use the detailed diagnosis of the *Internal mails received* measure to know the internal email IDs that received the emails, the number of emails that each ID received, and the total size of the emails received by an ID. This way, you can quickly identify the email ID that received the maximum number of emails and that which received mails of the maximum size.

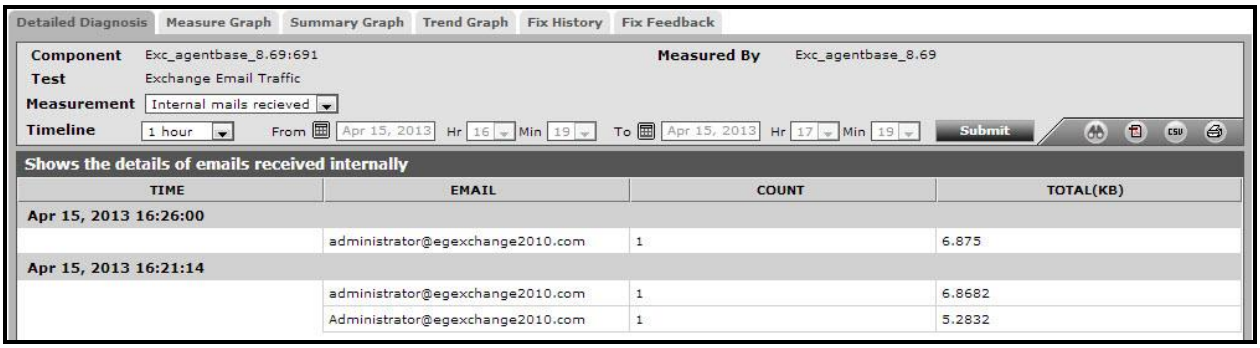

Figure 3.7: The detailed diagnosis of the Internal mails received measure

Use the detailed diagnosis of the *Internal mails sent* measure to know the internal email IDs that sent the emails, the number of emails that were sent from each ID, and the total size of the emails sent from an ID. This way, you can quickly identify the email ID that sent the maximum number of emails and that which sent mails of the maximum size.

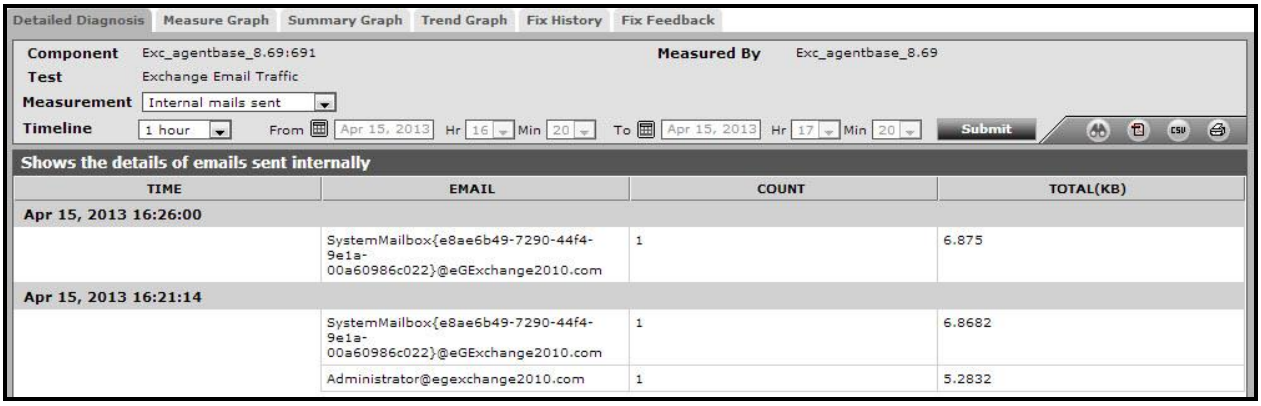

Figure 3.8: The detailed diagnosis of the Internal mails sent measure

3.3.13 Exchange Replication Health Test

This test checks all aspects of replication and replay and reports on the health of each aspect. It is designed for the proactive monitoring of continuous replication and the continuous replication pipeline, the availability of Active Manager, and the health and status of the underlying cluster service, quorum, and network components.

Purpose	Checks all aspects of replication and replay and reports on the health of each aspect. It is designed for the proactive monitoring of continuous replication and the continuous replication pipeline, the availability of Active Manager, and the health and status of the underlying cluster service, quorum, and network components
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **TRANSIENTEVENTSUPPRESSIONWINDOW** – This parameter specifies the number of minutes that the queue lengths can be exceeded before the queue length tests are considered to have failed. This parameter is used to reduce the number of failures due to transient load generation. By default, this parameter is set to *3* minutes. 6. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. 7. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for each aspect of replication that is tested		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td>test</td>
<td>Mailbox replication service status:
Indicates the connectivity status of this aspect of replication.</td>
<td></td>
<td>If the value of this measure is Success, it indicates that this replication aspect is in good health currently. If the value of this measure is Failure, it indicates problems in this replication aspect. The numeric values that correspond to these measure values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Success</td><td>1</td></tr>
<tr><td>Failure</td><td>0</td></tr>
</table>
If the status reported by this measure is Failure, you can use the detailed diagnosis of this test to determine the reason for the failure.
Note:
Typically, this measure reports the Measure Values listed in the table above to indicate status of each replication-related activity that is monitored. However, in the graph of this measure, the Numeric values are used to represent replication health.</td>
</tr>
</table>

If some aspect related to replication reports the status as *Failed*, you can use the detailed diagnosis of the **Exchange Replication Health** test to view the error that caused the failure.

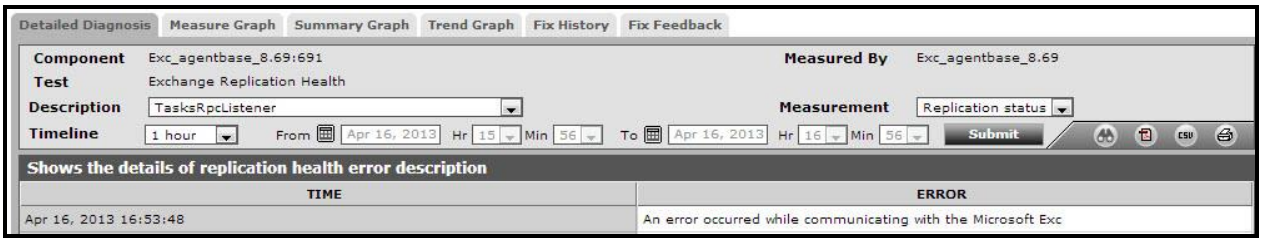

Figure 3.9: The detailed diagnosis of the Replication health test

3.3.14 Mailbox Replication Service Test

The Mailbox Replication Service (MRS), which resides on all Microsoft Exchange Server 2010 Client Access servers, is the service responsible for moving mailboxes, importing and exporting .pst files, and restoring disabled and soft-deleted mailboxes. MRS plays an integral role in migrations from Exchange 2003 or 2007 to Exchange 2010 because moving mailboxes is the only practical way to get user data into mailboxes. If you encounter errors during such migrations, you can use this test to make sure that MRS is running and that it responds to a remote procedure call (RPC) ping check.

Purpose	Use this test to make sure that MRS is running and that it responds to a remote procedure call (RPC) ping check		
Target of the test	A server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. 6. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for each aspect of mailbox replication service health that is tested by this test		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td>test</td>
<td>Mailbox replication service status:
Indicates the current status of this aspect of MRS health.</td>
<td></td>
<td>If the value of this measure is True, it indicates that this aspect is in good health currently. If the value of this measure is False, it indicates problems in this aspect. The numeric values that correspond to these measure values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>True</td><td>1</td></tr>
<tr><td>False</td><td>0</td></tr>
</table>
If the status reported by this measure is False, you can use the detailed diagnosis of this test to determine the reason for the failure of the corresponding aspect.
Note:
Typically, this measure reports the Measure Values listed in the table above to indicate the status of MRS. However, in the graph of this measure, the Numeric values are used to represent MRS health.</td>
</tr>
</table>

If the value of the *Mailbox replication service status* measure is *False* for any of the checks that are performed by the test, then, you can use the detailed diagnosis of this test to figure out why that check failed.

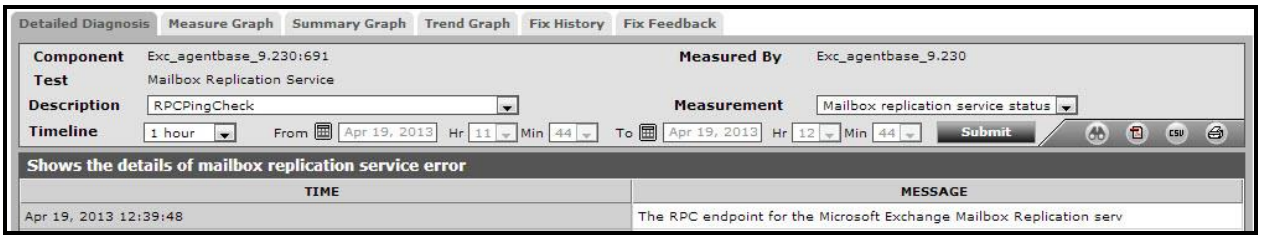

Figure 3.10: The detailed diagnosis of the Mailbox Replication Service test

3.3.15 Exchange Databases Test

Mailbox Database can be considered a container that stores and maintains all mailboxes of the users. When administrators install Microsoft Exchange Server 2010 one Mailbox Database is automatically created which enables the Exchange Server to create mailboxes for several users. Since mailboxes constantly grow in size and number, these mailbox databases should be adequately sized. Lack of space in an exchange database can cause the mailboxes in that database to stop receiving emails! To avoid such an outcome, administrators can use this test to continuously track the space usage in each of the Exchange mailbox databases configured on the Exchange server and rapidly isolate the mailbox database that is running out of space. The test also helps administrators differentiate between actual free space and white space.

This test applies only to Exchange 2010 servers.

Purpose	Continuously tracks the space usage in each of the Exchange mailbox databases configured on the Exchange server and rapidly isolates the mailbox database that is running out of space; the test also helps administrators differentiate between actual free space and white space		
Target of the test	An Exchange 2010 server configured with the Mailbox role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691.		
Outputs of the test	One set of results for each mailbox database on the Exchange 2010 server		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total drive space: Indicates the total size of this database.	GB	
	Free space in drive: Indicates the amount of space lying unused in this mailbox database.	GB	A high value is desired for this measure. **Note that this value does not include white space**.
	Disk free space in drive: Indicates the percentage of free space in this mailbox database.	Percent	Ideally, the value of this measure should be high. A low value or a consistent decrease in this value is a cause for concern, as it indicates that the space in the mailbox database is getting rapidly eroded. You can compare the value of this measure across mailbox databases to identify that database that has very low free space. If any database has less than 5% free space, it is worrysome, since this could mean that the mailbox database is experiencing a severe space crunch; this in turn can cause that mailbox to stop receiving mails. **Note that this percentage does not consider the whitespace in the Exchange database**.
	White space: Indicates the total amount of whitepspace in this mailbox database.	GB	Exchange 2010 constantly runs an online defragmentation process, which automatically deletes content that has passed the deleted item retention period that is configured in Exchange. The free space that is created in the Exchange mailbox database as a result of this deletion is known as whitespace. When new content comes in, this whitespace will be used by Exchange 2010 before expanding the database any larger.

	Total free space in drive: Indicates the total amount of space that is free in this mailbox database.	GB	This is the total free space in the database, inclusive of the whitespace.
	Total disk free space in drive: Indicates the percentage of total space in this mailbox database that is currently free.	Percent	Ideally, the value of this measure should be high. A low value or a consistent decrease in this value is a cause for concern, as it indicates that the space in the mailbox database is getting rapidly eroded. You can compare the value of this measure across mailbox databases to identify that database that has very low free space. If any database has less than 5% free space, it is worrysome, since this could mean that the mailbox database is experiencing a severe space crunch; this in turn can cause that mailbox to stop receiving mails. **Note that this percentage includes the whitespace in the Exchange database**.

3.3.16 Exchange ActiveSync Connectivity Test

Exchange ActiveSync lets you synchronize a mobile device with your Exchange 2010 mailbox, so that you can check your emails from your mobile phone itself! Whenever a mobile phone user complaints that he/she is unable to check or is experiencing slowness when checking emails on his/her mobile phone, Exchange administrators need to quickly determine what is causing the non-sync – is it because ActiveSync is unable to synchronize with the user's mailbox? Or is it because ActiveSync is taking too long to perform the synchronization? At which stage of the synchronization did the failure/delay occur? This test helps answer all these questions. The test periodically checks ActiveSync connectivity at every stage (a.k.a scenario) of the synchronization – eg., the Logon stage, the FolderSync stage, the Options stage, etc. - reports issues and latencies (if any) in connectivity, and leads you to the exact stage at which the failure/slowdown occurred.

Purpose	Periodically checks ActiveSync connectivity at every stage of the synchronization – eg., the Logon stage, the FolderSync stage, the Options stage, etc. - reports issues and latencies (if any) in connectivity, and leads you to the exact stage at which the failure/slowdown occurred
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Client Access server.
3. PORT – The port number of the client access server. By default, this is 110.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. CLIENT ACCESS SERVER – Specify the fully-qualified domain name of the Client Access server.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each <ClientAccessServer>/<LocalSiteNameofClientAccessServer>/<SynchronizationStage/Scenario tested> combination</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>ActiveSync connectivity status:
Indicates whether the ActiveSync connectivity check was successful or not at this stage/scenario of the synchronization.</td><td></td><td>If the value of this measure is Success, it indicates that the ActiveSync connectivity check was successful at this stage. If the value of this measure is Failure, it indicates that mailbox synchronization using ActiveSync failed at this stage. The numeric values that correspond to these measure values are as follows:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Success</td><td>1</td></tr><tr><td>Failure</td><td>0</td></tr></table>
Note:
Typically, this measure reports the Measure Values listed in the table above to indicate the ActiveSync connectivity status. However, in the graph of this measure, the Numeric values are used to represent the connectivity status.</td></tr>
</table>

	ActiveSync latency: Indicates the time taken by ActiveSync to successfully complete this stage/scenario of the synchronization.	Secs	A low value is desired for this measure. A high value indicates that this stage/scenario of the synchronization is taking too long to complete. Compare the value of this measure across stages/scenarios to know where the maximum delay occurred. This will greatly aid troubleshooting.

4

Monitoring the Hub Transport Servers

Deployed inside your Active Directory directory service forest, the Hub Transport server role handles all mail flow inside the organization, applies transport rules, applies journaling policies, and delivers messages to a recipient's mailbox. Messages that are sent to the Internet are relayed by the Hub Transport server to the Edge Transport server role that is deployed in the perimeter network. Messages that are received from the Internet are processed by the Edge Transport server before they are relayed to the Hub Transport server. If you do not have an Edge Transport server, you can configure the Hub Transport server to relay Internet messages directly. You can also install and configure the Edge Transport server agents on the Hub Transport server to provide anti-spam and antivirus protection inside the organization.

The Hub Transport server role stores all its configuration information in Active Directory. This information includes transport rules settings, journal rule settings, and connector configurations. Because this information is stored in Active Directory, you can configure settings one time, and then those settings are applied by every Hub Transport server in the organization.

The message-processing scenarios that you can manage on the Hub Transport server role are described in the following sections.

- **Internal Mail Flow**

  The Hub Transport server role processes all messages that are sent inside the Exchange 2007/2010 organization before the messages are delivered to a recipient's Inbox or are routed to users outside the organization. There are no exceptions to this behavior; messages are always passed through a server that runs the Hub Transport server role.

- **Messaging Policy and Compliance Features**

  A collection of transport agents lets you configure rules and settings that are applied as messages enter and leave the mail flow components. You can create messaging policy and rule settings that are designed to meet different regulations and that can easily be changed to adapt to your organization's requirements. The transport-based messaging policy and compliance features include server-based rules that you configure to enforce your organization's compliance scenarios and the Journaling agent that acts to enforce message retention.

- **Anti-Spam and Antivirus Protection**

  The Exchange 2007/2010 Built-in Protection features provide anti-spam and antivirus protection for messages. Although these Built-in Protection features are designed for use in the perimeter network on the Edge Transport server role, the Edge Transport agents can also be configured on the Hub Transport server. By default, these agents are not enabled on the Hub Transport server role. To use the anti-spam features on the Hub Transport server, you must register the agents in a configuration file and enable the features that you want to use by running a provided Exchange Management Shell script. You install and enable the antivirus agent in a separate operation.

The error-free functioning of the Hub Transport server is therefore essential to ensure uninterrupted mail flow within the Exchange organization and to insulate the Exchange organization from spam/virus attacks. By continuously monitoring the operations of the Hub Transport server, administrators can be promptly alerted to ineffectiveness of the anti-spam or anti-virus agents on the server and slowdowns in the processing of mail messages by the server.

eG Enterprise provides an *Microsoft Exchange Hub Transport* model that monitors the internal health of the Hub Transport server.

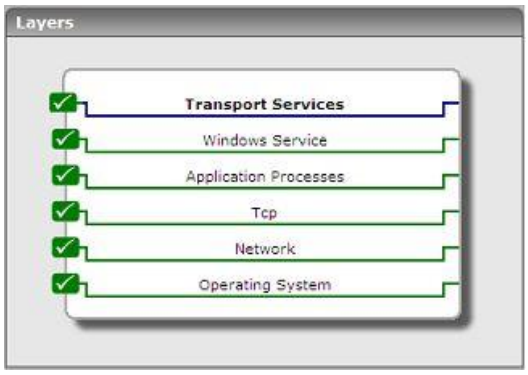

Figure 4.1: Layer model of the Microsoft Exchange Hub Transport server

The tests mapped to each layer report critical statistics that reveal the following:

- Is the server experiencing processing bottlenecks? Are there any lengthy message queues on the server? If so, which ones?
- How effective is the Recipient filter agent? How many requests per second were rejected by the Recipient Lookup and Recipient Block List data sources?
- How successful is the Sender Filter agent in evaluating and filtering out "suspect" senders?
- Is the Sender ID agent efficient?

- Has the Hub Transport server experienced latencies while connecting to the Exchange store? Which store interface can this delay be attributed to?
- How many messages are available in the delivery queue? Is the number very high?
- Do too many messages exist in the retry queue?
- Are too many messages awaiting delivery to an external recipient?
- Have messages been queued in the Unreachable queue?
- Does the poison queue contain messages?
- Which receive connector is overloaded with data and messages?

The sections to come discuss the **Transport Services** layer alone, as all other layers have been discussed elaborately in the *Monitoring Unix and Windows Servers* document.

4.1 The Transport Services Layer

The tests mapped to this layer monitor the transport services offered by the Hub Transport server.

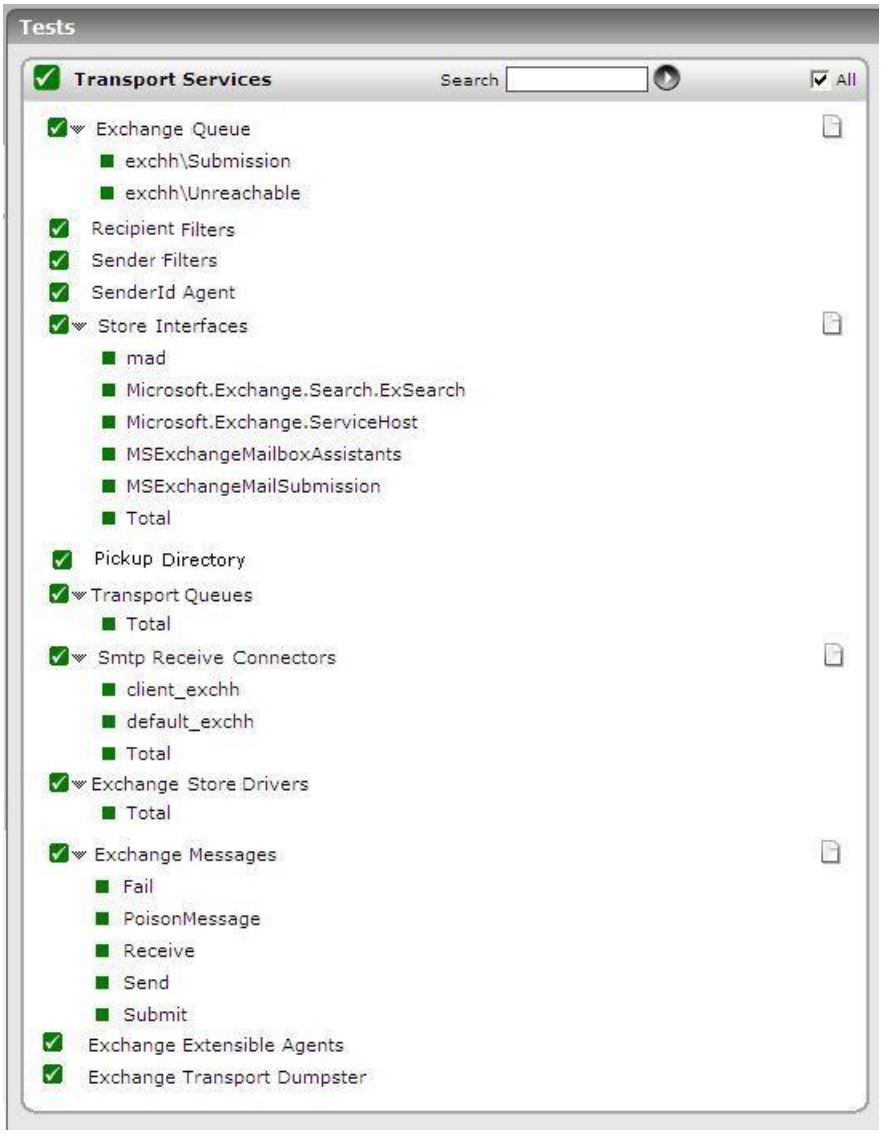

Figure 4.2: The tests mapped to the Transport Services layer

4.1.1 Exchange Queues Test

A *queue* is a temporary holding location for messages that are waiting to enter the next stage of processing. Each queue represents a logical set of messages that an Exchange transport server processes in a specific order. Queues exist only on computers that have the Hub Transport server role or Edge Transport server role installed.

This test reports the length of each message queue on the Microsoft Exchange Edge Transport server or the Microsoft Exchange Hub Transport server, so that queues experiencing processing bottlenecks can be accurately identified.,

Purpose	Reports the length of each message queue on the Microsoft Exchange Edge Transport server or the Microsoft Exchange Hub Transport server
Target of the	A server configured with the Hub/Edge Transport role

<table>
<tr><td>test</td><td></td></tr>
<tr><td>Agent deploying the test</td><td>An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td>1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST - Indicates the IP address of the Edge/Hub Transport server
3. PORT - The port number of the Edge/Hub Transport server. By default, this is 691.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td>One set of results for every queue on the Hub/Edge Transport server being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>
<table>
<tr><th>Measurement</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Messages in queue:
Indicates the number of messages currently found in this queue.</td><td>Number</td><td>A high number could indicate a processing bottleneck in the queue.</td></tr>
</table>
</td></tr>
</table>

4.1.2 Recipient Filters Test

The Recipient Filter agent is an anti-spam agent that is enabled on computers that have the Microsoft Exchange server 2007/2010 Edge Transport server role installed.

The Recipient Filter agent blocks messages according to the characteristics of the intended recipient in the organization. The Recipient Filter agent can help you prevent the acceptance of messages in the following scenarios:

- **Nonexistent recipients**: You can prevent delivery to recipients that are not in the organization's

address book. For example, you may want to stop delivery to frequently misused account names, such as administrator@contoso.com or support@contoso.com.

- **Restricted distribution lists**: You can prevent delivery of Internet mail to distribution lists that should be used only by internal users.

- **Mailboxes that should never receive messages from the Internet**: You can prevent delivery of Internet mail to a specific mailbox or alias that is typically used inside the organization, such as Helpdesk.

The Recipient Filter agent acts on recipients that are stored in one or both of the following data sources:

- **Recipient Block list:** An administrator-defined list of recipients for which inbound messages from the Internet should never be accepted.

- **Recipient Lookup**: Verification that the recipient is in the organization. Recipient Lookup requires access to Active Directory directory service information that is provided by EdgeSync to Active Directory Application Mode (ADAM).

You can use this test to monitor the effectiveness of the Recipient Filter Agent. This test reports the number of messages that were rejected based on the **Recipient Block List** and the **Recipient Lookup** data sources.

Purpose	Monitors the effectiveness of the Recipient Filter Agent by reporting the number of messages that were rejected based on the **Recipient Block List** and the **Recipient Lookup** data sources		
Target of the test	A server configured with the Hub/Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Hub/Edge Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Recipients rejected by recipient validation: Indicates the number of recipients rejected per second by recipient validation.	Rejects/Sec	One benefit of the Recipient Filter agent is the ability to verify that the recipients on an inbound message are in your organization before Exchange 2007/2010 transmits the message into your organization. The ability to verify recipients in your organization relies on a **Recipient Lookup** data source that is available to the Hub/Edge Transport server. The value of this measure indicates the number of recipients who were rejected by this data source.

	Recipients rejected by block list: Indicates the number of recipients rejected by block list per second.	Rejects/Sec	The Recipient Block list is a list that is maintained by the Edge Transport server administrators. The Recipient Block list data is stored in the Edge Transport server instance of ADAM. You must enter blocked recipients on each Edge Transport server computer. You can enter the recipients that you want the Recipient Filter agent to block in the Exchange Management Console on the **Blocked Recipients** tab of the **Recipient Filtering Properties** page. You use the **Set-RecipientFilterConfig** command in the Exchange Management Shell to enter recipients.

4.1.3 Sender Filters Test

The Sender Filter agent is an anti-spam filter that is enabled on computers that have the Microsoft Exchange server 2007/2010 Edge Transport server role installed. The Sender Filter agent relies on the MAIL FROM: Simple Mail Transfer Protocol (SMTP) header to determine what action, if any, to take on an inbound e-mail message.

The Sender Filter agent acts on messages from specific senders outside the organization. Administrators of Edge Transport servers maintain a list of senders who are blocked from sending messages to the organization. As an administrator, you can block single senders (kim@contoso.com), whole domains (*@.contoso.com), or domains and all subdomains (*@*.contoso.com). You can also configure what action the Sender Filter agent should take when a message that has a blocked sender is found.

Using this test, administrators can determine the overall health and effectiveness of the Sender Filter agent's operations.

Purpose	Determine the overall health and effectiveness of the Sender Filter agent's operations		
Target of the test	A server configured with the Hub/Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Hub/Edge Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages evaluated by sender filter: Indicates the number of messages evaluated by the Sender Filter agent per second.	Msgs/Sec	This is a good measure of the agent's processing ability.

	Messages filtered by sender filter: Indicates the number of messages filtered by the Sender Filter agent per second.	Msgs/Sec	

4.1.4 SenderId Agent Test

The Sender ID agent is an anti-spam agent that is enabled on computers that have the Microsoft Exchange server 2007/2010 Edge Transport server role installed. The Sender ID agent relies on the RECEIVED Simple Mail Transfer Protocol (SMTP) header and a query to the sending system's domain name system (DNS) service to determine what action, if any, to take on an inbound message.

Sender ID is intended to combat the impersonation of a sender and a domain, a practice that is frequently called *spoofing*. A *spoofed mail* is an e-mail message that has a sending address that was modified to appear as if it originates from a sender other than the actual sender of the message.

In essence, Sender ID asks a question: "Has this e-mail message been spoofed?" If the answer is "Yes, it has been spoofed," the Sender ID filter rejects or deletes the message immediately. If the answer is "No, we can confirm the sender's authenticity," the message is assigned a Sender ID status and transmitted to Intelligent Message Filter, if Intelligent Message Filter is enabled on the server, for additional anti-spam processing.

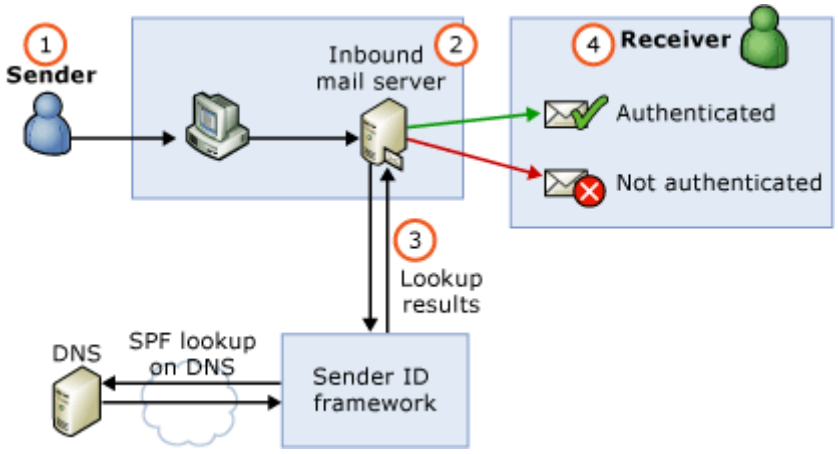

Figure 4.3: How the Sender ID filter works?

Here are the steps in the Sender ID verification process in Figure 4.3:

1. A sender sends an e-mail message to the receiver.
2. The receiver's inbound mail server receives the e-mail message and extracts the PRA.
3. The inbound mail server checks which domain claims to have sent the message, and examines the domain name system (DNS) for the sender policy framework (SPF) record of that domain. These SPF records identify authorized outgoing e-mail servers. The inbound server determines whether the sending e-mail server's IP address matches any of the IP addresses that are published in the SPF record.
4. If the IP addresses match, the e-mail message is authenticated and delivered to the receiver. If the IP

addresses do not match, the e-mail message fails authentication and is not delivered.

5. Based on the evaluation of the Sender ID record, every message is stamped with a Sender ID status. Intelligent Message Filter considers this status for the final assignment of an SCL rating, if Intelligent Message Filter is enabled on the server and the status is also available as an output from the Sender ID filter.

This test reports statistics related to the anti-spamming activities performed by the Sender ID agent, and reveals its overall efficiency.

Purpose	Reports statistics related to the anti-spamming activities performed by the Sender ID agent, and reveals its overall efficiency		
Target of the test	A server configured with the Hub/Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Hub/Edge Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages validated by Sender Id agent: Indicates the number of messages validated per second.	Msgs/Sec	

	Messages with no PRA: Indicates the number of messages per second that were detected as not having a valid PRA.	Msgs/Sec	When you enable Sender ID, each message contains a Sender ID status in the metadata of the message. When an e-mail message is received, the Edge Transport server queries the sender's DNS server to verify that the IP address from which the message was received is authorized to send messages for the domain that is specified in the message headers. The IP address of the authorized sending server is referred to as the purported responsible address (PRA). PRA is calculated based on the following message headers: • Resent-Sender: • Resent-From: • Sender: • From: A high value of this measure indicates that the Sender ID agent has rejected many messages owing to an invalid PRA.
	Messages with SoftFail result: Indicates the number of messages that were validated per second with a SoftFail result.	Msgs/Sec	Anti-spam stamps help you diagnose spam-related problems by applying diagnostic metadata, or "stamps," such as sender-specific information, puzzle validation results, and content filtering results, to messages as they pass through the anti-spam features that filter inbound messages from the Internet. The Sender ID (SID) stamp is based on the sender policy framework (SPF) that authorizes the use of domains in e-mail. The SPF is displayed in the message envelope as **Received-SPF**. The Sender ID evaluation process generates a Sender ID status for the message. If the status returned is **SoftFail** then it means that the IP address of the sender may not be in the SPF. **Softfail** is considered less trusted than **Neutral**, where the sender ID verification check is inconclusive.
	Messages with a fail – non-existent domain - result: Indicates the number of messages that were validated per second with a *Fail – Non-existent Domain* result.	Msgs/Sec	

	Messages with a fail – malformed domain result: Indicates the number of messages per second that were validated with a *Fail – Malformed Domain* result.	Msgs/Sec	
	Messages with a Fail Not Permitted result: Indicates the number of messages per second that were validated with a *Fail – Not Permitted* result.	Msgs/Sec	
	Messages with a None result: Indicates the number of messages per second that were validated with the result of *None*.	Msgs/Sec	The *None* result signifies that no published SPF data exists in the sender's Domain Name System (DNS).
	Messages with a TempError result: Indicates the number of messages per second that were validated with a *TempError* result.	Msgs/Sec	The *TempError* result denotes that a temporary DNS failure occurred, such as an unavailable DNS server.
	Messages with a Neutral result: Indicates the number of messages per second that were validated with a *Neutral* result.	Msgs/Sec	The *TempError* result implies that Sender ID verification check was inconclusive.
	Messages with a Pass result: Indicates the number of messages per second that were validated with a *Pass* result.	Msgs/Sec	A *Pass* result indicates that the IP Address and Purported Responsible Domain pair passed the Sender ID verification check.
	Messages missing originating IP: Indicates the number of messages for which the originating IP could not be determined.	Msgs/Sec	

	Messages with a PermError result: Indicates the number of messages per second that were validated with a *PermError* result.	Validates/Sec	A *PermError* result indicates that the DNS record is invalid, such as an error in the record format.

4.1.5 Store Interfaces Test

The core data storage repository for Microsoft Exchange server 2007/2010 is the Microsoft Exchange Information Store service. This test is useful in isolating and determining issues involving the interfaces between the Microsoft Exchange Information Store service on the Mailbox server and Edge/Hub Transport servers. Unlike Exchange Server 2003, Exchange 2007/2010 communicates with Hub Transport servers via RPC, not Simple Mail Transfer Protocol (SMTP), and therefore latency and queuing are a greater concern. This test isolates and determines issues involving the interface between the Microsoft Exchange Information Store service on the Mailbox server and Hub Transport servers.

Purpose	Isolates and determines issues involving the interface between the Microsoft Exchange Information Store service on the Mailbox server and Hub Transport servers		
Target of the test	A server configured with the Edge/Hub Trasport Server role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for each interface between the Exchange store and the Edge/Hub transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active connections to connection caches: Indicates the number of active connections in all connection caches for this interface.	Number	
	Idle connections to connection caches: Indicates the current number of idle connections in all connection caches for this interface.	Number	Ideally, the value of this measure should be zero, as idle connections in a cache are only resource drainers.

	Current RPC requests outstanding: Indicates the current number of outstanding RPC requests for this interface.	Number	Ideally, the value of this measure should be 0.
	Data transfer over an RPC call - average: Indicates the average number of bytes, sent to the server in one RPC call via this interface.	KB	These measures serve as good indicators of the data load generated by the RPC traffic to and from the Hub transport server.
	Avg. data received per RPC call: Indicates the average number of bytes, received from the server in one succeeded RPC call via this interface.	KB	
	Average RPC latency: Indicates the average latency in milliseconds, averaged across all RPC operations in the past 1024 RPC packets.	Msecs	Average is calculated over all RPCs since exrpc32 was loaded. Ideally, this value should be less than 25 ms. High RPC latencies often cause significant delays in the Mailbox server – Hub Transport server interactions, and hence need to be eliminated.
	Slow RPC requests: Indicates the percentage of slow RPC requests in the RPC queue, currently.	Percent	A slow RPC request is one that has taken more than 2 seconds. Any value higher than 5% for this measure is a cause for concern.
	RPC requests failed: Indicates the percentage of requests in the RPC queue that currently failed.	Percent	Ideally, this value should be 0. A non-zero value for this measure migt warrant an investigation.
	RPC requests succeeded: Indicates the percentage of RPC requests in the queue that currently succeeded.	Percent	

4.1.6 Transport Queues Test

A *queue* is a temporary holding location for messages that are waiting to enter the next stage of processing. Each queue represents a logical set of messages that an Exchange transport server processes in a specific order. Queues exist only on computers that have the Hub Transport server role or Edge Transport server role installed.

Long winding message queues or messages with long waiting times in a queue could indicate lapses in the processing ability of the transport servers. To verify this, it is essential for the queue length to be monitored continuously.

This test monitors the active, remote, and retry queues to figure out whether or not the transport servers are experiencing processing bottlenecks.

Purpose	Monitors the active, remote, and retry queues to figure out whether or not the transport servers are experiencing processing bottlenecks		
Target of the test	A server configured with the Edge Transport or Hub Transport server roles		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for every queue on the Edge/Hub Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages in the active mailbox delivery queue: Indicates the number of messages currently in active mailbox queues.	Number	Mailbox queues hold messages that are being delivered to a Mailbox server that is located in the same site as the Hub Transport server. Mailbox delivery queues exist only on Hub Transport servers. One mailbox delivery queue exists for each destination Mailbox server. If the number of messages in the queue grows continuously, it could indicate a processing bottleneck. You might want to investigate this condition further.

	Messages in the retry mailbox delivery queue: Indicates the number of messages currently in retry in the mailbox queues.	Number	Retry is a renewed connection attempt with the destination domain, smart host, or Mailbox server. Messages in this queue are in a retry state because an issue prevented their delivery. If the issue is transient, a subsequent reattempt to send the message may be successful. A high value for this measure could indicate any of the following: • A domain to which you send a large amount of e-mail is down or experiencing problems. • A computer on your network may be infected with a virus which is sending messages through your Exchange servers. • Your DNS server may have some issue resolving fully qualified domain names (FQDNs) to IP addresses. • There may be a network connectivity issue that is preventing your server from properly connecting to destination servers, or the internet. Some possible issues that could affect your connection are: • Router or routing issues between your server and the destination • Proxy or gateway server issues. • Internet Service providers (ISP) issues, such as a cut line, downed system, routing issues, global disturbance, or some other issue.

			The resolution to a high retry remote delivery queue length depends on the root cause of this problem. Try one or more of the following to identify and resolve the problem causing the high volume of messages in the remote delivery queue.: • Check the destination addresses for the messages in the retry queue. If the messages are all addressed to a single domain or small number of domains, verify that the specified domains are valid and functional. • Verify that there are no machines on your network that are infected with a virus which might be sending messages through your Exchange server(s). Take steps to remove the virus from the infected machine, or remove the machine from your network. • Check where the retry messages are being sent to, if there a large number of messages addressed to companies that you do not know, do not regularly work with, or with unusual subject lines that look to be spam in nature. • Confirm that your DNS server can resolve the FQDNs of the affected domains mail exchanger (MX) resource records to IP by using the NSLOOKUP command. Confirm that there are no network connectivity issues preventing your server from properly connecting to destination servers or the Internet.

	Messages in the active remote delivery queue: Indicates the number of messages currently in active remote delivery queues.	Number	Remote delivery queues hold messages that are being delivered to a remote domain or smart host by using the Simple Mail Transfer Protocol (SMTP). After all messages are delivered, these queues persist for three minutes and then are automatically deleted. A high value of this measure could indicate a processing bottleneck or a poor network link between the Hub Transport server and the remote domain to which messages are to be delivered.

	Messages in the retry remote delivery queue: Indicates the number of messages currently in the retry remote delivery queues.	Number	Retry is a renewed connection attempt with the destination domain, smart host, or Mailbox server. Messages in this queue are in a retry state because an issue prevented their delivery. If the issue is transient, a subsequent reattempt to send the message may be successful. A high value for this measure could indicate any of the following: • A domain to which you send a large amount of e-mail is down or experiencing problems. • A computer on your network may be infected with a virus which is sending messages through your Exchange servers. • Your DNS server may have some issue resolving fully qualified domain names (FQDNs) to IP addresses. • There may be a network connectivity issue that is preventing your server from properly connecting to destination servers, or the internet. Some possible issues that could effect your connection are: • Router or routing issues between your server and the destination • Proxy or gateway server issues. • Internet Service providers (ISP) issues, such as a cut line, downed system, routing issues, global disturbance, or some other issue.

			The resolution to a high retry remote delivery queue length depends on the root cause of this problem. Try one or more of the following to identify and resolve the problem causing the high volume of messages in the remote delivery queue.: • Check the destination addresses for the messages in the retry queue. If the messages are all addressed to a single domain or small number of domains, verify that the specified domains are valid and functional. • Verify that there are no machines on your network that are infected with a virus which might be sending messages through your Exchange server(s). Take steps to remove the virus from the infected machine, or remove the machine from your network. • Check where the retry messages are being sent to, if there a large number of messages addressed to companies that you do not know, do not regularly work with, or with unusual subject lines that look to be spam in nature. • Confirm that your DNS server can resolve the FQDNs of the affected domains mail exchanger (MX) resource records to IP by using the NSLOOKUP command. • Confirm that there are no network connectivity issues preventing your server from properly connecting to destination servers or the Internet.
	Messages in the active Non-SMTP delivery queue: Indicates the number of messages currently in the Drop directory that is used by a Foreign connector.	Number	This refers to the number of messages in the queue for which the **Delivery Type** has been set to **NonSmtpGatewayDelivery**. The messages in such a queue are typically queued for delivery to an external recipient by using a non-SMTP connector on the local server.

	Messages in the retry Non-SMTP delivery queue: Indicates the number of messages currently in retry in the non-SMTP gateway delivery queues.	Number	Messages in this queue are in a retry state because an issue prevented their delivery. If the issue is transient, a subsequent reattempt to send the message may be successful. The value of this measure could rise, owing to the following reasons: • A connector that connects to the Non-SMTP mail server might not be functioning properly. • A domain that you connect to via a Non-SMTP connector might be down or unreachable. • Your DNS server may have some issue resolving fully qualified domain names (FQDNs) to IP addresses. • There may be a network connectivity issue that is preventing your server from properly connecting to destination servers or the Internet. Some possible issues that could affect your connection are: • Router or routing issues between your server and the destination • Proxy or gateway server issues. • Internet Service providers (ISP) issues, such as a cut line, downed system, routing issues, global disturbance, or some other issue.
	Messages in the aggregate delivery queue: Indicates the number of messages currently queued for delivery in all queues.	Number	

	Largest delivery queue length: Indicates the number of messages that are currently queued to a given Exchange Hub Transport server or Edge Transport server.	Number	When this value is high, the server cannot establish a SMTP session to the other Hub Transport or Edge Transport server. Other symptoms you may experience when this threshold is reached are reduced intra-site, inter-site, and external mail flow. This alert may be caused by one or more of the following conditions: • Problem with a specific Hub Transport server or Edge Transport server. For example, one or more required services may not be running. • Issues with network connectivity, routers, or firewalls.
	Messages in the unreachable queue: Indicates the number of messages currently in the Unreachable queue.	Number	The categorizer sends messages to the unreachable queue when there is no known route to their destinations. Typically, an unreachable destination is caused by a configuration error that affects the delivery path By default, the messages in the unreachable queue have the status of Ready. Messages in the unreachable queue are never automatically resubmitted. Messages remain in the unreachable queue until they are manually resubmitted by an administrator, removed by an administrator, or the value specified in the MessageExpirationTimeOut parameter passes.
	Messages in the poison queue: Indicates the number of messages currently in the poison queue.	Number	The poison message queue contains messages that are determined to be potentially harmful to the Microsoft Exchange server 2007/2010 server after a server failure. The messages may be genuinely harmful in their content and format. Alternatively, they may be the results of a poorly written agent that has caused the Exchange server to fail when it processed the supposedly bad messages. Messages remain in the poison message queue until they are manually resumed or removed by an administrator. The messages in the poison message queue are never automatically resumed or expired.

	Messages in the submission queue: Indicates the number of messages currently in the submission queue.	Number	A sustained high Submission Queue Length value may indicate that an excessive amount of inbound messages have over-loaded the categorizer. It may also indicate that there is an issue with message categorization. Message resubmission sends undelivered messages back to the submission queue to be processed again by the categorizer. A sustained high Submission Queue Length may be caused by one or more of the following: • The server is being over-utilized and does not have enough resources to satisfy all of the current requests. This situation may occur if there are more messages being submitted for transport then the server can handle. Similarly, it may also occur if many messages are being resubmitted for categorization. • There is a problem with a custom event sink or rule, or a third-party event sink or rule.

4.1.7 SMTP Receive Connectors Test

Smtp Receive connectors are configured on computers that are running Microsoft Exchange server 2007/2010 and that have Hub Transport and Edge Transport server roles installed. The Smtp Receive Connector represents a logical gateway through which inbound messages are received. This test monitors the statistics of smtp receive connectors.

Using this test, you can periodically observe the traffic conducted by the Receive connectors, and be promptly alerted when the connector rejects messages.

Purpose	Periodically observe the traffic conducted by the Receive connectors, and be promptly alerted when the connector rejects messages
Target of the test	A server configured with the Hub Transport role
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Hub Transport server 3. **PORT** - The port number of the Hub Transport server. By default, this is 691.
Outputs of the test	One set of results for each receive connector on the Hub Transport server being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Current SMTP connections to the server: Indicates the current number of outbound connections from the SMTP Receive connectors.	Number	
	Messages received by the server: Indicates the number of messages received by the SMTP Receive connector each second.	Msgs / Sec	This is a good indicator of the load on the Receive connector.
	Data received by SMTP server: Indicates the number of bytes received per second.	Bytes/Sec	This is a good indicator of the load on the Receive connector.
	Messages refused due to size limit: Indicates the number of messages that were rejected currently because they were too big.	Number	Ideally, this value should be 0.
	Avg. size of messages: The average number of message bytes per inbound message received.	Bytes / Msg	
	Avg. recipients per message: Indicates the average recipients per message handled by this SMTP Receive connector.	Recipients/Msg	
	Avg. data transfer per connection: Indicates the average number of bytes received per connection.	Bytes/Conn	

	Avg. messages per connection: Indicates the average number of message bytes per inbound message received.	Msgs/Conn	

4.1.8 Exchange Store Drivers Test

The Store driver on the Hub Transport server places messages from the transport pipeline into the appropriate mailbox. The Store driver on the Hub Transport server also adds messages from the Outbox of a sender on the Mailbox server to the transport pipeline.

This test monitors the overall health of each of the Store Drivers.

Purpose	Monitors the overall health of the Store Driver		
Target of the test	A server configured with the Edge/Hub Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge/Hub Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Inbound data delivered: Indicates the number of requests processed from clients during the last measurement period.	KB	
	Inbound failed recipients: Indicates the number of failed deliveries during the last measurement period.	Number	Ideally, this value should be 0. A non-zero value requires further investigation.

	Inbound succeeded recipients: Indicates the number of successful deliveries during the last measurement period.	Number	
	Inbound local delivery calls: Indicates the number of local delivery attempts per second during the last measurement period.	Calls/Sec	
	Inbound message delivery attempts: Indicates the number of attempts for delivering messages per second during the last measurement period.	Attempts/Sec	
	Inbound delivering threads: Indicates the number of threads used in delivery currently.	Number	
	Inbound recipients delivered: Indicates the number of recipients to whom messages were delivered per second.	Recipients/Sec	
	Outbound submitted mail items: Indicates the number of mail messages per second being submitted for delivery.	Mails/Sec	

4.1.9 Pickup Directory Test

By default, the pickup directory exists on every Microsoft Exchange server 2007/2010 computer that has the Hub transport server role or the Edge Transport server role installed. Correctly formatted e-mail message files that you copy to the Pickup directory are submitted for delivery. The Pickup directory is used by administrators for mail flow testing or by applications that must create and submit their own messages. This test monitors the performance of the Pickup directory and reveals whether or not it has been able to submit all messages it contains for delivery.

Purpose	Monitors the performance of the Pickup directory and reveals whether or not it has been able to submit all messages it contains for delivery		
Target of the test	A server configured with the Edge Transport or Hub Transport server roles		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge/Hub Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages submitted to the pickup directory: Indicates the number of messages that were successfully submitted for delivery by the Pickup directory during the last measurement period.	Number	

	Messages to the pickup directory that caused NDR creations: Indicates the number of messages processed by the Pickup directory that caused NDRs to be created during the last measurement period.	Number	A correctly-formatted message file together with a valid sender that can't be successfully submitted for delivery by the Pickup directory generates a non-delivery report (NDR). Malformed content or Pickup directory message restriction violations could also cause the Pickup directory to generate an NDR. When an NDR is generated during Pickup directory message processing, the original message file is attached to the NDR message, and the message file is deleted from the Pickup directory. A correctly formatted message that is submitted by the Pickup directory may later experience a delivery failure and be returned to the sender with an NDR. This kind of failure may be caused by transmission issues that are unrelated to the Pickup directory, such as messaging server failures or routing failures along the delivery path of the message.
	Badmailed messages to the pickup directory: Indicates the number of messages that were submitted to the Pickup directory but were classified as badmail and not delivered, during the last measurement period.	Number	A message that is classified as badmail has serious problems that prevent the Pickup directory from submitting the message for delivery. The other condition that causes badmail is when the message is formatted correctly, but the recipients are not valid, and an NDR message can't be sent to the sender because the sender is not valid. Message files that are determined to be badmail are left in the Pickup directory and are renamed from *<filename>*.eml to *<filename>*.bad. If the *<filename>*.bad file already exists, the file is renamed to *<filename><datetime>*.bad. If badmail exists in the Pickup directory, an event log error is generated, but the same badmail messages do not generate repeated event log errors.

4.1.10 Exchange Messages Test

This test tracks the flow of messages through an Exchange 2007/2010 organization, and reports the number and size of messages that pertain to every key event type handled by the Exchange 2007/2010 server. These types include the following:

Type	Description
SEND	A message sent by Simple Mail Transfer Protocol (SMTP) to a different server.
RECEIVE	A message received and committed to the database.

SUBMIT	A message submitted by an Exchange 2007/2010 computer that has the Mailbox server role installed to an Exchange 2007/2010 computer that has the Hub Transport server role or Edge Transport server role installed.
POISON	A message added to the poison message queue or removed from the poison message queue.
FAIL	Message delivery failed

Whenever a user complaints of not being able to send or receive mails, the metrics reported by this test and the detailed diagnosis information provided therein will enable administrators to accurately determine the current status of the email sent by the user.

If need be, administrators can configure this test to additionally report the total number of messages on the Exchange 2007/2010 server and their total size, regardless of event type. Apart from the event types discussed above, this total will also include messages that belong to the following event types:

Type	Description
BADMAIL	A message submitted by the Pickup directory or the Replay directory that cannot be delivered or returned
DELIVER	A message delivered to a mailbox
DEFER	A message for which delivery was delayed
DSN	A message for which a delivery status notification (DSN) was generated
EXPAND	A distribution group was expanded
FAIL	Message delivery failed
REDIRECT	A message redirected to an alternative recipient after an Active Directory directory service lookup
RESOLVE	A message for which recipients were resolved to a different e-mail address after an Active Directory lookup
TRANSFER	Recipients were moved to a forked message because of content conversion, message recipient limits, or agents

Exchange administrators can use this total to accurately assess the overall message traffic on the server and the ability of the server to handle the inflow/outflow of messages.

Purpose	Tracks the flow of messages through an Exchange 2007/2010 organization, and reports the number and size of messages that pertain to every key event type handled by the Exchange 2007/2010 server
Target of the test	A server configured with the Mailbox role
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – Indicates the IP address of the Mailbox server. 3. **PORT** – The port number through which the Mailbox server communicates. By default, this is 691. 4. **XCHGEXTENSIONSHELLPATH -** The Exchange Management Shell is a command-line management interface, built on Windows PowerShell, which enables you to administer every part of Microsoft Exchange Server. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the **XCHGEXTENSIONSHELLPATH** is set to *none* by default. 5. **ALLEVENTS** – By default, this flag is set to **false**, indicating that this test will report metrics for only the following event types by default: SEND, RECEIVE, SUBMIT, FAIL, POISON. If you want the test to additionally report metrics across all event types – i.e., support an additional *All* descriptor, which will report the total number of emails handled by the server and their total size – then, set this flag to **true**. 6. **DD FOR RECEIVEMESSAGE** – In large, highly active Exchange environments, hundreds of emails may be received by the Exchange server within a short period of time. In such environments, the frequent collection of detailed diagnosis of the received emails may increase the processing overheads of the eG agent, and may even choke the eG database. To avoid this, the **DD FOR RECEIVEMESSAGE** flag is set to **No** by default; this implies that the test will not provide the detailed diagnosis for the **RECEIVE** descriptor – i.e., for the received messages – by default. To view detailed diagnosis for these messages as well, set this flag to **Yes**. 7. **DD FOR SENDMESSAGE** - In large, highly active Exchange environments, hundreds of emails may be sent by the Exchange server within a short period of time. In such environments, the frequent collection of detailed diagnosis information related to the sent emails may increase the processing overheads of the eG agent, and may even choke the eG database. To avoid this, the **DD FOR SENDMESSAGE** flag is set to **No** by default; this implies that the test will not provide the detailed diagnosis for the SEND descriptor – i.e., for the sent messages – by default. To view detailed diagnosis for these messages as well, set this flag to **Yes**. 8. **DD FOR SUBMITMESSAGE** - In large, highly active Exchange environments, hundreds of emails may be submitted to the transport pipeline within a short period of time. In such environments, the frequent collection of detailed diagnosis information related to the submitted emails may increase the processing overheads of the eG agent, and may even choke the eG database. To avoid this, the **DD FOR SUBMITMESSAGE** flag is set to **No** by default; this implies that the test will not provide the detailed diagnosis for the SUBMIT descriptor – i.e., for the sent messages – by default. To view detailed diagnosis for these messages as well, set this flag to **Yes**. 9. **ISPASSIVE** – If the value chosen is **YES**, then the Exchange server under consideration is a passive server in an Exchange cluster. No alerts will be generated if the server is not running. Measures will be reported as "Not applicable" by the agent if the server is not up.		
Outputs of the test	One set of results for each of the following event types: SEND, RECEIVE, FAIL, POISON, SUBMIT		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Number of emails:** Indicates the number of emails of this event type detected during this measurement period.	Number	By default, this measure provides detailed diagnosis for the FAIL and POISON messages only. Using the detailed diagnosis of these descriptors, you can view the complete details of the failed and poison messages. Optionally, users can turn on detailed diagnosis generation for the RECEIVE, SEND, and SUBMIT messages as well, so as to view the complete details of such messages. The *All* descriptor, even if displayed, will not provide detailed diagnosis information.
	Total traffic: Indicates the total size of messages of this event type, during the last measurement period.	Number	Since the value of this measure includes the size of attachments, an unusually high value could indicate that one/more messages carry large attachments. A high value could also indicate the availability of a large number of messages of a particular type.

4.1.11 Exchange Extensible Agents Test

Transport agents let you install custom software, created by Microsoft, by third-party vendors, or by your organization, on a computer that is running Microsoft Exchange server 2007/2010. This software can then process e-mail messages that pass through the transport pipeline on a Hub Transport server or Edge Transport server. Custom transport agents provide additional functionality to Exchange 2007/2010, such as anti-spam or antivirus programs or any transport function that your organization may require.

Delays in e-mail processing by the transport agents and flaws while performing anti-spam /anti-virus activities may affect the stability and security of Exchange. By periodically checking how the agents process e-mail messages, you can easily spot processing bottlenecks and security lapses. This test does just that.

Purpose	Periodically monitors how the agents process e-mail messages, so that you can easily spot processing bottlenecks and security lapses		
Target of the test	A server configured with the Edge Transport or Hub Transport server roles		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge/Hub Transport server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Agent processing time:** Indicates the time taken by the transport agent per event for processing e-mail messages.	Secs	Ideally, the value of this measure should be low. A high value indicates that the agent is taking too long a time to process e-mail messages. The reason for this needs to be investigated. Sustained higher latencies may indicate a hung agent.
	Total agent invocations: Indicates the total number of agent invocatins since last restart.	Number	

4.1.12 Exchange Transport Dumpster Test

The **transport dumpster** is a feature of the Hub Transport server role that submits recently delivered mail after an unscheduled outage. The transport dumpster should always be turned on when using CCR or local continuous replication (LCR). The transport dumpster is enabled organization wide by setting the amount of storage available per storage group and setting the time to retain mail in the transport dumpster. If either of these settings are violated in real-time, then the transport dumpster starts deleting mails.

Using this test, you can closely monitor the transport dumpster and the mails within, determine how often the Exchange server experienced outages, and also figure out whether any critical mails were lost during downtime.

Purpose	Closely monitor the transport dumpster and the mails within, determine how often the Exchange server experienced outages, and also figure out whether any critical mails were lost during downtime		
Target of the test	A server configured with the Edge Transport or Hub Transport server roles		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge/Hub Transport server 3. **PORT** - The port number of the Edge/Hub Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge/Hub Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Overall size of the Dumpster: Indicates the total size of the mail items(in bytes) that are currently in the transport dumpster on this server.	Bytes	A low value is typically desired for this measure. A very high value or a value that increases consistently could reveal a prolonged outage. If the value exceeds or is dangerously close to the value of the *MaxDumpsterSizePerStorageGroup* parameter, it could mean that mails will soon be deleted from the dumpster.
	Dumpster inserts rate: Indicates the rate at which mails were inserted in the dumpster.	Inserts/Sec	
	Number of current items: Indicates the total number of mail items that are currently available in the transport dumpster on this server.	Number	A low value is typically desired for this measure. A very high value or a value that increases consistently could reveal a prolonged outage.

	Dumpster deletions rate: Indicates the rate at which items were deleted from the dumpster.	Deletes/Sec	A high rate of deletion is typically indicative of the frequent violation of the *MaxDumpsterSizePerStorageGroup* and/or the *MaxDumpsterTime* setting. These parameters have been discussed below: • **MaxDumpsterSizePerStorageGroup:** This parameter specifies the maximum size of the transport dumpster queue for each storage group. It is recommended that you set this to a size that is 1.5 times the size of the maximum message that can be sent in the organization. If the organization has no size limits, we recommend you configure the *MaxDumpsterSizePerStorageGroup* parameter to a size that is 1.5 times the size of the average message size sent in the organization. For example, if the maximum size for messages is 10 megabytes (MB), you should configure the *MaxDumpsterSizePerStorageGroup* parameter with a value of 15 MB. • **MaxDumpsterTime:** This parameter specifies how long an e-mail message should remain in the transport dumpster queue, to a value of 07.00:00:00, which is 7 days. This amount of time is sufficient to allow for an extended outage to occur without loss of e-mail.

			When using the transport dumpster feature, additional disk space is needed on the Hub Transport server to host the transport dumpster queues. The amount of storage space required is approximately equal to the value of *MaxDumpsterSizePerStorageGroup* multiplied by the number of storage groups. If you do not configure the transport dumpster, the default values are used. The default value for the *MaxDumpsterSizePerStorageGroup* parameter is 18 MB and the default value for the *MaxDumpsterTime* parameter is 7 days. If either the size limit or time limit is reached, messages are removed from the transport dumpster queue by order of first in, first out.

4.1.13 Exchange Email Traffic Test

Periodic workload monitoring is imperative to evaluate the processing ability of the Exchange 2010 server and to proactively detect potential overload conditions. By continuously monitoring the email traffic to and from the Exchange server, this test turns a spotlight on the workload of the Exchange server, helps detect overload conditions, and also points you to the source of the overload – mails sent/received by users in the intranet? Or mail traffic over the internet?

Purpose	Turns a spotlight on the workload of the Exchange server, helps detect overload conditions, and also points you to the source of the overload – mails sent/received by users in the intranet? Or mail traffic over the internet?
Target of the test	A server configured with the Hub Transport role
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - Indicates how often this test needs to be executed.
2. HOST – Indicates the IP address of the Mailbox server.
3. PORT – The port number through which the Mailbox server communicates. By default, this is 691.
4. XCHGEXTENSIONSHELLPATH - The Exchange Management Shell is a command-line management interface, built on Windows PowerShell which enables you to administer every part of Microsoft Exchange. This test uses the Exchange management shell to run scripts and collect the desired performance metrics from the Exchange server. By default, the test auto-discovers the location of the Exchange management shell and thus, automatically loads the Exchange management shell snap-in (exshell.psc1) for script execution. This is why, the XCHGEXTENSIONSHELLPATH is set to none by default.
5. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
6. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
7. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Mailbox server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Internal mails received:
Indicates the number of mails received by the Exchange server from the intranet.</td><td>Number</td><td>In the event of an overload, you can compare the value of this measure with the value of the Internal mails sent, External mails received, and External mails sent measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet?
To know the email IDs that received the emails, the number of emails that each ID received, and the total size of the emails to an ID, use the detailed diagnosis of this test.</td></tr>
</table>

	Internal mails received size: Indicates the total size of the mails received by the Exchange server from the intranet.	KB	
	Internal mails sent: Indicates the number of mails sent by the Exchange server to the intranet.	Number	In the event of an overload, you can compare the value of this measure with the value of the *Internal mails received, External mails received*, and *External mails sent* measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet? To know the email IDs that sent the emails, the number of emails that each ID sent, and the total size of the emails from an ID, use the detailed diagnosis of this test.
	Internal mail sent size: Indicates the total size of the mails sent by the Exchange server to the intranet.	KB	
	External mails received: Indicates the number of mails received by the Exchange server from the internet.	Number	In the event of an overload, you can compare the value of this measure with the value of the *Internal mails received*, *Internal mails sent*, and *External mails sent* measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet? To know the email IDs that received the emails, the number of emails that each ID received, and the total size of the emails to an ID, use the detailed diagnosis of this test.
	External mail received size: Indicates the total size of the mails received by the Exchange server from the internet.	KB	

	External mails sent: Indicates the number of mails sent by the Exchange server to the internet.	Number	In the event of an overload, you can compare the value of this measure with the value of the *Internal mails sent*, *Internal mails sent*, and *External mails received* measures to determine what could have contributed to the overload – is it because of incoming or outgoing mail traffic? Is it because of the exchange of mails over the intranet or the internet? To know the email IDs that sent the emails, the number of emails that each ID sent, and the total size of the emails from an ID, use the detailed diagnosis of this test.
	External mail sent size: Indicates the total size of mails sent by the Exchange server to the internet.	KB	

Use the detailed diagnosis of the *Internal mails received* measure to know the internal email IDs that received the emails, the number of emails that each ID received, and the total size of the emails received by an ID. This way, you can quickly identify the email ID that received the maximum number of emails and that which received mails of the maximum size.

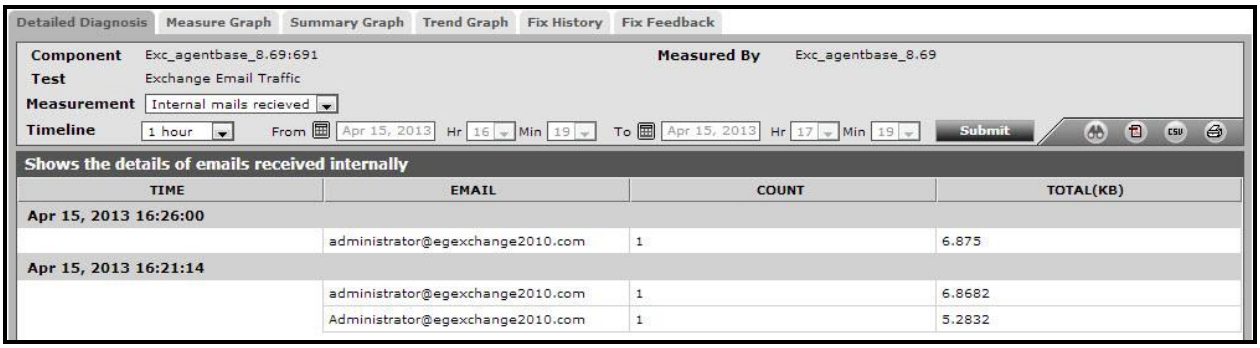

Figure 4.4: The detailed diagnosis of the Internal mails received measure

Use the detailed diagnosis of the *Internal mails sent* measure to know the internal email IDs that sent the emails, the number of emails that were sent from each ID, and the total size of the emails sent from an ID. This way, you can quickly identify the email ID that sent the maximum number of emails and that which sent mails of the maximum size.

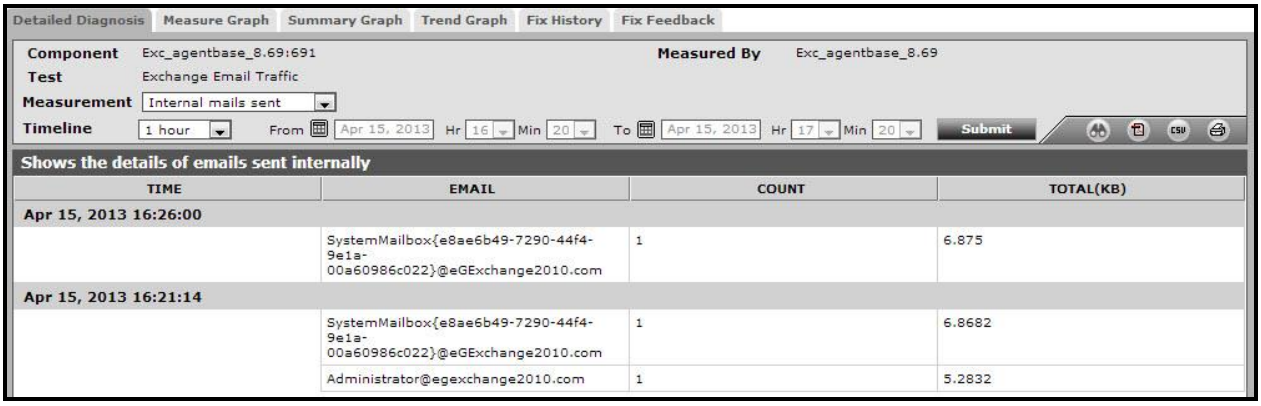

Figure 4.5: The detailed diagnosis of the Internal mails sent measure

5

Monitoring the Edge Transport Servers

In Exchange 2007/2010, the Edge Transport server role is deployed in your organization's perimeter network as a stand-alone server or as a member server of a perimeter-based Active Directory domain. Designed to minimize the attack surface, the Edge Transport server handles all Internet-facing mail flow, which provides Simple Mail Transfer Protocol (SMTP) relay and smart host services for the Exchange organization. Additional layers of message protection and security are provided by a series of agents that run on the Edge Transport server and act on messages as they are processed by the message transport components. These agents support the features that provide protection against viruses and spam and apply transport rules to control message flow.

The message-processing scenarios that you can manage on the Edge Transport server role are described in the following sections.

- **Internet Mail Flow**

  Servers that run the Edge Transport server role accept messages that come into the Exchange 2007/2010 organization from the Internet. After the messages are processed by the Edge Transport server, they are routed to Hub Transport servers inside the organization. All messages that are sent to the Internet from the organization are routed to Edge Transport servers after the messages are processed by the Hub Transport server.

- **Anti-Spam and Antivirus Protection**

  In Exchange 2007/2010, the anti-spam and antivirus features provide services to block viruses and spam, or unsolicited commercial e-mail, at the network perimeter. Most viruses use spam-like tactics to gain access to your organization and to entice users to open an e-mail message. If you can filter out most of your spam, you are also more likely to capture viruses before they enter your organization.

  Spammers use a variety of techniques to send spam into your organization. Servers that run the Edge Transport server role help prevent users in your organization from receiving spam by providing a collection of agents that work together to provide different layers of spam filtering and protection.

- **Edge Transport Rules**

  Edge Transport rules are used to control the flow of messages that are sent to or received from the Internet. The Edge Transport rules help protect corporate network resources and data by applying an action to messages that meet specified conditions. These rules are configured for each server. Edge Transport rule conditions are based on data, such as specific words or text patterns in the message subject, body, header, or From address, the spam confidence level (SCL), or attachment type. Actions determine how the message is processed when a specified condition is true. Possible actions include quarantine of a message, dropping or rejecting a message, appending additional recipients, or logging an event. Optional exceptions exempt particular messages from having an action applied.

- **Address Rewriting**

  You use address rewriting to present a consistent appearance to external recipients of messages from your Exchange 2007/2010 organization. You configure the Address Rewriting agent on the Edge Transport server role to enable the modification of the SMTP addresses on inbound and outbound messages.

If any of these critical services were to fail – for instance, say the Edge Transport server processes internet messages very slowly – it can cause significant delays in the delivery of important mails to specified recipients. In the world of business, such slip-ups are inexcusable, as prompt and effective email correspondence is essential to win orders and earn customer goodwill. Therefore, to prevent such adversities and their impact on corporate revenues, the Edge transport server will have to be monitored 24 x 7, and problems in its operations should be reported to administrators proactively.

eG Enterprise offers a specialized *Microsoft Exchange Edge Transport* model that provides real-time insights into the performance of Edge Transport servers.

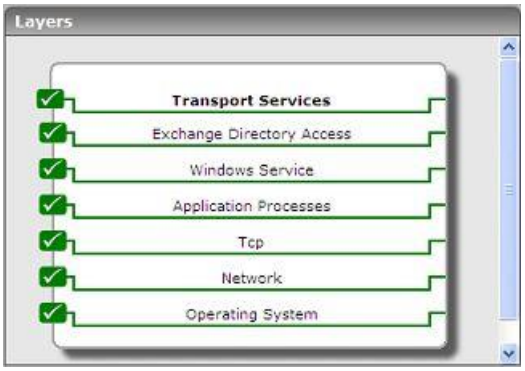

Figure 5.1: Layer model of the Microsoft Exchange Edge Transport server

Every layer of Figure 5.1 reports a wide variety of statistics that enable administrators to quickly find answers to the following critical performance queries:

- Is the Active Directory cache adequately sized to handle requests from the Edge transport server?
- search requests to any domain controller fail owing to a bad network link or the non-availability of the domain controller?
- Were any LDAP fatal errors experienced while communicating with a domain controller?
- Did too many bind calls to any domain controller fail?

- Is any domain controller responding too slowly to read and search requests?
- Is the server experiencing processing bottlenecks? Are there any lengthy message queues on the server? If so, which ones?
- How effective is the Recipient filter agent? How many requests per second were rejected by the Recipient Lookup and Recipient Block List data sources?
- How successful is the Sender Filter agent in evaluating and filtering out "suspect" senders?
- Is the Sender ID agent efficient?
- Has the Edge Transport server experienced latencies while connecting to the Exchange store? Which store interface can this delay be attributed to?
- How many messages are available in the delivery queue? Is the number very high?
- Do too many messages exist in the retry queue?
- Are too many messages awaiting delivery to an external recipient?
- Have messages been queued in the Unreachable queue?
- Does the poison queue contain messages?
- How efficient is Connection Filter agent? How many connection requests were rejected by the agent? Which data store used by the agent rejected the maximum requests - the IP Block list providers, IP allow list providers, or the IP Allow/Block lists defined by administrators?
- Were any spams detected by the content filter agent?
- Were any local/remote senders blocked by the Protocol Analysis agent?

The sections to come discuss the top 2 layers of Figure 5.1, as the remaining layers have already been dealt with in the *Monitoring Unix and Windows Servers* document.

5.1 The Exchange Directory Access Layer

The computer that has the Edge Transport server role installed does not have access to the Active Directory directory service. All configuration and recipient information is stored in the Active Directory Application Mode (ADAM) directory service. To perform recipient lookup tasks, the Edge Transport server requires data that resides in Active Directory. EdgeSync is a collection of processes that are run on a computer that has the Hub Transport server role installed to establish one-way replication of recipient and configuration information from Active Directory to the ADAM instance on an Edge Transport server. The Microsoft Exchange EdgeSync service copies only the information that is required for the Edge Transport server to perform anti-spam configuration tasks and the information about the connector configuration that is required to enable end-to-end mail flow. The Microsoft Exchange EdgeSync service performs scheduled updates so that the information in ADAM remains current.

The tests mapped to this layer measure the health of the interactions between the Active Directory and the Edge Transport server.

Figure 5.2: The tests mapped to the Exchange Directory Access layer

Both these tests have been dealt with elaborately in Chapter 2 of this document. Therefore, let us proceed to discuss the topmost layer – the **Transport Services** layer.

5.2 The Transport Services Layer

The tests linked to this layer monitor the effectiveness of the critical services offered by the Edge Transport server.

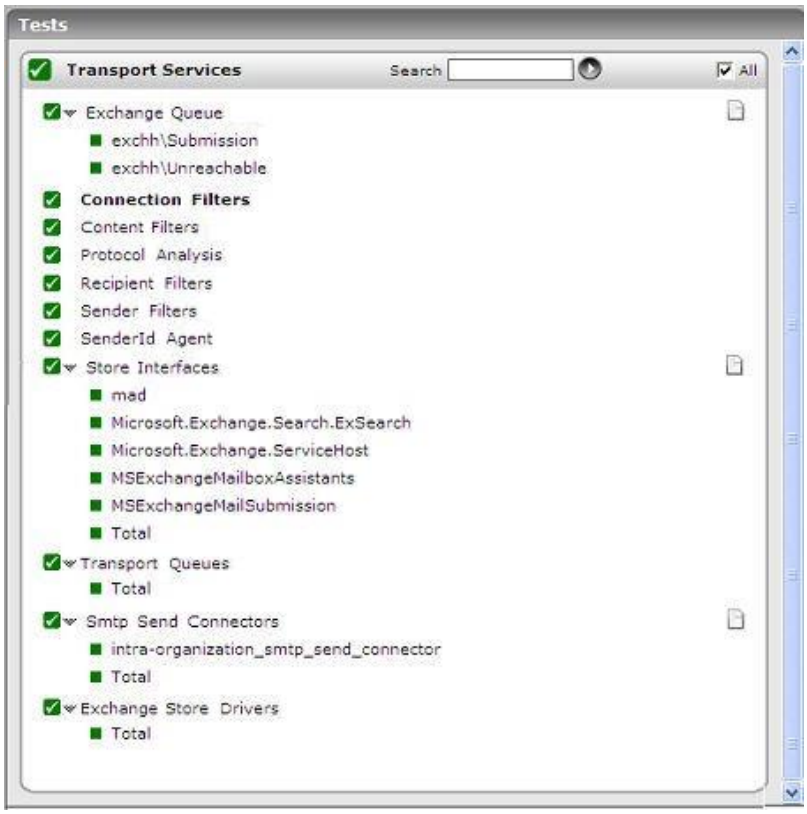

Figure 5.3: The tests mapped to the Transport Services layer

Since most of these tests are common to both the Edge Transport and Hub Transport servers, Chapter 4 of this document provides the details of the common tests. The sections to come therefore, discuss the tests that are specific to the Edge Transport server alone.

5.2.1 Connection Filters Test

The Connection Filter agent is an anti-spam agent that is enabled on computers that have the Microsoft Exchange server 2007/2010 Edge Transport server role installed. The Connection Filter agent relies on the IP address of the remote server that is trying to connect to determine what action, if any, to take on an inbound message. The remote IP address is available to the Connection Filter agent as a by-product of the underlying TCP/IP connection that is required for the Simple Mail Transfer Protocol (SMTP) session

When you enable the Connection Filter agent, the Connection Filter agent is the first anti-spam agent to run when an inbound message is evaluated. When an inbound message is submitted to an Edge Transport server on which the Connection Filter agent is enabled, the source IP address of the SMTP connection is checked against any of the following data stores of IP addresses:

- Administrator-defined IP Allow lists and IP Block lists
- IP Block List providers
- IP Allow List providers

You must configure at least one of these data stores of IP addresses for the Connection Filter agent to be operational.

The source P address is first compared to the administrator-defined IP Allow list and IP Block list. If the IP address does not exist on either the administrator-defined IP Allow list or IP Block list, the Connection Filter agent queries the IP Block List provider services according to the priority rating that is assigned to each provider. If the IP address appears on the IP Block list of an IP Block List provider, the Edge Transport server waits for and parses the RCPT TO header, responds to the sending system with an SMTP 550 error, and closes the connection. If the IP address does not appear on the IP Block lists of any one of the IP Block List providers, the next agent in the anti-spam chain processes the connection.

This test monitors the connection filtering agent's activities to reveal the number of connection requests/inbound messages that are in various stages of filtering.

Purpose	Monitors the connection filtering agent's activities to reveal the number of connection requests/inbound messages that are in various stages of filtering		
Target of the test	A server configured with the Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge Transport server 3. **PORT** - The port number of the Edge Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Connections to IP block list providers: Indicates the number of connections to the IP Block List providers during the last measurement period.	Number	IP Block List provider services compile lists of IP addresses from which spam has originated in the past. Additionally, some IP Block List providers provide lists of IP addresses for which SMTP is configured for open relay. There are also IP Block List provider services that provide lists of IP addresses that support dial-up access. You can configure multiple IP Block List provider configurations by using the Exchange Management Console or the Exchange Management Shell. When you configure the Connection Filter agent to use an IP Block List provider, the Connection Filter agent queries the IP Block List provider service to determine whether a match exists with the connecting IP addresses before the message is accepted into the organization. The value of this measure indicates the number of connections that the filtering agent has established with the IP Block List provider service to perform such queries.
			When you use the Connection Filter agent, it is a best practice to use one or more IP Block List providers to manage access into your organization. However, there may be some disadvantages to using an IP Block List provider. Because the Connection Filter agent must query an external entity for each unknown IP address, outages or delays at the IP Block List provider service can cause delays in the processing of messages on the Edge Transport server. In extreme cases, such outages or delays could cause a mail-flow bottleneck on the Edge Transport server. The other disadvantage of using an external IP Block List provider service is that legitimate senders are sometimes added to the IP Block lists of IP Block List providers by mistake. Legitimate senders can be added to the IP Block lists that are maintained by IP Block List provider as the result of an SMTP misconfiguration, where the SMTP server was unintentionally configured to act as an open relay is an example of such a misconfiguration.

	Connections to IP allow list providers: Indicates the number of connections on the IP Allow List providers during the last measurement period.	Number	IP Allow lists are sometimes referred to as IP safe lists or "white" lists elsewhere in the software industry. IP Allow List providers maintain lists of IP addresses that are definitively known not to be associated with any spam activity. When an IP Allow List provider returns an IP Allow match, which indicates that the sender's IP address is more likely to be a reputable or "safe" sender, the Connection Filter agent relays the message to the next agent in the anti-spam chain. The value of this measure indicates the number of connections the filtering agent has established with an IP allow list provider for checking whether the source IP address exists therein.
	Connections to IP block list: Indicates the number of connections on the IP Block List during the last measurement period.	Number	By using administrator-defined IP Allow lists and IP Block lists, you can configure connection filtering to support the following scenarios: • **To exempt IP addresses from the IP Block lists of IP Block List providers:** You may have to exempt IP addresses from the IP Block lists of IP Block List providers when legitimate senders are unintentionally put on an IP Block List provider's IP Block list. For example, legitimate senders could be unintentionally put on an IP Block list when an SMTP server was unintentionally configured to act as an open relay. In this scenario, the sender will probably try to correct the misconfiguration and remove their IP address from the IP Block List provider's IP Block list. For more information about IP Block List providers, see "IP Block List Providers" later in this topic. • **To deny access from IP addresses that are a source of unsolicited e-mail messages but are not found on an IP Block List provider's IP Block lists:** Sometimes, you may receive a large quantity of unsolicited messages from a source that was not yet identified by a *real-time block list (RBL)* service to which you subscribe.
	Connections to IP allow list: Indicates the number of connections on the IP allow list during the last measurement period.	Number	

5.2.2 Content Filters Test

Content filtering provides another tool to help manage the flow of messages entering and exiting your business's mail stream. Content filtering enables you to filter messages by using a variety of filtering tools. These include:

- **Sender-domains filtering (for Realtime and Manual scan jobs):** Sender-domains filtering enables you to filter messages from particular senders or domains.
- **Subject line filtering (for Realtime and Manual scan jobs):** Subject line filtering enables you to filter messages based on the content of the subject line of the message.
- **Filter set templates (simplify the creation and management of file and content filters on all scan jobs):** Filter set templates can be created for use with any Forefront Security for Exchange Server scan job. A single filter set template can be associated with any or all of the scan jobs and administrators can also create multiple filter set templates for use on different servers or different scan jobs.

The Content Filter agent is the last filter to scan inbound messages. While doing so, the Content Filter agent uses Microsoft SmartScreen technology to assess the contents of the messages and to assign a **spam confidence level (SCL)** rating to each message. By comparing the SCL threshold configuration with the assigned SCL rating, the content filter feature takes a specific action on a specific message, such as rejecting a message or deleting a message.

This test monitors the operations of the Content Filtering agent, reports the count of messages that have been assigned various SCL ratings, and also reveals the action the filter has taken on the messages.

Purpose	Monitors the operations of the Content Filtering agent, reports the count of messages that have been assigned various SCL ratings, and also reveals the action the filter has taken on the messages		
Target of the test	A server configured with the Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge Transport server 3. **PORT** - The port number of the Edge Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge Transport server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Messages at Spam Control Level 0:** Indicates the number of messages that were assigned a spam confidence level (SCL) rating of 0 during the last measurement period.	Number	Messages with an SCL rating of 0 are considered less likely to be spam.
	Messages at Spam Control Level 1: Indicates the number of messages assigned a spam confidence level (SCL) rating of 1 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 2: Indicates the number of messages assigned a spam confidence level (SCL) rating of 2 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 3: Indicates the number of messages assigned a spam confidence level (SCL) rating of 3 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 4: Indicates the number of messages assigned a spam confidence level (SCL) rating of 4 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 5: Indicates the number of messages assigned a spam confidence level (SCL) rating of 5 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.

	Messages at Spam Control Level 6: Indicates the number of messages assigned a spam confidence level (SCL) rating of 6 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 7: Indicates the number of messages assigned a spam confidence level (SCL) rating of 7 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 8: Indicates the number of messages assigned a spam confidence level (SCL) rating of 8 during the last measurement period.	Number	Higher the SCL rating, greater is the likelihood of the message to be spam.
	Messages at Spam Control Level 9: Indicates the number of messages assigned a spam confidence level (SCL) rating of 9 during the last measurement period.	Number	Messages with an SCL rating of 9 are considered more likely to be spam.
	Messages quarantined: Indicates the number of messages that were quarantined during the last measurement period.	Number	Quarantined messages are typically sent to the spam quarantine mailbox that you specified.
	Messages scanned: Indicates the number of messages that were scanned for viruses during the last measurement period.	Number	

	Messages rejected: Indicates the number of messages that were rejected during the last measurement period.	Number	If the connection filter rejects a message, it sends sends an SMTP error response to the sending server.
	Messages deleted: Indicates the number of messages that were deleted during the last measurement period.	Number	For deleted messages, the computer that has the Edge Transport server role installed sends a fake "OK" Simple Mail Transfer Protocol (SMTP) command to the sending server and then deletes the messages. Because the sending server assumes that the message was sent, the sending server does not retry to send the message in the same session.
	Messages with SCL unknown: Indicates the number of messages that could not be scanned by the filter during the last measurement period.	Number	Ideally, this value should be 0.
	Messages that bypassed scanning: Indicates the number of messages that bypassed scanning during the last measurement period.	Number	Forefront Security for Exchange Server can be configured to only scan file attachments that are more likely to contain viruses. It does this by first determining the file type and then by determining whether that file type can be infected with a virus. Determining the file type is accomplished by looking at the file header and not by looking at the file extension. This is a much more secure method because file extensions can be easily spoofed. This check increases Forefront Security for Exchange Server performance while making sure that no potentially infected file attachments pass without being scanned. If you would like Forefront Security for Exchange Server to bypass scanning for file types that are not commonly known to be capable of carrying a virus, set the registry key **ScanAllAttachments** to **0**.

5.2.3 Protocol Analysis Test

The Protocol Analysis / Sender Reputation agent is an anti-spam agent that is enabled on computers that are running Exchange 2007/2010 that have the Edge Transport server role installed. The Sender Reputation agent can block messages according to many characteristics of the sender. The Sender Reputation agent relies on persisted data about the sender to determine what action, if any, to take on an inbound message.

The **Sender Reputation Level (SRL)** is a number between 0 and 9 that predicts the probability that a specific sender is a spammer or malicious sender. A value of 0 indicates that the message is not likely to be spam. A value of 9 indicates that a message is likely to be spam. You can configure the threshold for sender blocking by SRL. This SRL block threshold defines the SRL value that must be exceeded for sender reputation to block a sender. If a message is

equal to or greater than the SRL block threshold, that sender will be added to the IP Block list from 0 to 48 hours. The default is 24 hours.

This test monitors the activities of the Sender Reputation agent and reveals how many senders were blocked for what reason.

Purpose	Monitors the activities of the Sender Reputation agent and reveals how many senders were blocked for what reason		
Target of the test	A server configured with the Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge Transport server 3. **PORT** - The port number of the Edge Transport server. By default, this is 691.		
Outputs of the test	One set of results for the Edge Transport server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Senders processed: Indicates the number of senders who were scanned for reputation level during the last measurement period.	Number	
	Senders blocked due to a local open proxy: Indicates the number of senders who were blocked because of a open local proxy check during the last measurement period.	Number	One of the characteristics that sender reputation evaluates is the result of a test for open proxy servers. Frequently, spammers route messages through open proxy servers on the Internet. By routing spam through open proxy servers, spammers can send messages that appear to originate from a different server than their own. A non-zero value for this measure indicates that that one/more senders were blocked because a local open proxy server was detected.
	Senders blocked due to a remote open proxy: Indicates the number of senders who were blocked because of a remote open proxy check during the last measurement period.	Number	One of the characteristics that sender reputation evaluates is the result of a test for open proxy servers. Frequently, spammers route messages through open proxy servers on the Internet. By routing spam through open proxy servers, spammers can send messages that appear to originate from a different server than their own. A non-zero value for this measure indicates that that one/more senders were blocked because a remote open proxy server was detected.

	Senders blocked due to local sender reputation level: Indicates the number of senders who were blocked because of local sender reputation level (SRL) threshold violation during the last measurement period.	Number	A high value for this measure indicates that many local senders violated the reputation level threshold. If the number is unreasonably high, you might want to review your SRL block threshold configuration. By default, the SRL threshold value is 7. Use caution when you set the SRL threshold. A threshold that is too low may unintentionally block legitimate senders. A threshold that is too high may not block malicious senders or spammers.
	Senders blocked due to remote sender reputation level: Indicates the number of senders who were blocked because of remote sender reputation level (SRL) threshold violation during the last measurement period.	Number	A high value for this measure indicates that many remote senders violated the reputation level threshold. If the number is unreasonably high, you might want to review your SRL block threshold configuration. By default, the SRL threshold value is 7. Use caution when you set the SRL threshold. A threshold that is too low may unintentionally block legitimate senders. A threshold that is too high may not block malicious senders or spammers.

5.2.4 SMTP Send Connectors Test

SMTP Send connectors are configured on computers that are running Microsoft Exchange server 2007/2010 and that have Hub Transport and Edge Transport server roles installed. The Smtp Send Connector represents a logical gateway through which outbound messages are sent.

Using this test, you can periodically observe the traffic conducted by the Send connectors.

Purpose	Periodically observe the traffic conducted by the Send connectors		
Target of the test	A server configured with the Edge Transport role		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - Indicates how often this test needs to be executed. 2. **HOST** - Indicates the IP address of the Edge Transport server 3. **PORT** - The port number of the Edge Transport server. By default, this is 691.		
Outputs of the test	One set of results for each send connector supported on the Edge Transport server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Current SMTP connections:** Indicates the current number of outbound connections from the SMTP Send connectors.	Number	
	Messages sent: Indicates the number of messages received by the SMTP Send connector each second.	Msgs / Sec	This is a good indicator of the load on the Send connector.
	Data send in SMTP messages: Indicates the number of bytes sent per second.	KB/Sec	This is a good indicator of the load on the Send connector.
	Recipients per message – average: Indicates the average recipients per message handled by this SMTP Send connector.	Recipients/Msg	
	Data transferred per connection - average: Indicates the average number of bytes sent via this connector per connection.	KB/Conn	

	Messages per connection – average: Indicates the average number of message bytes per outbound message sent.	Msgs/Conn	

6

The Integrated Exchange 2007 and Exchange 2010 Models

In addition to the models discussed above, eG Enterprise provides an integrated *Microsoft Exchange 2007* and an *Microsoft Exchange 2010* monitoring model, which includes the monitoring capabilities of each of the models discussed previously. Typically, if a single server is configured with all the server roles – i.e., Mailbox, Hub Transport, Edge Transport, and Client Access roles – you can use one of the integrated models mentioned above, depending upon the version of Exchange (whether 2007/2010) in use. Figure 6.1 depicts the integrated model. **Note that both the models will comprise of the same set of layers. The difference will be in the names of Exchange services that are monitored by default by both the models - these service names will vary from one model to another.**

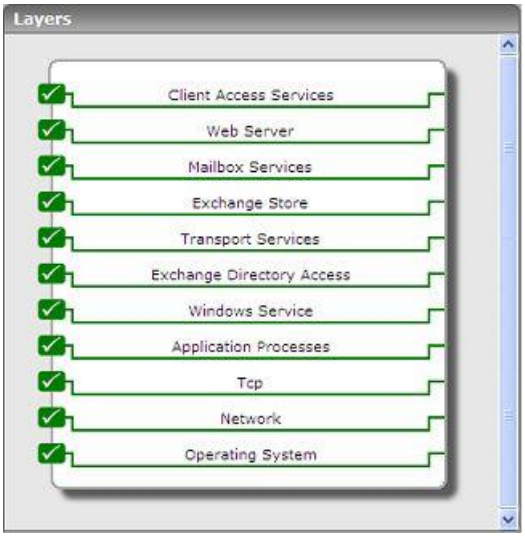

Figure 6.1: Layer model of Microsoft Exchange 2007 and Microsoft Exchange 2010

This model provides administrators an overview of the health of the entire Exchange 2007/2010 environment, and enables them to accurately pinpoint that server role which is serving as a road-block to the timely delivery of Exchange 2007/2010 services

While you can find the details of the bottom 5 layers in the *Monitoring Unix and Windows Servers* document, the top 6 layers (except the **Web Server** layer) have been dealt with in the previous chapters; for details related to the **Web Server** layer, refer to the *Monitoring Web Servers* document.

7

Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to the **Exchange 2007 and 2010 servers**. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com. We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.