



# Monitoring Microsoft Skype for Business

eG Innovations Product Documentation

# Table of Contents

---

CHAPTER 1: INTRODUCTION .....	1
CHAPTER 2: HOW TO MONITOR MICROSOFT SKYPE FOR BUSINESS USING EG ENTERPRISE? .....	2
2.1 Managing the Microsoft Skype for Business .....	2
CHAPTER 3: MONITORING MICROSOFT SKYPE FOR BUSINESS .....	4
3.1 The Windows Service Layer .....	6
3.1.1 Ls Event Log Test .....	7
3.2 The Lync Core Services Layer .....	18
3.2.1 Call Park Services Test .....	19
3.2.2 Data Collections Test .....	21
3.2.3 LS Archiving Test .....	22
3.2.4 LS Audio/Video Details Test .....	24
3.2.5 LS Routing Applications Test .....	28
3.2.6 Lync Policy Decision Test .....	31
3.2.7 Lync Registrar Test .....	33
3.2.8 Mediation Service Details Test .....	37
3.2.9 Response Group Test .....	40
3.2.10 SIP Peers Test .....	41
3.2.11 Call Details Reporting Test .....	44
3.2.12 Quality Of Experience Test .....	45
3.3 The Lync Conferencing Services Layer .....	46
3.3.1 Application Sharing Test .....	47
3.3.2 Data MCU Details Test .....	50
3.3.3 Instant Messaging Details Test .....	53
3.3.4 Lync CCCP Test .....	55
3.3.5 Lync Media Planning Test .....	57
3.3.6 Lync SIP EPS Connections Test .....	58
3.3.7 User service Conference Test .....	59
3.3.8 LS CAA Test .....	62
3.4 The Lync Web Services Layer .....	63
3.4.1 Address Book Query Test .....	64
3.4.2 Communicator Web Application Test .....	66
3.4.3 Expand Distribution List Test .....	67
3.4.4 Join Lancher Test .....	69
3.4.5 LS Replications Test .....	70
3.4.6 LS Web Information Test .....	71
3.4.7 Lync Provisioning Test .....	77

---

3.5 The Lync User Services Layer .....	77
3.5.1 LS DB Store Test .....	78
3.5.2 LS User Services Test .....	80
3.5.3 Lync LDM Info Test .....	82
3.5.4 LYSS Details Test .....	84
3.5.5 Lync Client Connections Test .....	85
3.5.6 Lync Users/Clients Test .....	86
3.5.7 Lync HTTPS Transport Details Test .....	87
ABOUT EG INNOVATIONS .....	89

## Table of Figures

---

Figure 2.1: Adding an Microsoft Skype for Business .....	2
Figure 3.1: The layer model of the Microsoft Skype for Business .....	4
Figure 3.2: The tests mapped to the Windows Service layer .....	7
Figure 3.3: Configuring an ApplicationEvents test .....	14
Figure 3.4: List of policies .....	15
Figure 3.5: Adding a new filter policy .....	16
Figure 3.6: Results of the configuration .....	18
Figure 3.7: The tests mapped to the Lync Core Services layer .....	19
Figure 3.8: The tests mapped to the Lync Conferencing Services layer .....	47
Figure 3.9: The tests mapped to the Lync Web Services layer .....	64
Figure 3.10: The tests mapped to the Lync User Service layer .....	78

## Chapter 1: Introduction

Microsoft Skype for Business Server 2013 communications software offers instant messaging (IM), presence, conferencing, and telephony solutions that can support enterprise-level collaboration requirements.

Microsoft Skype for Business is an instant messaging client used with Microsoft Skype for Business Server or Skype for Business Online available with Microsoft Office 365 and Live@Edu. If users experience slowdowns when using this client for their communication, conferencing, and collaboration purposes, they may lose confidence in the technology. To avoid this, it is imperative to watch out for issues in the operations of the server on a regular basis. This is where eG Enterprise helps administrators.

## Chapter 2: How to Monitor Microsoft Skype for Business Using eG Enterprise?

eG Enterprise is capable of monitoring the Microsoft Skype for Business server in both agent-based and agentless manners. In case of the agentless approach, the remote agent used to monitor the server should be deployed on a remote Windows host in the environment. Manage the Microsoft Skype for Business server using eG administrative interface to start monitoring the server. The procedure for achieving this is explained in the following section.

### 2.1 Managing the Microsoft Skype for Business

eG Enterprise cannot automatically discover the Microsoft Skype for Business. This implies that you need to manually add the component for monitoring. Remember that the eG Enterprise automatically manages the components that are added manually. To manage a Microsoft Skype for Business component, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENT** page that appears next, select *Microsoft Skype for Business* as the **Component type**. Then, click the **Add New Component** button. This will invoke Figure 2.1.

Figure 2.1: Adding an Microsoft Skype for Business

4. Specify the **Host IP** and the **Nick name** of the Microsoft Skype for Business in Figure 2.1. Then, click the **Add** button to register the changes.
5. Next, signout of the eG administrative interface.

## Chapter 3: Monitoring Microsoft Skype for Business

To avoid this, the eG Enterprise Suite monitors the Microsoft Skype for Business Server and captures anomalies before users notice and complain.

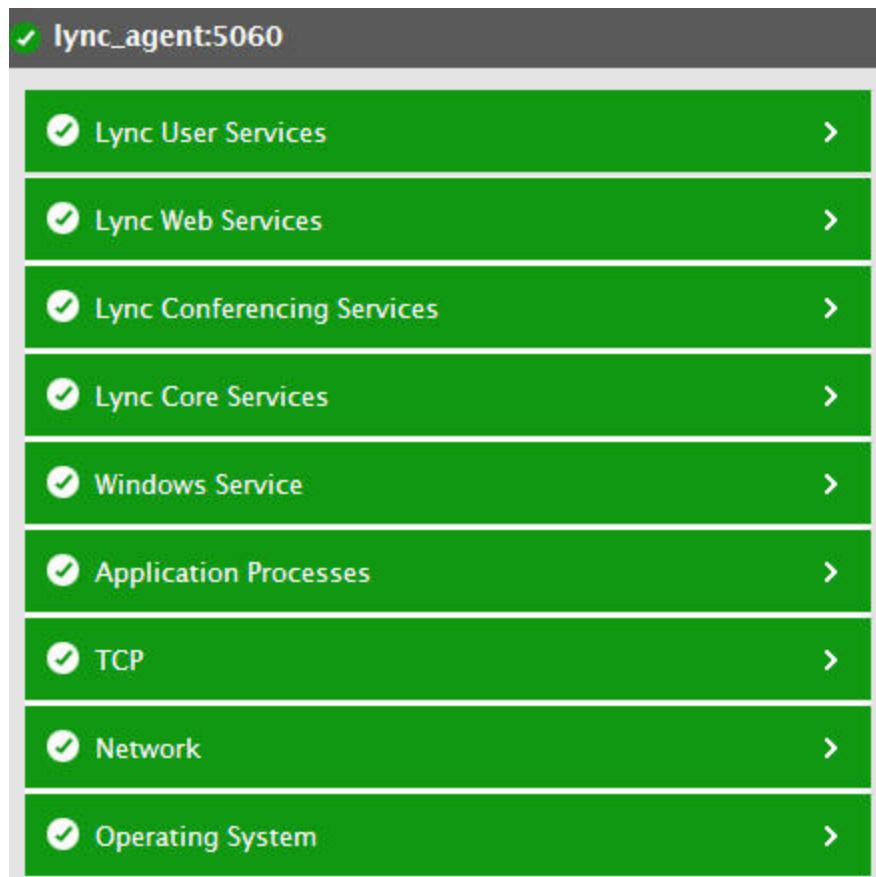


Figure 3.1: The layer model of the Microsoft Skype for Business

Using the metrics reported, administrators can find quick and accurate answers for the following performance questions:

- Observe address book accesses, measure the speed of these accesses, and report slowdowns;
- Monitor the message processing ability of the server and report deficiencies;
- Capture failed messages and when failures occurred – during message validation? in the MSMQ queue? or when written to the database?
- Measure the quality of the A/V conference experience with the Skype for Business server;
- Understand how well the server processes call park requests and report slowdowns;

- Know how many users/clients are currently connected to the server, and thus gauge the current load on the server;
- Identify dropped conference activities and unfinished tasks
- Monitor the server's interactions with the database and in the process, reveal the requests queued for processing by the database and the time these requests spent in queue; database slowdowns can thus be captured;
- Rapidly detect client and server authentication failures and DNS resolution failures;
- Messages that could not be sent to the server
- DNS resolution failures
- The current health and draining state of the Application sharing Conference unit, the Data MCU server and the Instant Messaging multipoint control unit;
- Know how well the conferences are handled by the Application sharing conference unit
- How many RDP connections failed in the Application sharing Control unit?
- How many whiteboards were served by the Data MCU server and how many conferences are currently active on the Data MCU server?;
- How many conferences were active on the Instant Messaging multipoint control unit and how many message delivery failure notification messages were sent from the server?;
- Know how many add user and add conference responses failed in the server
- How many conferences were initiated from each service and how many times conference processing was actually delayed?;
- How many times the application endpoint creation failed for the Conferencing Auto Application service and how many incoming calls were actually incomplete?
- Monitor the message processing ability of the server and report deficiencies;
- How many SIP connections failed and how many are currently active?;
- How long the distribution list take to process the requests?
- How many times the Join launcher service failed and how many incoming requests were received by the service?
- How many replication requests were received by the replication service and how long does the service take to process?;

- Continuously monitor the mobility services and figure out the successful Get Location requests, the push notifications that failed and throttled;
- How many times the user provisioning and publish calls failed?
- Monitor the backend database server and figure out how well requests are processed and how long does it take to process the requests?;
- How well the stored procedure calls were executed by the user service?;
- How many stale items were in the queue of the Lync storage service and how much space did the service utilize to store the items to the backend database server?;
- For each client version, how many unique users were actually connected?;
- How many users were actually connected to the server?
- For how many users were voice calls enabled?
- How many times the HTTPS connections failed?

The **Operating System, Network, TCP and Application Process** layers of the Microsoft Skype for Business monitoring model is similar to that of a Windows Generic server model. Since the tests pertaining to this layer have been dealt with in the Monitoring Unix and Windows Servers document, the section to come focuses on the **Windows Service** layer.

### 3.1 The Windows Service Layer

Using the test mapped to this layer, you can obtain statistical information about the event log related events generated by the target Microsoft Skype for Business server.

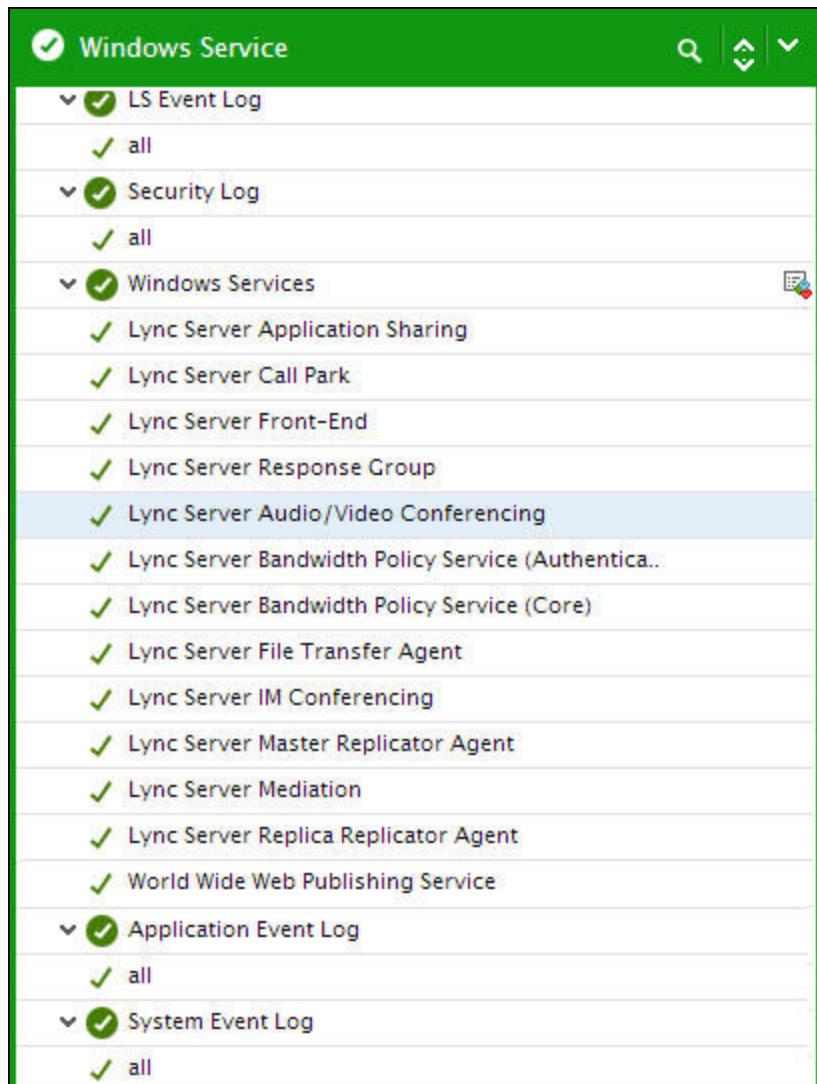


Figure 3.2: The tests mapped to the Windows Service layer

Since the **Windows Services** test has already been discussed in the *Monitoring Unix and Windows Servers* document and the rest of the tests have been discussed in the *Monitoring Event Logs* document, let us now discuss the **Ls Event Log** test in the forthcoming section.

### 3.1.1 Ls Event Log Test

This test reports the statistical information about the event log related events generated by the target Microsoft Skype for Business server.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the Filter configured for the App-V Management server that is to be monitored.

### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.
Policy Based Filter	<p>Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options:</p> <ul style="list-style-type: none"> <li>Manually specify the event sources, IDs, and descriptions in the Filter text area, or,</li> <li>Select a specification from the predefined filter policies listed in the Filter box</li> </ul> <p>For explicit, manual specification of the filter conditions, select the <b>No</b> option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the <b>Yes</b> option against the Policy Based Filter field.</p>
Filter	<p>If the Policy Based Filter flag is set to <b>No</b>, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: <i>{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}</i>. For example, assume that the Filter text area takes the value, <i>OS_events:all:Browse,Print:all:none:all:none</i>. Here:</p> <ul style="list-style-type: none"> <li><i>OS_events</i> is the display name that will appear as a descriptor of the test in the monitor UI;</li> <li><i>all</i> indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify <i>none</i>.</li> <li>Next, to ensure that specific event sources are excluded from monitoring,</li> </ul>

Parameters	Description
	<p>provide a comma-separated list of source names. Accordingly, in our example, <i>Browse</i> and <i>Print</i> have been excluded from monitoring. Alternatively, you can use <i>all</i> to indicate that all the event sources have to be excluded from monitoring, or <i>none</i> to denote that none of the event sources need be excluded.</p> <ul style="list-style-type: none"> <li>• In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The <i>all</i> in our example represents that all the event IDs need to be considered while monitoring.</li> <li>• Similarly, the <i>none</i> (following <i>all</i> in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying <i>all</i> makes sure that all the event IDs are excluded from monitoring.</li> <li>• The <i>all</i> which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use <i>none</i>. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - <i>desc*</i>, or <i>desc</i>, or <i>*desc*</i>, or <i>desc*</i>, or <i>desc1*desc2</i>, etc. <i>desc</i> here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.</li> <li>• In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: <i>desc*</i>, or <i>desc</i>, or <i>*desc*</i>, or <i>desc*</i>, or <i>desc1*desc2</i>, etc. <i>desc</i> here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use <i>all</i> instead, it would mean that all event descriptions are to be excluded from monitoring.</li> </ul> <p>By default, the Filter parameter contains the value: <i>all</i>. Multiple filters are to be separated by semi-colons ( ; ).</p>

Parameters	Description
	<p><b>Note:</b></p> <p>The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window.</p> <p>On the other hand, if the Policy Based Filter flag is set to <b>Yes</b>, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format:</p> <pre data-bbox="442 629 1434 741">{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_ids_to_be_included}:{event_ids_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}</pre> <p>To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a  icon appears near the Filter list box, once the <b>Yes</b> option is chosen against Policy Based Filter. Clicking on the  icon leads you to a page where you can modify the existing policies or create a new one (refer to Section 3.1.1.1). The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.</p>
UseWMI	<p>The eG agent can either use <b>WMI</b> to extract event log statistics or directly parse the event logs using <b>event log APIs</b>. If the UseWMI flag is <b>Yes</b>, then <b>WMI</b> is used. If not, the <b>event log APIs</b> are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to <b>No</b>. <b>On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the USEWMI flag should always be set to 'Yes'.</b></p>
Stateless Alerts	<p>Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a <b>CRITICAL</b> email alert with the details of the error event to</p>

Parameters	Description
	<p>configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as <b>CRITICAL</b>, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the <b>stateless alerting</b> capability. To enable this capability for this test, set the Stateless Alertsflag to <b>Yes</b>. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.</p>
Events During Restart	<p>By default, the Events During Restart flag is set to <b>Yes</b>. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to <b>No</b> ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.</p>
DDforInformation	<p>eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to <b>Yes</b>, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforInformation flag to <b>No</b>.</p>
DDforWarning	<p>To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to <b>No</b>.</p>
DD Frequency	<p>Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <b>1:1</b>. 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 <b>none</b> against DD Frequency.</p>
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> </ul>

Parameters	Description
<ul style="list-style-type: none"> <li>Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>	

### Measurements made by the test

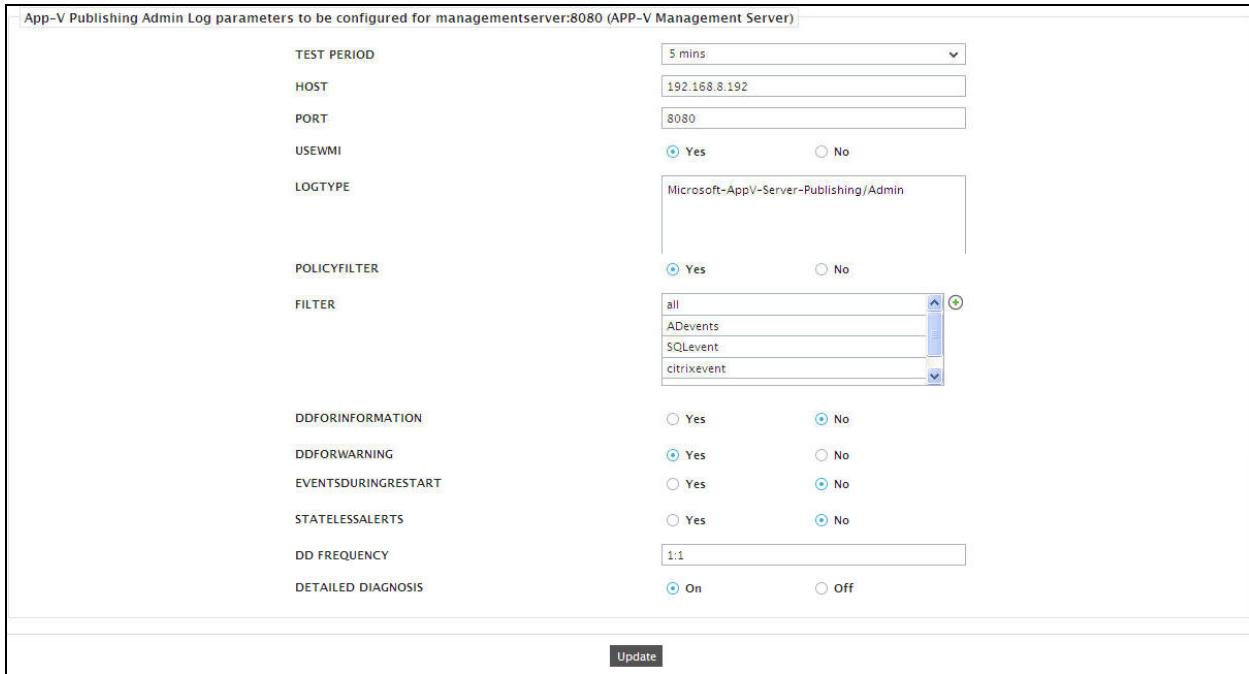
Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of server publishing admin information events generated when the test was last executed.	Number	<p>A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications.</p> <p>Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.</p>
Warnings	Indicates the number of server publishing admin warnings that were generated when the test was last executed.	Number	<p>A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications.</p> <p>Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.</p>
Error messages	Indicates the number of server publishing admin error events that were generated when the test was last executed.	Number	<p>A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems.</p> <p>An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications.</p> <p>Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.</p>
Critical messages	Indicates the number of server publishing admin critical error events that	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly

Measurement	Description	Measurement Unit	Interpretation
	were generated when the test was last executed.		<p>without any potential problem.</p> <p>An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications.</p> <p>Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.</p>
Verbose messages	Indicates the number of server publishing admin verbose events that were generated when the test was last executed.	Number	<p>Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better.</p> <p>The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.</p> <p>Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.</p>

### 3.1.1.1 Adding a New Policy

To add a new policy, do the following:

1. Click on the  icon available next to the **FILTER** list box. (see Figure 3.3).



The screenshot shows a configuration dialog for an ApplicationEvents test. The 'FILTER' list box is currently selected, as indicated by a green border around the list item 'all'. Other items in the list include 'ADevents', 'SQLevent', and 'citrixevent'. The dialog includes fields for 'TEST PERIOD' (5 mins), 'HOST' (192.168.8.192), 'PORT' (8080), 'USEWMI' (Yes), 'LOGTYPE' (Microsoft-AppV-Server-Publishing/Admin), 'POLICYFILTER' (Yes), and various policy settings like 'DDFORINFORMATION' (No), 'DDFORWARNING' (Yes), 'EVENTSDURINGRESTART' (Yes), 'STATELESSALERTS' (Yes), 'DD FREQUENCY' (1:1), and 'DETAILED DIAGNOSIS' (On). A 'Update' button is at the bottom.

Figure 3.3: Configuring an ApplicationEvents test

2. Figure 3.4 will then appear listing the policies that pre-exist.

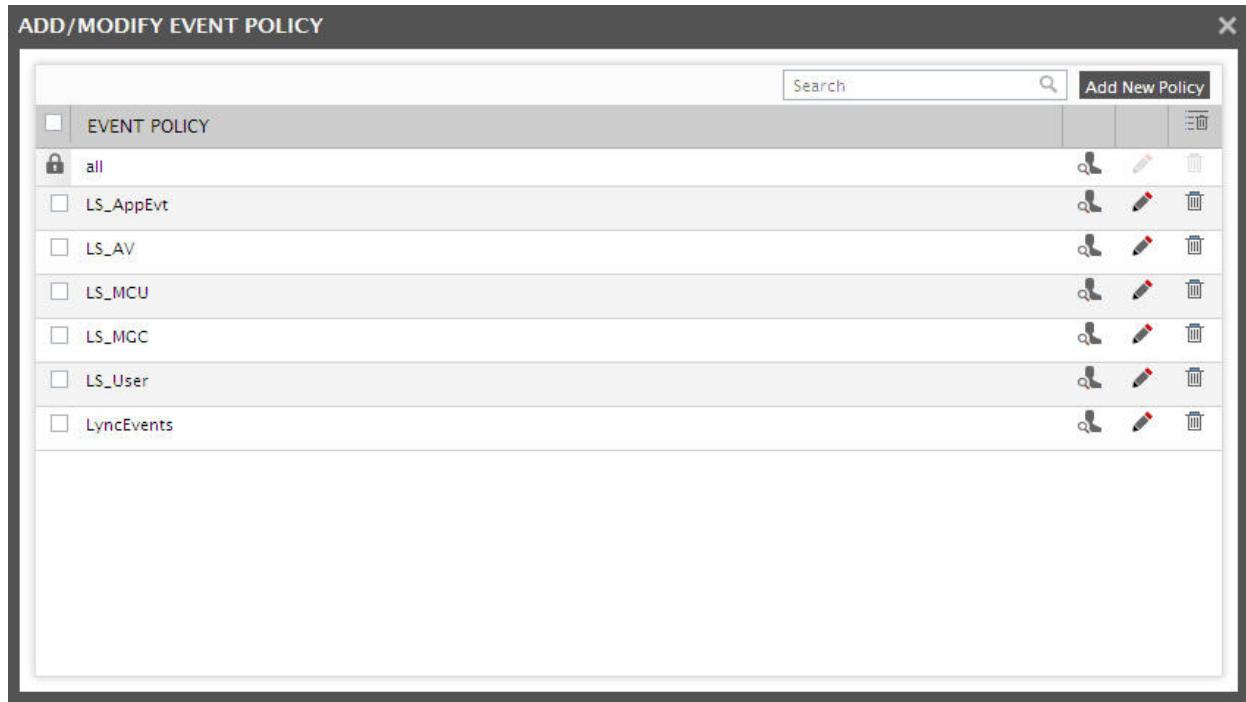


Figure 3.4: List of policies

3. To view the contents of a policy, click on the icon against the policy name. While a policy can be modified by clicking on the , it can be deleted using the icon. The default policy is all, which can only be viewed and not modified or deleted. The specification contained within this policy is: `all:none:all:none:all:none`.
4. To create a new policy, click on the **Add New Policy** button in Figure 3.4. Doing so invokes Figure 3.5, using which a new policy can be created.

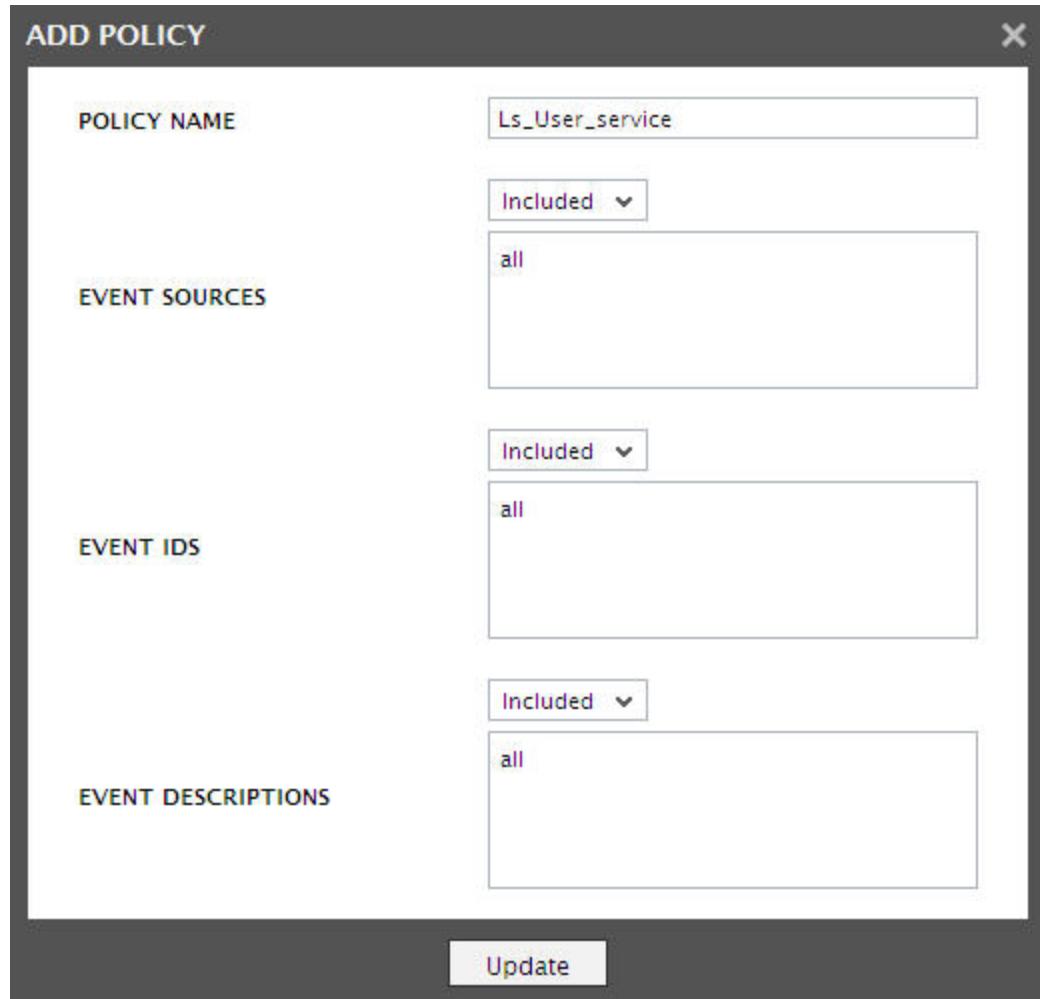


Figure 3.5: Adding a new filter policy

5. In Figure 3.5, first, provide a unique name against **POLICY NAME**.
6. To include one/more event sources for monitoring, select **Included** from the **EVENT SOURCES** drop-down list, and then specify a comma-separated list of event sources to be included in the list box available below the drop-down list.
7. To exclude specific event sources from monitoring, select **Excluded** from the **EVENT SOURCES** drop-down list, and then specify a comma-separated list of event sources to be excluded in the list box available below the drop-down list.

**Note:**

At any given point in time, you can choose to either **Include** or **Exclude** event sources, but you cannot do both. If you have chosen to include event sources, then the eG Enterprise system automatically assumes that no event sources need to be excluded. Accordingly, the `{event_`

*sources\_to\_be\_excluded}* section of the filter format mentioned above, will assume the value *none*. Similarly, if you have chosen to exclude specific event sources from monitoring, then the *{event\_sources\_to\_be\_included}* section of the format above will automatically take the value *all*, indicating that all event sources except the ones explicitly excluded, will be included for monitoring.

8. In the same way, select **Included** from the **EVENT IDS** list and then, provide a comma-separated list of event IDs to be monitored.
9. If you, on the other hand, want to exclude specific event IDs from monitoring, then first select **Excluded** from the **EVENT IDS** list box, and then provide a comma-separated list of event IDs to be excluded.

**Note:**

At any given point in time, you can choose to either **Include** or **Exclude** event IDs, but you cannot do both. If you have chosen to include event IDs, then the eG Enterprise system automatically assumes that no event IDs need be excluded. Accordingly, the *{event\_ids\_to\_be\_excluded}* section of the filter format mentioned above, will assume the value *none*. Similarly, if you have chosen to exclude specific event IDs from monitoring, then the *{event\_ids\_to\_be\_included}* section of the format above will automatically take the value *all*, indicating that all event IDs except the ones explicitly excluded, will be included for monitoring.

10. Likewise, select **Included** from the **EVENT DESCRIPTIONS** list and then, provide a comma-separated list of event descriptions to be monitored.
11. For excluding specific event descriptions from monitoring, first select **Excluded** from the **EVENT DESCRIPTIONS** list box, and then provide a comma-separated list of event descriptions to be excluded.

Instead of the complete event descriptions, wild card-embedded event description patterns can be provided as a comma-separated list in the **Included** or **Excluded** text boxes. For instance, to include all events that start with *st* and *vi*, your **Included** specification should be: *st\*,vi\**. Similarly, to exclude all events with descriptions ending with *ed* and *le*, your **Excluded** specification should be: *\*ed,\*le*.

**Note:**

At any given point in time, you can choose to either **Include** or **Exclude** event descriptions/users, but you cannot do both. If you have chosen to include event descriptions/users, then the eG Enterprise system automatically assumes that no event descriptions/users need be excluded. Accordingly, the *{event\_descriptions\_to\_be\_excluded}* section or the *{users\_to\_be\_excluded}*

section (as the case may be) of the filter formats mentioned above, will assume the value *none*. Similarly, if you have chosen to exclude specific event descriptions/users from monitoring, then the `{event_descriptions_to_be_included}` section or the `{users_to_be_included}` section (as the case may be) of the formats above will automatically take the value *all*. This indicates that all event descriptions/users except the ones explicitly excluded, will be included for monitoring.

- Finally, clicking the **Update** button will display a pop up window as depicted by Figure 3.6.

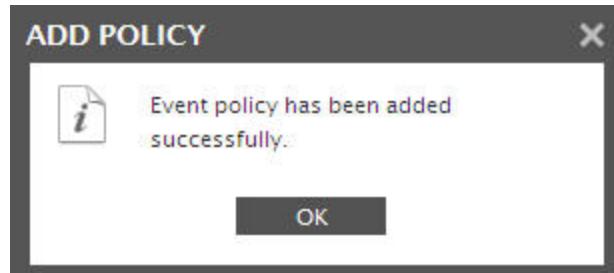


Figure 3.6: Results of the configuration

If you have configured a policy to **Include** a few/all events (sources/IDs/descriptions/users), and **Exclude** *none*, then, while reconfiguring that policy, you will find that the **Include** option is chosen by default from the corresponding drop-down list in Figure 3.5. On the other hand, if you have configured a policy to **Exclude** a few specific events and **Include all** events, then, while modifying that policy, you will find the **Exclude** option being the default selection in the corresponding drop-down list in Figure 3.5.

## 3.2 The Lync Core Services Layer

Using the tests mapped to this layer, administrators can proactively detect the following:

- Capture failed messages and when failures occurred – during message validation? in the MSMQ queue? or when written to the database?
- Measure the quality of the A/V conference experience with the Skype for Business server;
- Understand how well the server processes call park requests and report slowdowns;
- Know how many users/clients are currently connected to the server, and thus gauge the current load on the server;
- Identify dropped conference activities and unfinished tasks
- Know how many messages were dropped from the message queue and how many actually failed to be written to the database;



Figure 3.7: The tests mapped to the Lync Core Services layer

Let us discuss each test of this layer (see Figure 3.7) in the forthcoming sections.

### 3.2.1 Call Park Services Test

Enterprise Voice call management features control how incoming calls are routed and answered. Skype for Business Server 2013 provides the following call management features:

- **Call Park:** Enables voice users to temporarily park a call and then pick it up from the same phone or another phone.
- **Group Pickup:** Enables users to answer calls made to another user who is assigned to a pickup group by dialing the call pickup group number.
- **Response Group:** Routes incoming calls to groups of agents by using hunt groups or interactive voice response (IVR) questions and answers.
- **Announcement:** Plays a message for calls made to an unassigned number, or routes the call elsewhere, or both.

The Call Park application enables an Enterprise Voice user to put a call on hold from one telephone and then retrieve the call later from any telephone. When the user parks a call, Skype for Business Server transfers the call to a temporary number, called an orbit, where the call is held until someone retrieves it or it times out. In environments where audio/video calls play a major role, the administrators may find it difficult to track the number of calls routed to the users at any given point of time. The **Call Park Services** test addresses such issues. Using this test, administrators can easily identify the number of calls that were parked, the call park requests that failed, the call park requests that failed due to unavailability of the orbit, number of times a call was failed to be retrieved etc. This

way, administrators may be able to identify the load on the Skype for Business server and act accordingly.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the Microsoft Skype for Business server being monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Park requests failed	Indicates the number of call park requests that failed during the last measurement period.	Number	
Park requests failed due to unavailable orbit	Indicates the number of call park requests that failed due to the unavailability of an orbit during the last measurement period.	Number	The Call Park application uses extension numbers in the Call Park orbit table to park calls. You need to configure the Call Park orbit table with the ranges of extension numbers that your organization reserves for parked calls. These extensions need to be virtual extensions (that is, extensions that have no user or phone assigned to them). Each Skype for Business Server pool where a Call Park application is deployed and configured can have one or more orbit ranges. Orbit ranges must be globally unique across the Skype for Business Server

Measurement	Description	Measurement Unit	Interpretation
			deployment. An orbit may be unavailable if there are too many calls that are currently put on hold.
Failed fallback attempts	Indicates the number of fallback attempts that failed during the last measurement period.	Number	A high value for this measure is a cause of concern. This may be due to the fallback destination being unreachable.
Total SIP requests received	Indicates the number of SIP requests received since the start of the call park service.	Number	
Parked calls	Indicates the number of calls that are currently parked.	Number	
Total number of time the SQL server was called	Indicates the number of times the backend SQL server was called by this service during the last measurement period.	Number	
Failed orbit requests	Indicates the number of times the request to retrieve a call from the orbit failed.	Number	Ideally, the value of this measure should be zero.

### 3.2.2 Data Collections Test

In large environments where multiple conference calls are made through Microsoft Skype for Business server, it is not possible for the administrators to monitor the quality of each conference call and figure out what exactly caused a joinee to drop out of the conference call – is it due to database connectivity issues or validation errors?. Administrators are required to figure out how many times the conference calls were dropped over a period of time so that issues can be sorted out before end users start complaining. The **Data Collections** test exactly serves this purpose! Using this test, administrators may figure out the number of dropped conference join and leave activities from the server and also the number of unfinished tasks in the server.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data Collections	Indicates the number of conference leave and join activities that were dropped from this server.	Number	The conference join and leave activities may be dropped due to archiving database connectivity issues or validation errors.
Unfinished tasks	Indicates the number of times tasks that were unfinished were found on this server.	Number	

### 3.2.3 LS Archiving Test

Archiving in Skype for Business Server 2013 provides a way for you to archive communications that are sent through Skype for Business Server 2013. The following types of content can be archived:

- Peer-to-peer instant messages
- Conferences (meetings), which are multiparty instant messages
- Conference content, including uploaded content (for example, handouts) and event-related content (for example, joining, leaving, uploading sharing, and changes in visibility)
- Whiteboards and polls shared during a conference

Archiving is automatically installed on each Front End Server when you deploy the server, but Archiving is not enabled until you configure it. The Archiving feature includes the following components:

**Archiving agents.** Archiving agents (also known as unified data collection agents) are installed and activated automatically on every Front End pool and Standard Edition server. Although archiving agents are activated automatically, no messages are actually captured until Archiving is enabled and appropriately configured.

**Archiving data storage.** Data storage for Skype for Business Server 2013 can be either of the following:

- **Archiving using Microsoft Exchange integration** - If you have users who are homed on Exchange 2013 and their mailboxes have been put on In-Place Hold, you can select the option to integrate Skype for Business Server 2013 storage with Exchange storage. If you choose the Microsoft Exchange integration option, you use Exchange 2013 policies and configurations to control the archiving of Skype for Business Server 2013 data for those users.
- **Archiving using Skype for Business Server Archiving databases** - If you have users who are not homed on Exchange 2013 or who have not had their mailboxes put on In-Place Hold, or if you don't want to use Microsoft Exchange integration for any or all users in your deployment, you can deploy Skype for Business Server Archiving databases using SQL Server to store Archiving data for those users. In this case, Skype for Business Server 2013 Archiving policies and configurations determine whether Archiving is enabled and how it is implemented. To use Skype for Business Server 2013, you must add the appropriate SQL Server databases to your topology and publish the topology.

This test monitors the archiving pattern of the Microsoft Skype for Business and provides you with the numerical statistics of the following information:

- The messages that failed validation;
- The messages that were dropped;
- The messages that failed to be written to the backend database server and
- The messages that were sent for archiving.

Using this test, administrators can easily identify technical glitches in the Microsoft Skype for Business that caused issues which eventually led to message processing deficiencies.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

#### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

#### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Messages for which validation failed	Indicates the number of messages that failed validation.	Number	A low value is desired for this measure. A rapid increase in the value of this measure may indicate the presence of malformed messages in the message queue.
Messages dropped from MSMQ queue	Indicates the number of messages that were dropped from the MSMQ queue.	Number	
Messages failed to be written to DB	Indicates the number of messages that failed to be written to the backend database server.	Number	A low value is desired for this measure. A sudden/gradual increase in this value is a cause of concern which clearly indicates the failure of the service that is responsible for writing the messages to the database.
Archived messages	Indicates the number of messages that were sent for archiving to the Archiving Message queue.	Number	A high value is desired for this measure.

#### 3.2.4 LS Audio/Video Details Test

A/V conferencing enables real-time audio and video communications between the Skype for Business users. In environments where users use A/V conferences for communication, it becomes

mandatory for the administrator to check the quality of the A/V Conferencing experience and the load on the Microsoft Skype for Business server so that administrators may proactively be alerted to abnormalities/technical glitches in the conferences. This is where the **LS Audio/Video Details** test helps. Using this test, you can easily figure out the health of the A/V Conferencing server, the number of conferences that were created and those that are currently active, the number of RTP endpoint connectivity status events in the A/V Conferencing server etc. This way, administrators can easily identify the load on the Microsoft Skype for Business server.

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Active conferences on the A/V conferencing server	Indicates the number of conferences that were active on the A/V Conferencing server.	Number	This measure is a good indicator of the load on the Microsoft Skype for Business server.
Created conferences	Indicates the number of conferences created during the last measurement period.	Number	
SIP Bye messages sent by A/V conferencing server	Indicates the number of SIP BYE messages sent by the A/V Conferencing server during the last measurement period.	Number	To deliver a conference invitation, the inviting user sends a special sort of SIP INVITE message to the user who is being invited. Instead of Session Description Protocol (SDP) content, this INVITE message has an XML document with conference details as

Measurement	Description	Measurement Unit	Interpretation										
			its body. The invitee sends a 200 OK acknowledging this INVITE, and as soon as the SIP handshake has finished it immediately ends the SIP dialog with a BYE message.										
AVMCU health state	Indicates the current health state of the A/V conferencing server.		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Normal</td><td>0</td></tr> <tr> <td>Loaded</td><td>1</td></tr> <tr> <td>Full</td><td>2</td></tr> <tr> <td>Unavailable</td><td>3</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the current health of the A/V Conferencing server. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 3.</p>	Measure Value	Numeric Value	Normal	0	Loaded	1	Full	2	Unavailable	3
Measure Value	Numeric Value												
Normal	0												
Loaded	1												
Full	2												
Unavailable	3												
Draining status of the AVMCU	Indicates the draining status of the A/V Conferencing server.		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Not requesting to drain</td><td>0</td></tr> <tr> <td>Requesting to drain</td><td>1</td></tr> <tr> <td>Draining</td><td>2</td></tr> </tbody> </table>	Measure Value	Numeric Value	Not requesting to drain	0	Requesting to drain	1	Draining	2		
Measure Value	Numeric Value												
Not requesting to drain	0												
Requesting to drain	1												
Draining	2												

Measurement	Description	Measurement Unit	Interpretation
			<p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the draining status of the A/V Conferencing server. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 2.</p> <p>When a server is drained, it stops taking new connections and calls. These new connections and calls are routed through other servers in the pool. A server being drained allows its sessions on existing connections to continue until they naturally end. When all existing sessions have ended, the server is ready to be taken offline.</p>
Active media processing contexts	Indicates the number of active media processing contexts in the A/V Conferencing server.	Number	
CreateContext requests received	Indicates the number of CreateContext requests received by the A/V Conferencing server during the last measurement period.	Number	
Total RTP Endpoint connectivity status events	Indicates the number of RTP Endpoint connectivity status events on the A/V Conferencing server during the last measurement period.	Number	

### 3.2.5 LS Routing Applications Test

The VoIP components located on Front End Servers of the Microsoft Skype for Business server are as follows:

- Translation Service
- Inbound Routing component
- Outbound Routing component
- Exchange UM Routing component
- Intercluster Routing component
- Mediation Server component in Skype for Business Server 2013

When a call request is received by the front end servers, the components mentioned above processes the call requests and routes it to the destination. Whenever the requests stagnate and are unable to be routed, the requests time out. Administrators may want to identify the exact technical glitch that caused a large number of call requests to time out frequently in the registrar pools of the front end servers. This is where the **LS Routing Applications** test helps! This test reports on how well the call requests are processed and how many requests timed out in the primary and backup registrar pools. Alongside, administrators may be alerted to the number of emergency calls received and the telephone numbers that were translated.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

## Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Inter cluster routing requests	Indicates the number of call requests processed by the intercluster routing component.	Number	The Intercluster routing component is responsible for routing calls to the callee's primary Registrar pool. If that is unavailable, the component routes the call to the callee's backup Registrar pool. If the callee's primary and backup Registrar pools are unreachable over the IP network, the Intercluster routing component reroutes the call over the PSTN to the user's telephone number.
Primary registrar timeouts	Indicates the number of call requests that timed out in the primary registrar pool.	Number	<p>The Skype for Business Server Registrar is a server role that enables client registration and authentication and provides routing services. It resides along with other components on a Standard Edition server, Front End Server, Director, or Survivable Branch Appliance. A Registrar pool consists of Registrar Services running on the Front End pool and residing at the same site. A Skype for Business client discovers the Front End pool through the following discovery mechanism:</p> <ul style="list-style-type: none"> <li>• DNS SRV record</li> <li>• Autodiscovery Web Service (new in Skype for Business Server 2013)</li> <li>• DHCP option 120</li> </ul> <p>After the Skype for Business client connects to the Front End pool, it is directed by the load balancer to one of the Front End Servers in the pool. That Front End Server, in turn, redirects the client to a preferred Registrar in the</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>pool.</p> <p>Each user enabled for Enterprise Voice is assigned to a particular Registrar pool, which becomes that user's primary Registrar pool.</p> <p>To assure voice resiliency in the event of a central site failure, the primary Registrar pool must have a single designated backup Registrar pool located at another site. Assuming a resilient WAN link between the two sites, users whose primary Registrar pool is no longer available are automatically directed to the backup registrar pool.</p>
Backup registrar timeouts	Indicates the number of call requests that timed out in the backup registrar pool.	Number	
Received emergency calls	Indicates the number of emergency calls received in the pool.	Number	
Emergency call failure responses	Indicates the number of emergency calls that failed to receive a response at the gateway.	Number	Ideally, the value of this measure should be zero.
Phone numbers translated	Indicates the total number of telephone numbers translated by the Translation Service.	Number	The Translation Service is the server component that is responsible for translating a dialed number into the E.164 format or another format, according to the normalization rules that are defined by the administrator. The Translation Service can translate to formats other than E.164 if your organization uses a private numbering system or uses a gateway or PBX that does not support E.164.

### 3.2.6 Lync Policy Decision Test

Real-time communications are sensitive to the latency and packet loss that can occur on congested networks. Call admission control (CAC) determines, based on available network bandwidth, whether to allow real-time communications sessions such as voice or video calls to be established. The Call Admission Control (CAC) uses the Policy Decision Point (PDP) to distribute policies defining some of the CAC rules. CAC controls real-time traffic for voice and video only. It does not control data traffic.

Administrators define CAC policies, which are enforced by the Bandwidth Policy Service that is installed with every Front End pool. CAC settings are automatically propagated to all Skype for Business Server Front End Servers in your network.

This test helps you figure out how many client authentications failed and how many client connections were authenticated. Using this test, administrators can easily figure out the number of active audio and video call requests for which bandwidth is checked. This way, administrators can figure out the client connection that is consuming too much of bandwidth.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Rate of client authentications failures due to	Indicates the rate at which client authentication failed due to timeout.	Failures/Sec	A low value is desired for this measure. If there are too many authentication failures, then check for

Measurement	Description	Measurement Unit	Interpretation
timeout			certification issues. The value of this measure may be high due to an attack by the client on the machine through the creation of connections and consumption of resources from the Bandwidth Policy Service. If the rise in the value is due to an attack, then trace the machine and block the machine through firewall policies.
Rate of first packet timeouts	Indicates the rate at which the sessions timed out before the arrival of the first packet.	Timeouts/Sec	
Rate of server authentication timeouts	Indicates the rate at which the server connections timed out before an authentication message was received.	Failures/Sec	A low value is desired for this measure.
Active authenticated client connections	Indicates the number of active client connections that are authenticated.	Number	
BW check requests for video	Indicates the number of video call requests for which bandwidth is checked.	Number	
Active video BW reservations	Indicates the number of active bandwidth reservations for video calls.	Number	
BW check requests for audio	Indicates the number of audio call requests for which bandwidth is checked.	Number	
Active audio BW reservations	Indicates the number of active bandwidth reservations for audio calls.	Number	

Measurement	Description	Measurement Unit	Interpretation
Total core BW check requests	Indicates the overall number of core (both audio and video) call requests for which bandwidth is checked.	Number	
Active core BW reservations	Indicates the total number of active reservations for audio and video calls.	Number	
TCP segments received rate	Indicates the rate at which TCP segments were received by this service.	Received/Sec	

### 3.2.7 Lync Registrar Test

Registrars are used to authenticate logon requests, and to maintain information about user status and availability. The Registrar is perhaps the most important component in Skype for Business Server; without a Registrar, users would not be able to log on to the system, and Skype for Business Server would not be able to keep track of users and their current status. When a user logs on to Skype for Business Server, the endpoint the user is logging on from (be it a computer, a mobile phone, or some other device) sends a REGISTER request to the registration server; in turn the server responds by challenging the client device for authentication credentials. If the client passes the challenge (that is, if the client presents a valid set of credentials), then the user is authenticated and endpoint information such as IP address, port, and user name is logged in the registration database. When a user logs off, this information is then removed from the database. In between log on and log off, the Registrar keeps status information up-to-date and helps to route messages to and from the user. Often administrators may want to identify how well the information i.e., messages are received and processed by the server from the user and rate the user experience on the server. This is where the **Lync Registrar** test helps!

This test monitors the messages sent to the server and reports the following:

- The number of queued messages that were dropped, the number of stored security associations and security records;
- The number of incoming messages that timed out and those that were held for a longer duration;

- The number of requests that were received by the registrar and the number of requests that were dropped;
- The responses that were received for the incoming messages;
- The DHCP requests and SIP protocol requests received etc.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Queued messages that could not be sent	Indicates the number of messages in the message queue that could not be sent to the monitoring server and were eventually dropped.	Number	Ideally, the value of this measure should be zero. If the value of this measure is high, check if the message queuing is working properly. Also, ensure that the QoE Monitoring agent has permissions to access the MSMQ.
Rate of authentication system errors	Indicates the rate at which authentication failed due to system errors.	Errors/Sec	A high value indicates that this server may be experiencing problems when using Windows Authentication subsystem.
Stored security associations	Indicates the total number of security associations that are currently stored in the server.	Number	

Measurement	Description	Measurement Unit	Interpretation
Stored security records	Indicates the total number of security records that are currently stored in the server.	Number	A security record is created for each user authenticated on the Skype for Business server.
Avg holding time for incoming messages	Indicates the average time the incoming messages that were being processed by the server were put on hold.	Secs	The value of this measure should be in the range of 1-3 seconds. The server will throttle new incoming messages after going above the high watermark and until the number of messages falls below the low watermark. The server starts rejecting new connections when the average holding time is greater than overload time of 15 seconds.  A high value for this measure indicates that the server is too busy and is unable to process the messages on time.
Incoming messages that timed out	Indicates the number of incoming messages that timed out.	Number	The incoming messages may time out when the server is processing the messages beyond the maximum tracking interval.
Incoming messages held longer than permitted	Indicates the number of incoming messages that were processed by the server were put on hold beyond the high watermark time threshold.	Number	A value close to 0 is desired for this measure.
Incoming messages held above high overload watermark	Indicates the number of incoming messages that were processed by the server were put on hold beyond the overload watermark time threshold.	Number	
Avg incoming message processing time	Indicates the average time taken to process incoming SIP protocol messages.	Secs	
Rate of incoming	Indicates the rate at which	Messages/Sec	

Measurement	Description	Measurement Unit	Interpretation
messages	the incoming SIP protocol messages are received.		
Rate of events in processing	Indicates the number of SIP transactions or dialog state change events that are currently processed.	Number	
Received requests rate	Indicates the rate at which SIP protocol requests were received.	Requests/Sec	
Incoming requests dropped rate	Indicates the rate at which incoming SIP protocol requests were dropped.	Requests/Sec	The incoming requests/incoming responses may be dropped due to bad headers, insufficient routing information, server resource allocation failure etc.
Incoming responses dropped rate	Indicates the rate at which incoming SIP protocol responses were dropped.	Requests/Sec	
Connections refused due to server overload	Indicates the number of connections that were refused due to the overload condition of the server.	Number	The connections may be refused with the Service Unavailable response when the server is overloaded.
Opened sockets	Indicates the number of opened sockets.	Number	
Rate of 500 responses	Indicates the rate at which Local 500 Responses were generated by the server.	Responses/Sec	A Local 500 Response indicates that a server component of the Skype for Business server is not functioning properly.
Rate of 503 responses	Indicates the rate at which Local 503 Responses were generated by the server.	Responses/Sec	A Local 503 Response corresponds to the server being unavailable. On a healthy server, you should not receive this code at a steady rate. However, during ramp up, after a server has been brought back online, there may be some 503 responses. Once all users get back in and the server returns to a stable state, there should no longer be any 503 responses returned.

Measurement	Description	Measurement Unit	Interpretation
Rate of 504 responses	Indicates the rate at which Local 504 Responses were generated by the server.	Responses/Sec	A Local 504 Response indicates connectivity issues with other servers. It can indicate connection failures or delays connecting to remote servers.
DHCP requests	Indicates the total number of DHCP requests.	Number	
Ack requests received	Indicates the total number of incoming SIP protocol ACK requests received.	Number	
Routed requests	Indicates the number of SIP protocol requests that were routed using a routing table rule.	Number	
Database queue depth	Indicates the average number of database requests that are currently waiting to be run.	Number	The value of this measure should be less than 500. A value above 500 indicates that the backend database server is busy and is currently unable to process the requests.
Database queue latency	Indicates the average time spent by the database request in the database queue.	Secs	If for any reason resources (such as disk, memory, network, and processor) on the Back End Server are too high, the queue latency will also be high. Queue latency can also be high if there is too much network latency between the Front End Server and the Back End Server.
Sproc latency	Indicates the average time taken by the backend database server to process an sproc call.	Secs	

### 3.2.8 Mediation Service Details Test

The Mediation Server is considered the last point of contact for the Skype for Business environment before communicating to the telephony world for audio communication, whether its ingress or egress VoIP calls to the Public Switched Telephone Network (PSTN) world. The Mediation Server

is required for leveraging for inbound calls, outbound calls, and dial-in conferencing in a Skype for Business 2013 environment.

The Mediation Server listens for calls from both the server and the gateway. To listen for calls from the server, it uses Session Initiation Protocol (SIP) listening ports. Typically, port 5070 is used for Transport Layer Security (TLS) traffic coming from the Skype for Business server. To listen for calls from the gateway, it uses port 5067 for IP and PSTN traffic.

The Mediation Server is considered a Back-to-Back User Agent (B2BUA), which is responsible for handling communication between two endpoints in a SIP call. All SIP signaling will traverse through the Front End Server and Mediation Server when in route to a PSTN gateway. Only when media bypass is enabled will calls bypass the Mediation Server, which means that the calls go from the Skype for Business 2013 client straight to the PSTN gateway. When there's a PSTN gateway device local to where the call is being made, organizations will typically enable the Skype for Business environment for media bypass to avoid having the media traverse the WAN to make a call. This improves call quality because it reduces the chance for latency, jitter, and packet loss, which naturally occurs when a VoIP call travels through multiple hop points before the call reaches the endpoint receiving the call.

Some of the Mediation Server's key functions include:

- Encrypting and decrypting Secure Real-Time Transport Protocol (SRTP) data on the Skype for Business server side
- Translating SIP over TCP to SIP over mutual TLS (MTLS)
- Translating media streams between the Skype for Business server and the gateway peer of the Mediation Server

Using the **Mediation Server Details** test, it is possible to figure out how many audio channels are enabled with Phase Shift Modulation quality reporting and how many calls actually failed due to heavy load as well as due to the unexpected interaction from the proxy.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

## Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

## Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
PSM quality enabled audio channels	Indicates the number of audio channels that are currently enabled with Phase Shift Modulation (PSM) quality reporting.	Number	A low value is desired for this measure. Calculating PSM may cause processing overhead which in turn may directly affect the performance of the server. Therefore, administrators need to keep a check on the number of audio channels for which PSM reporting is enabled.
Calls that failed due to unexpected interaction from the Proxy	Indicates the number of calls that failed due to the unexpected interaction from proxy on the mediation server during the last measurement period.	Number	Ideally, the value of this measure should be zero.
Media connectivity check failures	Indicates the number of calls that failed when the media connectivity between the Mediation server and the remote endpoints could not be established.	Number	Ideally, the value of this measure should be zero.
Missing candidates	Indicates the number of times the media stack does not have media relay components.	Number	
Call failures due to heavy load	Indicates the scaled index between zero and 100 that	Number	

Measurement	Description	Measurement Unit	Interpretation
	is related to all call failures due to Global Health Index as a heavy load.		

### 3.2.9 Response Group Test

When a caller calls a response group, the call is routed to an agent based on a hunt group or the caller's answers to interactive voice response (IVR) questions. The Response Group application uses standard response group routing methods to route the call to the next available agent. Call routing methods include serial, longest-idle, parallel, round robin, and Attendant routing (that is, all agents are called at the same time for every incoming call, regardless of their current presence). If no agents are available, the call is held in a queue until an agent is available. While in the queue, the caller hears music until an available agent accepts the call. If the queue is full, or if the call times out while in the queue, the caller might hear a message and then is either disconnected or transferred to a different destination. When an agent accepts the call, the caller might or might not be able to see the agent's identity, depending on how the administrator configures the response group. Agents can either be formal, which means that they must sign in to the group before they can accept calls routed to the group, or informal, which means that they do not sign into and out of the group to accept calls. The Response Group application uses an internal service, called Match Making, to queue calls and find available agents. Each computer that runs the Response Group application runs the Match Making service, but only one Match Making service per Skype for Business Server pool is active at a time--the others are passive. In order to figure out the number of calls that were handled by the response group server pool, the number of calls received and the duration of the calls, use the **Response Group** test. This way, administrators may be proactively alerted to overload condition of the server pool and help them rectify the issue before end users are affected.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.

Parameters	Description
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current calls	Indicates the number of calls that are currently handled by the Response Group server pool.	Number	A high value for this measure indicates overload condition of the pool which may lead to performance degradation of the Skype for Business server. To avoid this, administrators may need to consider adding an additional Response Group pool.
Invoked SQL statements	Indicates the number of SQL statements that were invoked by the Response Group server pool during the last measurement period.	Number	
Received calls	Indicates the number of calls received by the server pool during the last measurement period.	Number	
Call duration	Indicates the total duration of the calls in the server pool.	Secs	
Active calls on the RGS	Indicates the number of calls that are currently active on the Response Group server pool.	Number	

### 3.2.10 SIP Peers Test

Session Initiation Protocol (SIP) is an ASCII-character-based signaling protocol designed for real-time transmission using Voice over IP (VoIP). In Speech Server, SIP establishes sessions with

requests and responses. When Speech Server is connected to a client, session initiation and call control are handled over SIP.

The SIP protocol is hosted through a SIP peer, which connects Speech Server to the caller's endpoint. The SIP peer can be a VoIP client, SIP client, or traditional telephony client. In Microsoft Skype for Business server, SIP Peer communication includes internal SIP communications between servers (MTLS), between Server and Client (TLS) and between Front End Servers and Mediation Servers (MTLS). Administrators may want to monitor the SIP Peer communication on the target Microsoft Skype for Business server so that they can figure out where exactly there is a delay in sending/receiving the messages. The **SIP Peers** test exactly helps you analyze the same! This test provides the administrators with a detailed insight on the numerical statistics of the active connections, active TLS connections, outstanding messages, connections that are flow controlled etc. In addition, this test helps you to figure out the time delay for each message in the outgoing queue and the rate at which outgoing messages were dropped.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

#### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

#### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Active connections	Indicates the number of established connections that are currently active.	Number	A connection is considered established when peer credentials are verified or the peer receives a 2xx response.
Active TLS connections	Indicates the number of established TLS (Transport)	Number	A TLS connection is considered established when the peer certificate,

Measurement	Description	Measurement Unit	Interpretation
	Layer Security) connections that are currently active.		and possibly the host name, are verified for a trust relationship.
Outstanding messages	Indicates the number of messages that are currently present in the outgoing queue.	Number	
Average delay in outgoing queue	Indicates the average time for which the messages were delayed in the outgoing queue.	Secs	A high value for this measure is a cause of concern which may be due to the accumulation of messages in the outgoing queue. Generally, the server may drop client connections if it is in a throttled state and therefore, messages may be delayed in the outgoing queue for more than 32 seconds.
Flow-controlled connections	Indicates the number of connections that are currently flow-controlled.	Number	Whenever the number of transactions on a specific connection exceeds the above threshold, the connection is marked as flow controlled and the server doesn't post any receives on it.
Flow-controlled connections dropped	Indicates the number of connections that are dropped due to excessive flow-control.	Number	A low value is desired for this measure.
Avg processing delay due to flow-control	Indicates the average time delay in message processing when the socket is flow-controlled.	Number	A low value is desired for this measure.
Incoming requests	Indicates the rate at which requests were received during the last measurement period.	Requests/Sec	This measure is a good indicator of the user load on the server.
Outgoings dropped rate	Indicates the rate at which outgoing messages timed out.	Timeouts/Sec	The outgoing messages may time out when they stay in the outgoing queue for a duration longer than required.

Measurement	Description	Measurement Unit	Interpretation
			A low value is desired for this measure.

### 3.2.11 Call Details Reporting Test

Call Detail Recording (CDR) enables you to track usage of the following:

- Peer-to-peer instant messaging sessions,
- Voice over Internet Protocol (VoIP) phone calls, and
- Conferencing calls.

This usage data includes information about who called whom, when they called, and how long they talked.

This test enables you to capture the failed messages during Call Details Recording and when exactly the failures occurred – is it during message validation? or in the MSMQ queue? Or when written to the database?

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Messages that failed	Indicates the number of	Number	A low value is desired for this

Measurement	Description	Measurement Unit	Interpretation
validation	messages that failed validation during Call Details Recording.		measure.
Messages dropped from MSMQ queue	Indicates the number of messages dropped from the MSMQ queue during Call Details Recording.	Number	A low value is desired for this measure.
Messages failed to be written to DB	Indicates the number of messages that were not written to the backend database server during Call Details Recording.	Number	

### 3.2.12 Quality Of Experience Test

Quality of Experience (QoE) records numeric data that indicates the media quality and information about participants, device names, drivers, IP addresses, and endpoint types involved in calls and sessions. Using this test, you can identify the number of reports that were dropped due to database insertion failures and the message queue/database transactions that failed.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Reports dropped due to database failure	Indicates the number of reports dropped due to database insertion failure during the last measurement period.	Number	Ideally, the value of this measure should be zero. A sudden/gradual increase in the value of this measure may be due to the unavailability of the database or due to the incompatibility of the versions between the database and data collection components.
Message transactions failed	Indicates the number of message queue/database transactions that failed during the last measurement period.	Number	If the value of this measure is abnormally high, then ensure that the SQL server is configured properly and started.

### 3.3 The Lync Conferencing Services Layer

Using the tests mapped to this layer, you can easily identify the conferencing activities occurring of the Microsoft Skype for Business server and figure out the following:

- The current health and draining state of the Application sharing Conference unit, the Data MCU server and the Instant Messaging multipoint control unit;
- Know how well the conferences are handled by the Application sharing conference unit
- How many RDP connections failed in the Application sharing Control unit?
- How many whiteboards were served by the Data MCU server and how many conferences are currently active on the Data MCU server?;
- How many conferences were active on the Instant Messaging multipoint control unit and how many message delivery failure notification messages were sent from the server?;
- Know how many add user and add conference responses failed in the server
- How many conferences were initiated fro each service and how many conference processing was actually delayed?;
- How many times the application endpoint creation failed for the Conferencing Auto Application service and how many incoming calls were actually incomplete?

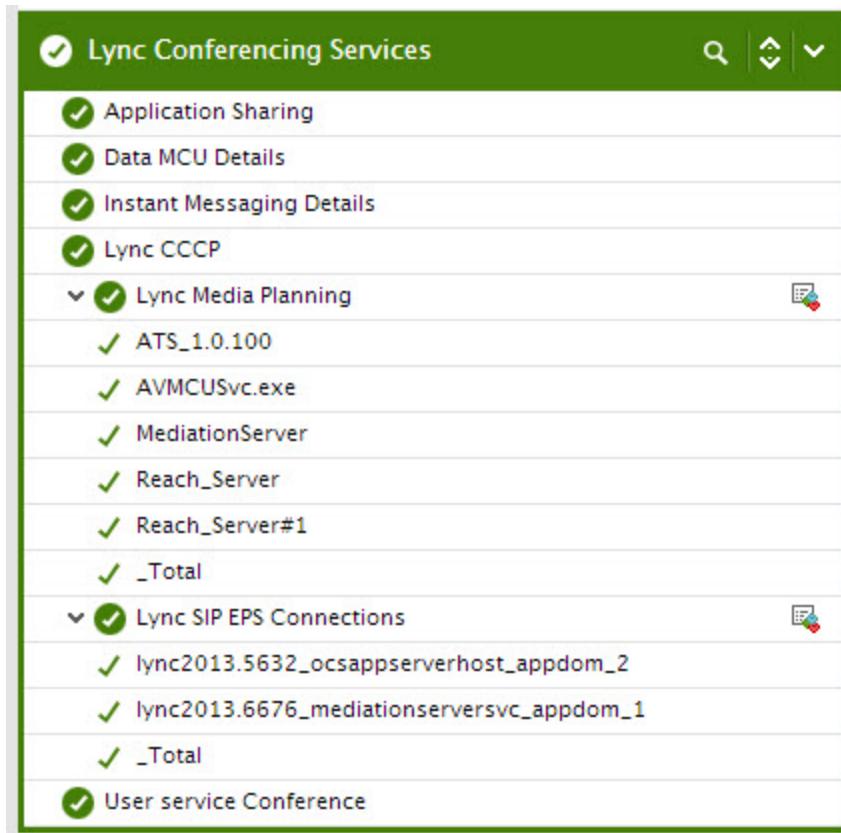


Figure 3.8: The tests mapped to the Lync Conferencing Services layer

Let us discuss each test of this layer (see Figure 3.8) in detail in the forthcoming sections.

### 3.3.1 Application Sharing Test

This test monitors the Application Sharing Messaging Conferencing Unit (ASMCU) and reports the current health and draining status of the server. In addition, this test helps you to figure out the number of users connected to the server, the number of active conferences, the number of active sessions and media connections. Using this test, administrators can identify the load on the server and the failures encountered during media transmission and RDP connections.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the Microsoft Skype for Business server that is to be monitored.

## Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

## Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation										
Current MCU health	Indicates the current health of the Skype for Business Application sharing Conference Unit (ASMcu).		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Normal</td> <td>0</td> </tr> <tr> <td>Loaded</td> <td>1</td> </tr> <tr> <td>Full</td> <td>2</td> </tr> <tr> <td>Unavailable</td> <td>3</td> </tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned Measure Values while indicating the current health of the ASMcu. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 3.</p>	Measure Value	Numeric Value	Normal	0	Loaded	1	Full	2	Unavailable	3
Measure Value	Numeric Value												
Normal	0												
Loaded	1												
Full	2												
Unavailable	3												
Draining state	Indicates the draining status of the ASMcu.		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Not requesting to drain</td> <td>0</td> </tr> </tbody> </table>	Measure Value	Numeric Value	Not requesting to drain	0						
Measure Value	Numeric Value												
Not requesting to drain	0												

Measurement	Description	Measurement Unit	Interpretation						
			<table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Requesting to drain</td><td>1</td></tr> <tr> <td>Draining</td><td>2</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned Measure Values while indicating the draining status of the ASMCu. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 2.</p>	Measure Value	Numeric Value	Requesting to drain	1	Draining	2
Measure Value	Numeric Value								
Requesting to drain	1								
Draining	2								
Estimated time to process all pending items in HTTP stack	Indicates the time taken to process all pending items in HTTP stack.	Secs	A low value is desired for this measure. An abnormally high value for this measure indicates that the server is processing an item that is unending.						
Active conferences	Indicates the number of active conferences in this server.	Number	A high value is desired for this measure.						
Connected users	Indicates the number of users currently connected to this server.	Number	This measure is a good indicator of the load on the server.						
Active sessions	Indicates the number of sessions that are currently active on this server.	Number							
Active media connections	Indicates the number of media connections that were active on this server.	Number							
Packet loss	Indicates the total number of packets lost in this server.	Number	Ideally, the value of this measure should be zero.						
Failed SIP dialogs	Indicates the number of SIP dialogs that failed in this server.	Number							

Measurement	Description	Measurement Unit	Interpretation
Allocated bandwidth	Indicates the amount of bandwidth allocated for this server during the last measurement period.	MB	
Media failures	Indicates the number of times media transmission timed out on this server.	Number	
Failed RDP connections	Indicates the number of times the RDP connections failed i.e., timed out on this server.	Number	A high value for this measure indicates connectivity issues which when ignored leads to the degradation of the server.

### 3.3.2 Data MCU Details Test

The Data MultiPoint Control Unit (Data MCU) on the front end server serves the Whiteboarding and presentation options to the users. When a Powerpoint presentation is uploaded to a meeting using Skype for Business, the file is actually uploaded to the Data MCU. This Data MCU returns a broadcast URL representing the uploaded file, causing the Skype for Business server to start an embedded browser file inside the client. The URL returned to the client is the URL of the Office Web Apps server, that retrieves the file from the Data MCU using the Web Open Platform Interface (WOPI) protocol and renders the slides into web-based presentable data. Every second, the browser frame pings the Office Web App server using AJAX to confirm which slide should be previewed. This information is sent back to Skype for Business 2013 Client's embedded browser frame, thus moving the presentation along. For the Skype for Business client to receive whiteboards and presentations without any technical glitch, it is essential to monitor the health of the Data MCU server. This is where the **Data MCU Details** test helps! Using this test, you can figure out the health and draining status of the Data MCU server along with the numerical statistics of the MCU users, active conferences and the whiteboards served by the Data MCU server.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

## Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

## Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation										
Session queues state	Indicates the state of the session queues.	Number	If the value reported by this measure is 1, then it indicates an overload condition.										
MCU users	Indicates the number of users in the Data MCU server.	Number											
Active conferences	Indicates the number of conferences that are currently active on the server.	Number											
DataMCU health state	Indicates the current state of the Data MCU server.		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Normal</td> <td>0</td> </tr> <tr> <td>Loaded</td> <td>1</td> </tr> <tr> <td>Full</td> <td>2</td> </tr> <tr> <td>Unavailable</td> <td>3</td> </tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the current health of the Data MCU server. However, in the graph of this measure, states will be</p>	Measure Value	Numeric Value	Normal	0	Loaded	1	Full	2	Unavailable	3
Measure Value	Numeric Value												
Normal	0												
Loaded	1												
Full	2												
Unavailable	3												

Measurement	Description	Measurement Unit	Interpretation								
			represented using the corresponding numeric equivalents i.e., 0 to 3.								
Draining status of the DataMCU	Indicates the current draining status of the Data MCU server.		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Not requesting to drain</td><td>0</td></tr> <tr> <td>Requesting to drain</td><td>1</td></tr> <tr> <td>Draining</td><td>2</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the draining status of the Data MCU server. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 2.</p> <p>When a server is drained, it stops taking new connections and calls. These new connections and calls are routed through other servers in the pool. A server being drained allows its sessions on existing connections to continue until they naturally end. When all existing sessions have ended, the server is ready to be taken offline.</p>	Measure Value	Numeric Value	Not requesting to drain	0	Requesting to drain	1	Draining	2
Measure Value	Numeric Value										
Not requesting to drain	0										
Requesting to drain	1										
Draining	2										
Annotations	Indicates the number of annotations i.e., whiteboards served by the Data MCU server.	Number									
Content objects	Indicates the number of content objects in the Data MCU server.	Number									

### 3.3.3 Instant Messaging Details Test

Instant messaging (IM) enables your users to communicate with each other in real time on their computers using text-based messages. Both two-party and multiparty IM sessions are supported. A participant in a two-party IM conversation can add a third participant to the conversation at any time. When this happens, the Conversation window changes to support conferencing features. Skype for Business and Communicator clients when involved in a one to one communication, is often referred to as peer-to-peer. Technically, the two clients are communicating in a one to one conversation, with the Instant Messaging multipoint control unit (IMMCU) in the middle. The IMMCU is a component of Front End Server. Placing the IMMCU in the required communication workflow allows call detail recording and other features that the Front End Server enables. Communication is from a dynamic source port on the client to the Front End Server port TLS/TCP/5061 (assuming the use of the recommended transport layer security). By design, peer-to-peer communication (as well as multiparty IM) is possible only when Skype for Business Server and the IMMCU is active and available. Therefore, administrators may constantly require to monitor the IMMCU. The **Instant Messaging Details** test helps the administrators accomplish this task. This test reports the current health and draining status of the IMMCU server and also helps administrators identify the number of active conferences, users connected to the server and the SIP connections that were throttled.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Current MCU health	Indicates the current		The values that this measure can report

Measurement	Description	Measurement Unit	Interpretation										
	health of the ImMCU server.		<p>and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Normal</td><td>0</td></tr> <tr> <td>Loaded</td><td>1</td></tr> <tr> <td>Full</td><td>2</td></tr> <tr> <td>Unavailable</td><td>3</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned Measure Values while indicating the current health of the ImMCU server. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 3.</p>	Measure Value	Numeric Value	Normal	0	Loaded	1	Full	2	Unavailable	3
Measure Value	Numeric Value												
Normal	0												
Loaded	1												
Full	2												
Unavailable	3												
Draining state	Indicates the current draining state of the ImMCU server.		<p>The values that this measure can report and their corresponding numeric values have been listed in the table below.</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Not requesting to drain</td><td>0</td></tr> <tr> <td>Requesting to drain</td><td>1</td></tr> <tr> <td>Draining</td><td>2</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the draining status of the ImMCU server. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 0 to 2.</p>	Measure Value	Numeric Value	Not requesting to drain	0	Requesting to drain	1	Draining	2		
Measure Value	Numeric Value												
Not requesting to drain	0												
Requesting to drain	1												
Draining	2												

Measurement	Description	Measurement Unit	Interpretation
			When a server is drained, it stops taking new connections and calls. These new connections and calls are routed through other servers in the pool. A server being drained allows its sessions on existing connections to continue until they naturally end. When all existing sessions have ended, the server is ready to be taken offline.
Throttled SIP connections	Indicates the number of SIP connections that were throttled on the server.	Number	If the value of this measure is greater than 1, then it indicates that the peer is not processing the requests on time which may be due to an overload condition of the peer.
Active conferences	Indicates the number of conferences that are currently active on the ImMCU server.	Number	
Connected users	Indicates the number of users currently connected to the ImMCU server.	Number	
IMDN failure rate	Indicates the rate at which the Instant Message Delivery Notification failure messages were sent from the server.	Sent/Sec	

### 3.3.4 Lync CCCP Test

Centralized Conference Control Protocol (C3P or CCCP) is used by protocol clients, front-end servers, and conferencing servers to establish and maintain the state of a conference. This test enables administrators to identify the number of CCCP messages in the send queue, the number of add user and add conference responses that failed, the messages received in the ImMCU and MCU.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Retried CCCP message rate	Indicates the rate at which the CCCP messages are retried.	Retries/sec	
Messages currently in the send queue	Indicates the total number of CCCP messages that are currently in the send queue.	Number	
Add-conference failed responses	Indicates the total number of add conference responses that failed.	Number	A low value is desired for this measure. A gradual/sudden increase in the value of this measure indicates connectivity issues.
Add-user failed responses	Indicates the total number of addUser responses that failed.	Number	
Audio/Video CCCP message received	Indicates the number of Audio/Video MCU messages received during the last measurement period.	Number	
Message received in MCUF	Indicates the number of MCUF CCCP messages received during the last measurement period.	Number	
Message received in	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
IMMCU	CCCP messages received in the Skype for Business Instant Messaging Conference Unit (ImMcu) during the last measurement period.		
Add-conference failed responses in IMMCU	Indicates the number of addConference responses that failed in the Skype for Business Instant Messaging Conference Unit (ImMcu).	Number	A low value is desired for this measure.

### 3.3.5 Lync Media Planning Test

This test auto discovers the services of the target Microsoft Skype for Business server and for each service, reports the number of times conferencing processing was delayed and the number of A/V conferences that were started. This way, administrators can identify the slowdowns in the conferences and rectify the issues before end users start complaining of the same!

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for each service of the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Delays in conference processing	Indicates the number of times the conference processing was delayed since the start of the service.	Number	Ideally, the value of this measure should be zero.
Started A/V conferences	Indicates the number of A/V conferences that have been initiated for this service.	Number	

**3.3.6 Lync SIP EPS Connections Test**

This test auto discovers the SIP EPS connections on the Microsoft Skype for Business server and for each connection, reports the number of DNS resolution failures encountered and the amount of data received.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
DNS resolution	Indicates the number of	Number	Ideally, the value of this measure

Measurement	Description	Measurement Unit	Interpretation
failures	DNS resolution failures encountered by this connection.		should be zero. The DNS resolution failure may happen when there are no records specified in DNS for the host name being resolved.
Received data	Indicates the amount of data received by this connection.	MB	

### 3.3.7 User service Conference Test

Every user in the Microsoft Skype for Business server is assigned a Registrar pool, backup registrar pool and a User Services pool. Administrators may want to determine how well the User Services pool is servicing the requests made to it. The **User service Conference** test exactly serves this purpose. This test helps administrators figure out the number of conferences in the pool, the Invite and ACK requests sent and received, the requests that were sent to the MCU server and the responses that failed, the number of conferences that were successful etc. This way administrators can identify the load on the user services pool and identify any notable performance issues before end users start complaining!

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Conferences scheduled in the pool	Indicates the number of conferences scheduled in the user services pool.	Number	
Rate of calls to the ConfJoinParticipant sproc	Indicates the rate at which the users joined the conference through sproc calls.	Calls/Sec	
Invite requests dispatched	Indicates the number of invite requests sent to the Focus component to join the conference.	Number	<p>The Focus is the conference state server. It is implemented as a SIP user agent that is addressable by using a conference URI. The Focus runs in the User Services module of all Front End Servers. All group IM, multiparty A/V, and data collaboration sessions are managed on the server by the Focus.</p> <p>The Focus is responsible for the following tasks:</p> <ul style="list-style-type: none"> <li>• Initiating conferences</li> <li>• Enlisting required conferencing servers</li> <li>• Authenticating participants before allowing them to enter a conference</li> <li>• Enforcing the policy that specifies whether the meeting organizer is authorized to invite external users</li> <li>• Maintaining SIP signaling relationships between conference participants and conferencing servers</li> </ul>

Measurement	Description	Measurement Unit	Interpretation
			<ul style="list-style-type: none"> <li>• Managing conference state</li> <li>• Accepting subscription to conferences and notifying users of changes in conference state, such as the arrival and departure of participants and the addition or removal of media</li> <li>• Maintaining and enforcing conference policies and rosters</li> </ul> <p>The Focus also enables the organizer to lock a meeting so that no more participants can enter after a certain time.</p>
Total ACK requests dispatched	Indicates the number of ACK requests dispatched to the Focus component.	Number	
Failed GetMcu requests	Indicates the number of requests failed to be sent to the MCU server.	Number	
Timed out GetMcu responses	Indicates the number of requests that timed out while waiting for a response from the MCU server.	Number	
Failed Create conference requests	Indicates the number of create conference requests that failed to reach the MCU.	Number	
Timed out Create conference requests	Indicates the number of create conference requests that timed out while waiting for a response from the MCU.	Number	

Measurement	Description	Measurement Unit	Interpretation
Conference latency	Indicates the average time taken to create a conference call.	Secs	A low value is desired for this measure.
Call latency	Indicates the average time taken to complete a MCU factory call.	Secs	
Allocation latency	Indicates the average time taken to complete a full allocation request.	Secs	
Successful conference requests	Indicates the number of add conference requests that were successful.	Number	A high value is desired for this measure.
Media specific command requests	Indicates the number of requests received for media specific commands.	Number	
C3P failure responses	Indicates the number of CCCP failure responses that were generated locally.	Number	Ideally, the value of this measure should be zero.
GetConference requests	Indicates the number of getconference requests received by this service.	Number	
Conference notifications in processing	Indicates the number of notifications for the conferences that are in process.	Number	

### 3.3.8 LS CAA Test

This test monitors the Conferencing Auto Application (CAA) service of the Microsoft Skype for Business server and reports the number of times the application endpoint creation failed along with the rate of incomplete incoming calls.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

#### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

#### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total application endpoint creation failures	Indicates the total number of times the application endpoint creation failed during the last measurement period .	Number	Ideally, the value of this measure should be zero. A gradual/sudden increase in the value of this measure is a cause of concern which may be due to deletion of the contact object or due to the application not been properly activated. Deactivating the application from the pool and reactivating it may sometimes help administrators to contain the shoot up of this measure.
Incomplete calls	Indicates the number of calls that were incomplete per second.	Calls/Sec	

## 3.4 The Lync Web Services Layer

Using the metrics reported, administrators can find quick and accurate answers for the following performance questions:

- Observe address book accesses, measure the speed of these accesses, and report slowdowns;
- Monitor the message processing ability of the server and report deficiencies;
- How many SIP connections failed and how many are currently active?;
- How long the distribution list take to process the requests?

- How many times the Join launcher service failed and how many incoming requests were received by the service?
- How many replication requests were received by the replication service and how long does the service take to process?;
- Continuosly monitor the mobility services and figure out the successful Get Location requests, the push notifications that failed and throttled;
- How many times the user provisioning and publish calls failed?



Figure 3.9: The tests mapped to the Lync Web Services layer

Let us see details of each test of this layer in the upcoming sections.

### 3.4.1 Address Book Query Test

The Address Book server provides global address list information from Active Directory Domain Services to Skype for Business Server clients. It also retrieves user and contact information from the RTC database, writes the information to the Address Book files, and then stores the files on a shared folder where they are downloaded by Skype for Business clients. The Address Book Server writes the information to the RTCAb database, which is used by the Address Book Web Query service to respond to user search queries from the Microsoft Skype for Business mobile application. It optionally normalizes enterprise user phone numbers that are written to the RTC database for the purpose of provisioning user contacts in Skype for Business. The Address Book service is installed by default on all Front End Servers. The Address Book Web Query service is installed by default with the Web services on each Front End Servers. Administrators may often want to figure out how well the address book is utilized by the Skype for Business clients and how well the address book is servicing the requests received by it. This is where the **Address Book Query** test helps!

Using this test, you can easily figure out the rate at which the address book is searched upon receiving a request, the time taken to process the requests, the requests that were successful and

the requests to the address book that failed. This way, administrators can easily analyze the effectiveness of the address book and if there are too many failed requests, administrators can figure out that the address book is outdated and requires immediate updation.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Rate of failed search requests	Indicates the rate at which the requests failed to search the address book.	Requests/Sec	A high value for this measure indicates that the address book is outdated and required to be updated.
Avg processing time for search requests	Indicates the average time taken to search for a contact from the address book upon a user request.	Secs	Backend Database Performance issues may occur if the value of this measure is greater than 0.1 second. In order to optimize this measure, verify the CPU load on the Database of the Skype for Business server and upgrade the hardware of the server if needed.
Address book search requests rate	Indicates the rate at which the requests to search the address book are processed.	Requests/Sec	
Rate of failed file requests	Indicates the rate at which the requests failed to fetch	Requests/Sec	A very high value of this measure may be due to authentication issues in the

Measurement	Description	Measurement Unit	Interpretation
	the files from the address book.		server or network connectivity issues. You may check the Skype for Business Server Event log to identify the exact cause of the request failure and rectify the issue.
Rate of succeeded file requests	Indicates the rate at which the requests to fetch the files from the address book were successful.	Requests/Sec	A high value is desired for this measure.

### 3.4.2 Communicator Web Application Test

If data collaboration is enabled in your conferencing policy then you (and anyone else who participates in a meeting that you organized on the Microsoft Skype for Business server) will have the following three capabilities at your disposal: PowerPoint Presentation, New Whiteboard, and New Poll. This test helps you to figure out the data collaboration authentication requests that failed, the SIP connections that failed and the data collaboration client connections that were closed due to throttling. In addition, you can figure out the number of active SIP connections and the amount of data that is queued for data collaboration. This way, administrators may figure out how well the data collaboration enabling is helpful for users to access powerpoints, polls etc.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Rate of failed data collaboration authentication requests	Indicates the rate at which data collaboration authentication requests failed.	Requests/Sec	
Data collaboration connection failures	Indicates the number of data collaboration connections that failed on the server.	Number	
Rate of SIP connection failures	Indicates the rate at which the SIP connections failed on the server.	Failures/Sec	A low value is desired for this measure. A sudden/gradual increase in the value of this measure indicates that the SIP connection was terminated unexpectedly due to network issues or due to the unavailability of remote servers.
Rate of data collaboration client connections closed due to throttling	Indicates the rate at which the data collaboration client connections were closed due to throttling on this server.	Conns/Sec	A low value is desired for this measure. Often data collaboration of the client connections failed due to the inability of the client to read data efficiently on the server. This may be due to the failure of network or organized attack.
Queued data for data Collaboration server connections	Indicates the amount of data that is queued for data collaboration server connections.	MB	A high value for this measure indicates the failure to read outgoing data by the Web Conferencing servers.
Active SIP connections	Indicates the number of SIP connections that are currently active.	Number	

#### 3.4.3 Expand Distribution List Test

Using Skype for Business, service users can search the address book for other users in their organization. They can also add email addresses and distribution lists to their Contacts lists. Instant messages can be sent to individual members of a distribution list or the entire distribution list.

Distribution lists can be expanded to see the individual members if there are 100 or fewer individual members (regardless of the members' presence or sign-in status). This test monitors the Distribution list and reports the time taken to process the requests, the number of requests that passed user input validation and the rate at which the AD requests time out. This way, administrators may oversee the effectiveness of the distribution list and if there are too many timed out requests, they can easily figure out the problematic list and rectify the same before users start complaining.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

#### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

#### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Request processing time	Indicates the average time taken to process a request successfully.	Secs	A high value for this measure may indicate performance issues with the Active Directory servers such as CPU overload.
Rate of SOAP exceptions	Indicates the rate at which SOAP exceptions occurred.	Exceptions/Sec	
User input validated requests	Indicates the number of requests that passed user input validation.	Number	
Timed out AD requests rate	Indicates the rate at which Active Directory requests timed out.	Requests/Sec	

### 3.4.4 Join Lancher Test

To enhance the quality of conferencing, the Microsoft Skype for Business server has introduced a new join launcher feature. The Skype for Business server updates the join launcher to validate each meeting before launching a client. When the meeting URL has been activated the Skype for Business Server runs the join launcher service. The main steps performed by the join launcher are:

- Determine language to use based on preferred language(s) in the browser and the language packs installed on the server. If no direct or indirect match (parent) can be found between preferred language(s) and the language packs installed on the server the server language will be used;
- It will convert the SimpleUrl based URL to the real URI used to join the meeting;
- Determine if the web request came from the Internet or from the internal network by looking at which web site was used;
- Determine which client to launch based on the installed client(s).

Using this test, administrators can figure out how many times the join launcher service failed and how many incoming requests were received. This way, administrators can be proactively alerted to performance/connectivity issues in the Microsoft Skype for Business server.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Join failures	Indicates the number of times the join launcher service failed.	Number	Ideally, the value of this measure should be zero. A sudden/gradual increase in the value of this measure indicates that the join launcher service is not initiated properly or there is a performance/connectivity issue in the Skype for Business server.
Incoming join requests	Indicates the number of incoming join requests i.e., meeting validation requests to the join launcher service.	Number	

### 3.4.5 LS Replications Test

This test monitors the replication service of the Skype for Business server and reports the number of replication requests received, the number of status update calls received along with the time taken to receive a response for a request.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Request latency	Indicates the average time taken to receive a response for a request.	Secs	
Sproc latency	Indicates the average time taken to process a sproc call.	Secs	
Received replication requests	Indicates the number of replication requests received by this service during the last measurement period.	Number	
Status update calls by Replica Replicator Agent	Indicates the number of status update calls received by the Replica replication agent.	Number	

### 3.4.6 LS Web Information Test

To support Skype for Business mobile applications on mobile devices, Skype for Business Server 2013 provides three services: Skype for Business Server 2013 Mcx Mobility Service, Skype for Business Server 2013 Autodiscover Service, and Skype for Business Server 2013 Push Notification Service.

The services that support mobility are as follows:

- **Skype for Business Server 2013 Unified Communications Web API (UCWA)** Provides services for real-time communications with mobile and web clients in Skype for Business Server 2013. When you deploy the Cumulative Updates for Skype for Business Server 2013, the installation creates a virtual directory in the internal and external web services (Ucwa). A web component that is part of the Ucwa virtual directory accepts calls from UCWA-enabled clients. The client apps communicate over a REST interface for presence, contacts, instant messaging, VoIP, video conferencing, and collaboration. UCWA uses a P-GET based channel to send events, such as an incoming call, incoming instant message, or a message to the client app.

- **Skype for Business Server 2013 Mobility Service (Mcx)** This service supports Skype for Business functionality, such as instant messaging (IM), presence, and contacts, on mobile devices. The Mobility Service is installed on every Front End Server in each pool that is to support Skype for Business functionality on mobile devices. When you install Skype for Business Server 2013, a new virtual directory (Mcx) is created under both the internal website and the external website on your Front End Servers.
- **Skype for Business Server 2013 Autodiscover Service** This service identifies the location of the user and enables mobile devices and other Skype for Business clients to locate resources—such as the internal and external URLs for Skype for Business Server 2013 Web Services, and the URL for the Mcx or UCWA—regardless of network location. Automatic discovery uses hardcoded host names (lyncdiscoverinternal for users inside the network; lyncdiscover for users outside the network) and the SIP domain of the user. It supports client connections that use either HTTP or HTTPS.
- The Autodiscover Service is installed on every Front End Server and on every Director in each pool that is to support Skype for Business functionality on mobile devices. When you install the Autodiscover Service, a new virtual directory (Autodiscover) is created under both the internal website and the external website, on both Front End Servers and Directors.
- **Push Notification Service** This service is a cloud-based service that is located in the Skype for Business Online data center. When the Skype for Business mobile application on a supported Apple iOS device or Windows Phone is inactive, it cannot respond to new events, such as a new instant messaging (IM) invitation, a missed instant message, a missed call, or voice mail, because these devices do not support mobile applications running in the background. In these cases, a notification of the new event—called a **push notification**—is sent to the mobile device. The Mobility Service sends the notification to the cloud-based Push Notification Service, which then sends the notification either to the Apple Push Notification Service (APNS) (for supported Apple iOS devices) or to the Microsoft Push Notification Service (MPNS) (for Windows Phone), which then sends it on to the mobile device. The user can then respond to the notification on the mobile device to activate the application.

Whenever users login to the Microsoft Skype for Business server from their Skype for Business mobile client, they must be able to do so without any technical glitch. The **LS Web Information** test helps administrators continuously monitor the mobility services of the Microsoft Skype for Business server and proactively alert them on failures encountered by the users when logging into the server. Using this test, administrators can figure out the following:

- The successful Get Location requests;
- The time taken to process the Get Location requests;

- The Get Location requests that failed;
- The PIN calls that were validated;
- The Push notification that failed and throttled;
- The requests to the Mobile Communication service that failed and rejected.

This way, administrators can be proactively alerted to issues faced by the users logging into the Skype for Business server through mobile clients before end users start complaining about the inaccessibility of the server.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Failed local Lookup user requests	Indicates the number of local lookup user requests that failed.	Number	
Successful GetLocation request rate	Indicates the rate at which Get Location requests were processed successfully.	Requests/Sec	The maximum value that is acceptable for this measure is 12. Requests to initiate a session consume the most CPU on the server. Peak supported load is 12/second. Sustainability depends on other loads on the server. Initiate a session typically means a sign-in for a user

Measurement	Description	Measurement Unit	Interpretation
			that has been signed out for an extended period of time.
Avg processing time for a successful GetLocation request	Indicates the average time taken to successfully process a Get Location request to retrieve a city.	Secs	
Failed GetLocation request rate	Indicates the rate at which Get Location requests failed.	Requests/Sec	
Received requests	Indicates the number of requests received for all device specific updates during the last measurement period.	Number	<p>Device Update Web service supports Skype for Business Phone Edition.</p> <p>The first time a user starts Skype for Business Phone Edition and signs in, the device gets in-band provisioning information from the server or Front End pool hosting the device user account. The information contains the internal and external URL of the server running Device Update Web service.</p> <p>If the device is turned on, but no user signs on and no user has ever previously signed on to the device, the device sends a Domain Name System (DNS) lookup request to ucupdates-r2.&lt;DNS domain name that was provided by DHCP&gt; and obtains the internal and external URL of the server running Device Update Web service.</p> <p>Thereafter, when the device is turned on, when the user signs in, and every 24 hours by default, Skype for Business Phone Edition checks for updates by sending an HTTP request over port 443 to the server hosting Device Update Web service. The request includes the current version</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>that Skype for Business Phone Edition is running.</p> <p>If a user has not signed on to the device, and the device is outside the organization's firewall, Device Update Web service denies the request.</p> <p>Otherwise, the Device Update Web service returns a response that contains one of the following:</p> <p>If there are no approved updates for the current version of the firmware, the response contains downloads=0. For test devices, updates must be pending rather than approved for this to occur.</p> <p>If there is an approved update for the current version, the response contains an internal and external URL for Device Update Web service. For test devices, updates must be pending rather than approved for this to occur.</p> <p>In the latter case, Skype for Business Phone Edition sends an HTTPS update request over port 443 to Device Update Web service. If the device is outside the organization's firewall, this request must include a user name and password for a valid user account, which means a user must be logged on to the device.</p> <p>The update image is downloaded to the device.</p> <p>The device waits for five minutes of idle activity, and then restarts to complete the update.</p>
Resolve user calls to	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
the web service	resolve user calls to the web service during the last measurement period.		
Received HTTP UPRP messages	Indicates the number of HTTP UPRP messages received during the last measurement period.	Number	
Validate pin calls to the pin auth provider	Indicates the number of PIN calls that were sent to the PIN auth provider for validation during the last measurement period.	Number	
Failed validate cert calls to the cert auth provider	Indicates the cert calls to the cert auth provider that failed validation.	Number	Ideally, the value of this measure should be zero.
Push notifications failure rate	Indicates the rate at which the Push notifications failed.	Failures/Sec	<p>Push notification refers to the sound alerts, on-screen alerts (text), and badges that are pushed by applications to the Apple iPhone, iPad, and Windows Phone, when the mobile device is inactive.</p> <p>Nokia Symbian, Android based Skype for Business Mobile clients and Skype for Business Lobile clients for Apple devices do not use Push notifications. Therefore the value of this measure exclusively indicates the failure of push notifications for the Windows based phones alone. Ideally, the value of this measure should be zero.</p>
Push notifications throttling rate	Indicates the rate at which the push notifications were throttled.	Requests/Sec	
Failed requests rate	Indicates the rate at which requests to the Mobile Communication Service failed.	Failures/Sec	Ideally, the value of this measure should be zero.

Measurement	Description	Measurement Unit	Interpretation
Rejected requests rate	Indicates the rate at which requests to the Mobile Communication Service were rejected.	Rejects/Sec	

### 3.4.7 Lync Provisioning Test

This test monitors the Provisioning service of the Microsoft Skype for Business server and reports the number of times user provisioning failed and the number of publish calls that failed.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Provision failures	Indicates the number of times user provisioning failed in the server.	Number	Ideally, the value of this measure should be zero.
Failed publish calls	Indicates the number of publish calls that failed in the server.	Number	Ideally, the value of this server should be zero.

## 3.5 The Lync User Services Layer

This layer helps you identify the following information:

- Monitor the backend database server and figure out how well requests are processed and how long does it take to process the requests?;
- How well the stored procedure calls were executed by the user service?;
- How many stale items were in the queue of the Lync storage service and how much space did the service utilize to store the items to the backend database server?;
- For each client version, how many unique users were actually connected?;
- How many users were actually connected to the server?
- For how many users were voice calls enabled?
- How many times the HTTPS connections failed?



Figure 3.10: The tests mapped to the Lync User Service layer

Let us discuss each test of this layer (see Figure 3.10) in the forthcoming sections.

### 3.5.1 LS DB Store Test

Microsoft Skype for Business server performance is affected by various factors such as user profiles, system architecture, software, hardware components, third- party integration points such as gateways and telephony equipment, network connectivity and performance, Windows Active Directory service configuration and performance in addition to the Windows operating system functionality. At the core of the Skype for Business server deployments' performance is the server software and hardware it is implemented on. As an example, a front-end server must have sufficient hardware resources to cope with the expected (short-term) user load. If a respective front-end server is required to provide services to 10 thousand users, then an adequately configured server must meet the expected load requirements to ultimately help ensure the best possible end-user experience.

Monitoring server performance is therefore extremely important to gauge whether the implemented server infrastructure have suitable hardware resources for the day-to-day peak-load requirements.

Monitoring server performance helps identify system bottlenecks allowing administrators to apply corrective action before the end-user experience is affected. The performance data obtained is then used for long-term capacity planning. For a server performance to be uniform, it is mandatory to track the performance of the backend SQL server database server used by the front end pool of the Microsoft Skype for Business server deployment. This is exactly where the **LS DB Store** test helps! This test monitors the backend database server and proactively alerts the administrators if the following exceeds/decreases far beyond the optimal range:

- The time spent by the requests in the queue while waiting to contact the backend server;
- The time taken by the backend server to process a request;
- The rate at which the requests were throttled, etc.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

#### Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

#### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Queue depth	Indicates the number of requests that are waiting in the queue to be run.	Number	The value of this measure should be less than 500. A value above 500 indicates that the backend database is busy and is unable to respond to requests quickly.
Queue latency	Indicates the time spent by a request in the queue while waiting to contact	Secs	A low value is desired for this measure. A gradual/sudden increase in

Measurement	Description	Measurement Unit	Interpretation
	the backend server.		the value of this measure indicates that the resource utilization of the backend server is high. Also, there may be a high network latency between the frontend and backend servers.
Sproc latency	Indicates the time taken by the backend server to process a request.	Secs	A low value is desired for this measure. If both the Sproc latency and Queue latency measures are high, then it indicates a bottleneck in the SQL Server back-end database. The disk IO contention and the SQL Server database active transactions are good indicators to identify the database that is having more contention than other databases.
Throttled requests rate	Indicates the rate at which the requests were rejected after being retried.	Request/Sec	The value of this measure will be high if the value of the Queue latency measure is abnormally high.

### 3.5.2 LS User Services Test

This test monitors the user service of the Microsoft Skype for Business server and reports how well the stored procedure calls were made, how many notifications and subscribes were sent through this service, how many active endpoints were connected etc. Using this test, administrators may figure out the slowdowns if any, in the user service by identifying the stored procedures that are taking too long to execute and take remedial measures to sort out the issue that is causing the slowdown.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

## Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

## Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
RPC calls currently in progress	Indicates the number of RPC calls that are currently in progress in the User Services RPC component.	Number	
Rate of calls to the CertStoreGetPublishedCert sproc	Indicates the rate at which calls were made to the CertStoreGetPublishedCert sproc.	Calls/Sec	
Rate of calls to the ExpireEndpoint sproc	Indicates the rate at which calls were made to the ExpireEndpoint sproc.	Calls/Sec	
Notifications sent	Indicates the number of NOTIFY notifications sent by the Notify component of the Skype for Business User Services module.	Number	
Aggregation requests received	Indicates the number of aggregation requests received since the start of the User services.	Number	
Failed publisher lock requests	Indicates the number of publisher lock requests that have failed in the module.	Number	
Number of times primary cluster was detected as down	Indicates the number of times primary cluster was detected as down.	Number	

Measurement	Description	Measurement Unit	Interpretation
Data loss events with state change	Indicates the number of data loss events that were generated due to state change of the front end server.	Number	
Data loss events without state change	Indicates the number of data loss events that were generated even when there was no state change detected on the front end server.	Number	
Replication operation failed rate	Indicates the rate at which replication operations failed on this service.	Number	
Is server connected to fabric pool manager?	Indicates whether/not the server is connected to the fabric pool manager.	Number	If the value of this measure is 1, then it indicates that the server is connected to the fabric pool manager.
Rate of calls to the RtcSetPresence sproc	Indicates the rate at which the RtcSetPresence stored procedure call was executed.	Calls/Sec	
Dispatched subscribes	Indicates the number of SUBSCRIBEs sent to the Subscribe component.	Number	
Active endpoints connected	Indicates the number of active endpoints connected to the server.	Number	
Publications rate	Indicates the number of category publications per second.	Publications/Sec	

### 3.5.3 Lync LDM Info Test

The Lync Dialog Manager (LDM) service enables logging of Web Conferencing Server Connection Management which manages TLS/MTLS connections from clients and other servers such as the

Web Conferencing Edge Server. The Lync Dialog Listener (LDL) module, when installed on every frontend server, tracks all INVITE/BYE-Commands and parses the SDP to find out the negotiated candidates. These commands/messages are sent to the Lync Dialog Manager (LDM) using a HTTP-PORST. This is another service, which has to be installed somewhere within your network. The LDM will then talk to a network management backend to deliver the information. Whenever large number of messages are delivered, it becomes tedious for the administrators to keep track of the messages moving to and from the LDM. This is where the **Lync LDM Info** test helps! This test monitors the LDM server and reports the number of TCP connections that are active and the number of messages received by the server. Using this test administrators can easily figure out overload conditions of the server proactively before any serious performance issue is detected.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Active TCP connections	Indicates the number of active TCP connections on the server.	Number	
Buffer allocation call rate	Indicates the number of buffer allocation calls per second.	Buffers/Sec	
Received messages	Indicates the number of messages that were received on the server during the last measurement period.	Number	This measure is a good indicator on the load of the server.

### 3.5.4 LYSS Details Test

Skype for Business Storage Service (LYSS) is a storage framework in Skype for Business Server 2013 intended to be used by different Skype for Business Storage Service consumers for accessing storage platforms in a Skype for Business Server 2013 system. Using this test, you can figure out the number of stale items in the queue of the service and the percentage of space used by the service to store items in the backend database.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Storage service queue items	Indicates the number of stale items in the queue of this storage service.	Number	A low value is desired for this measure. If the value of this measure is abnormally high, say for example 50000, then you should run the CleanUpStorageServiceData.exe tool from the Skype for Business Server 2013 Resource Kit which will delete all orphaned data from the pool.
Space used by storage service database	Indicates the percentage of space used by this storage service to store the items in the backend database.	Percent	

### 3.5.5 Lync Client Connections Test

Users connecting to the Microsoft Lync server may not often have the latest version of Skype for Business client installed on their systems. When issues relating to licenses, upgrades etc crop up, administrators may often find it difficult to figure out the users who are connected with different client versions. The **Lync Client Connections** test helps administrators figure out the same! Using this test, administrators can figure out the number of unique users who are currently connected to the server through each client version. This way, you can figure out the client version that is most commonly used by the users.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for each Skype for Business client version connecting to the target Microsoft Skype for Business server being monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

**Measurements made by the test**

Measurement	Description	Measurement Unit	Interpretation
Client connections	Indicates the total number of unique users currently connected to the server through this client version.	Number	<p>Comparing the value of this measure across versions will help you identify the client version that is most commonly used by the users.</p> <p>The detailed diagnosis of this measure lists the users who are currently connected to the Skype for Business server through this client version.</p>

### 3.5.6 Lync Users/Clients Test

This test monitors the interaction of the users through the Skype for Business client with the Microsoft Skype for Business server and reports the following:

- The number of users connected to the server;
- The number of users enabled on the server and the users for whom voice call is enabled;
- The number of Skype for Business client versions connected to the server;

Using this test, administrators can identify the load on the server and proactively take remedial measures to balance the load before end users realize connectivity issues.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current users	Indicates the number of users currently connected to the server.	Number	This value is a good indicator of the load on the server.  The detailed diagnosis of this measure if enabled, lists the users who are currently connected to the server.
Enabled users in server	Indicates the number of users enabled on the server.	Number	
Voice enabled users	Indicates the number of voice call enabled users on the server.	Number	
Percentage of enabled users connected	Indicates the percentage of users currently connected to the server.	Number	
Client versions connected	Indicates the number of clients of different version types connected to the server.	Number	

### 3.5.7 Lync HTTPS Transport Details Test

This test reports the rate at which the HTTPS connection attempts failed on the target Microsoft Skype for Business server.

**Target of the test :** A Microsoft Skype for Business server

**Agent deploying the test :** An internal/remote agent

**Outputs of the test :** One set of results for the target Microsoft Skype for Business server that is to be monitored.

**Configurable parameters for the test**

Parameters	Description
Test Period	How often should the test be executed.

Parameters	Description
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. The default port is 5060.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Rate of failed connection attempts	Indicates the rate at which HTTPS connection attempts failed.	Attempts/sec	A low value is desired for this measure.

## About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit [www.eginnovations.com](http://www.eginnovations.com).

### Contact Us

For support queries, email [support@eginnovations.com](mailto:support@eginnovations.com).

To contact eG Innovations sales team, email [sales@eginnovations.com](mailto:sales@eginnovations.com).

Copyright © 2018 eG Innovations Inc. All rights reserved.

This document may not be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the prior written consent of eG Innovations. eG Innovations 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. The information contained in this document is subject to change without notice.

All right, title, and interest in and to the software and documentation are and shall remain the exclusive property of eG Innovations. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.