Monitoring Tibco EMS Server

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

TIBCO Enterprise Message Service (EMS) is fully compliant Java Message Service (JMS) implementation from TIBCO with some enterprise-class enhancements. The Java Message Service makes it easy to write business applications that asynchronously send and receive critical business data and events.

Issues in the performance of the Tibco EMS can therefore obstruct the exchange of critical data across the environment, stalling key business transactions in the process and causing significant loss of revenue and reputation. It is therefore imperative to monitor the Tibco EMS.

eG Enterprise provides end-to-end monitoring of an IT infrastructure. While monitoring a the Tibco EMS sever, eG Enterprise analyzes the performance of the server in the light of the relationship it shares with other components in the environment, and accordingly determines the state of the server. This approach to monitoring enables eG Enterprise to accurately pinpoint the root-cause of problems that might occur in the environment.

This document focuses on the monitoring capabilities of eG Enterprise with regard to the Tibco EMS server.

Chapter 2: How does eG Enterprise Monitor Tibco EMS?

This chapter guides users in configuring and manage the Tibco EMS servers to work with eG Agent

Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are as follows:

```
tcp://{IPAddressorHostName_of_TibcoEMS}:{PortNumber_of_TibcoEMS} show server
tcp://{IPAddressorHostName_of_TibcoEMS}:{PortNumber_of_TibcoEMS} show durables
tcp://{IPAddressorHostName_of_TibcoEMS}:{PortNumber_of_TibcoEMS} show queues
tcp://{IPAddressorHostName_of_TibcoEMS}:{PortNumber_of_TibcoEMS} show topics
tcp://{IPAddressorHostName_of_TibcoEMS}:{PortNumber_of_TibcoEMS} show connections
```

For instance, if the IP address of your Tibco EMS server is 192.168.10.28 and its port is say, 9090, then a sample command in the .bat or .sh file would be:

```
tcp://192.168.10.28:9090 show server
```

The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the **COMMANDPATH** text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.

2.1 Managing the Tibco EMS Server

To manually add the Tibco EMS Server component for monitoring. Remember that the eG Enterprise automatically manages the components that are added manually. To manage a NexentaStor 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 Tibco EMS as the **Component type**. Then, click the **Add New Component** button. This will invoke Figure 2.1.

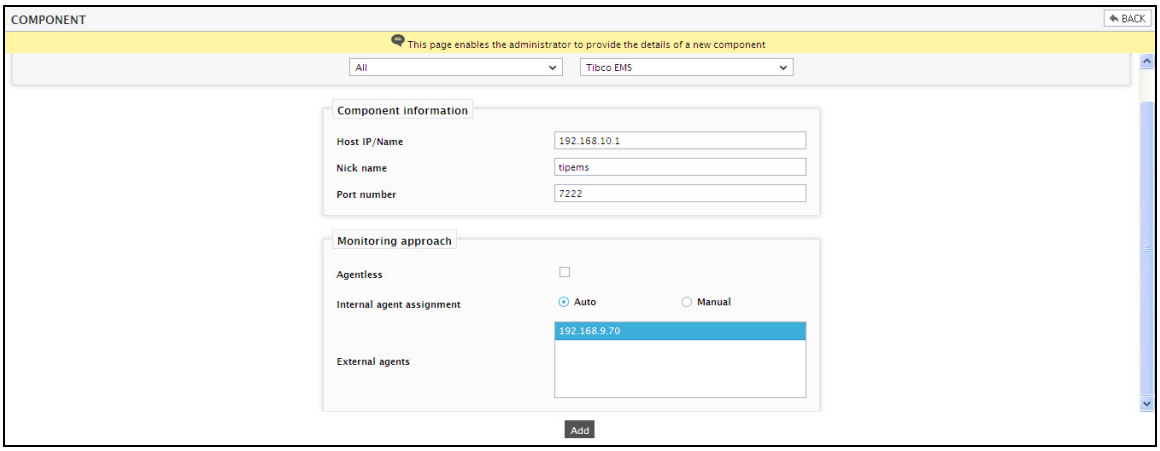

Figure 2.1: Adding a Tibco EMS Server for monitoring

3. Specify the **Host IP/Name** and the **Nick name** of the Tibco EMS in Figure 2.1. Then, click the **Add** button to register the changes.
4. Then, try to sign out of the eG administrative interface. Doing so will invoke a list of unconfigured tests for the Tibco EMS server.

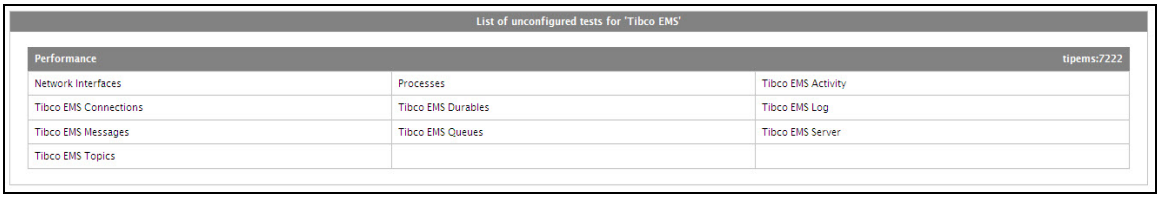

Figure 2.2: The list of unconfigured tests for the Tibco EMS server

5. Click on any test in the list of unconfigured tests to configure it. To know how to configure these tests, refer to **Monitoring Tibco EMS Servers**.
6. Then, proceed to sign out again. This time you will be prompted to configure the **Network Interfaces** and **Processes** tests of the target Tibco EMS server. Configure these tests one after another. The details on configuring these tests have been discussed in the *Monitoring Unix and Windows Server* document.
7. Finally, signout of the eG administrative interface.

Chapter 3: Monitoring Tibco EMS Servers

eG Enterprise offers a dedicated Tibco EMS monitoring model that continuously monitors the Tibco EMS, and proactively alerts administrators to potential performance bottlenecks.

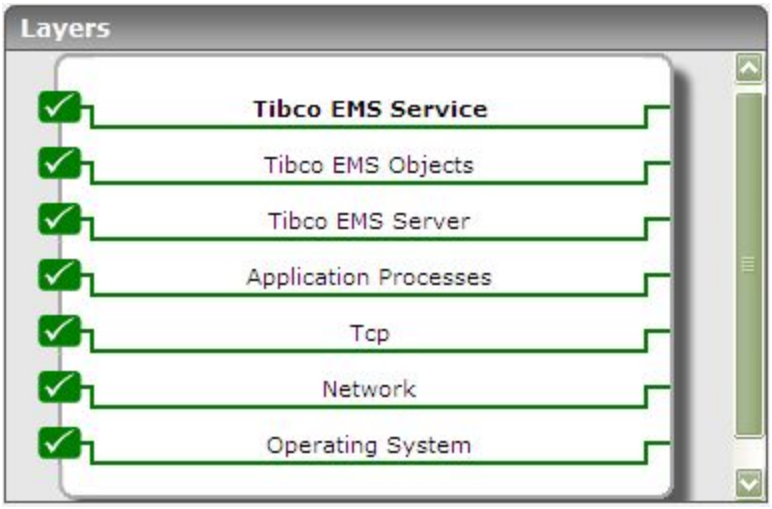

Figure 3.1: The layer model of the Tibco EMS Server

Using the metrics reported by this model, the following questions can be accurately answered:

- What is the current mode of operation of the server - active, inactive, or fault-tolerant standby mode?
- What is the total queue and topic load imposed by applications on the EMS Server? What type of queues and topics are maximum on the server?
- Are too many connections to the server idle?
- How busy is the EMS server in terms of the number of sessions it is handling?
- Are there any messages for durable subscriptions on the server?
- How many applications are sending messages and retrieving messages from the server?
- Are there too many pending messages on the server? Which queue/topic has the maximum number of pending messages?
- Are the pending messages too heavy? Which queue/topic contains the heaviest pending messages?
- Is there very little free message memory on the server?

- Which queue/topic on the server is static? Which queues/topics are dynamic or temporary?
- Which user to the server has the maximum number of sessions and connections open on the server?
- Are any durable subscribers offline currently?
- Are there too messages awaiting delivery to any durable subscriber?
- The heaviest pending message is meant for which durable subscriber?
- Has the log file utilized its maximum allocated space?

The sections that will follow discuss each of the top-3 layers of the layer model, as the other layers have already been discussed in *Monitoring Unix and Windows Servers* document.

3.1 The Tibco EMS Server Layer

The tests mapped to this layer reveal the overall health of the Tibco EMS server, and promptly alert administrators to the following problem conditions:

- Excessive space usage by the server log files;
- A large number of pending messages on the server;
- Excessive usage of message memory;
- Too many idle connections on the server;
- A server overload caused by a large number of topics/queues on the server;
- Non-availability of the server

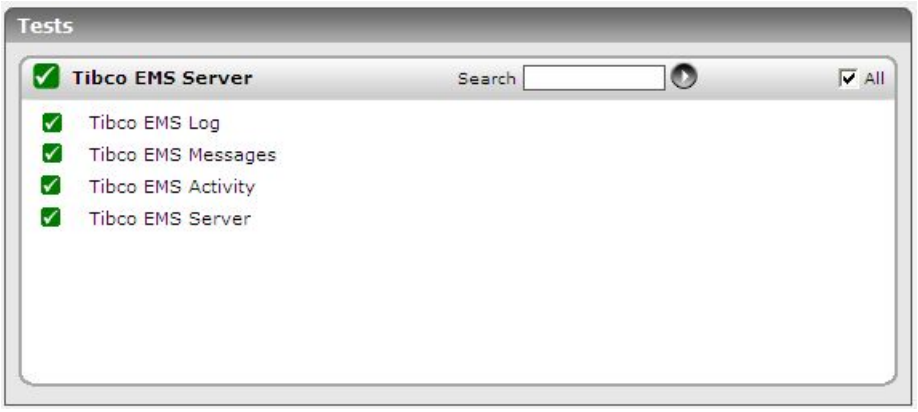

Figure 3.2: The tests mapped to the Tibco EMS Server Layer

3.1.1 Tibco EMS Log Test

The EMS server can be configured to write a variety of information to a log file. Server activities can be efficiently tracked using the information logged in the log files. Several parameters and commands control the location, logged information, and general configuration of the log files. One such parameter is the *logfile_max_size* configuration parameter that governs the maximum size upto which a log file can grow. To make sure that the log file does not grow boundlessly, this test periodically monitors the size of the log file and promptly alerts administrators if the log file is about to exceed the size limit set.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Log file size	Indicates the current size of the log file.	KB	

Measurement	Description	Measurement Unit	Interpretation
Max log file size	Indicates the maximum allocated size of the log file.	KB	This measure typically returns the value set for the logfile_max_size configuration parameter.
Log file usage	Indicates the percentage of allocated space that is currently used by the log file.	Percent	If the value of this measure is 100%, it indicates that the logfile_max_size threshold has been violated - i.e., the log file has reached its maximum size. In this case, the contents of the log file are copied to a file with the same name as the current log file, except that a sequence number is appended to the name of the backup file. The server queries the directory and determines the first available sequence number. For example, if the current log file is named tibems.log, the first copy is named tibems.log.1, the second is named tibems.log.2, and so on. To further regulare the space usage of the log files, you can also dynamically force the log file to be backed up and truncated using the rotatelog command in tibemsadmin.

3.1.2 Tibco EMS Messages Test

Messages are structured data that one application sends to another. Typically, the messages sent by an application are either delivered to a queue on the EMS server or published to a topic on the EMS server. In addition, you may have one/more messages that just wait on the EMS server pending delivery to a queue or a topic - this could be because of the messages being too large in size or the queue/topic being too crowded already to accomodate more messages.

This test takes stock of all the messages on an EMS server, regardless of where they are (whether queue/topic/pending delivery to queue or topic), and reports a wealth of performance statistics pertaining to the memory and storage space used by these messages on the server.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Pending messages	Indicates the number of pending messages on the server.	Number	If a consistent and significant increase is noticed in the value of this measure over a period of time, it could indicate either or both the following: a. A number of messages are still waiting to be delivered to queues or topics; b. Consumers are not / are unable to retrieve many messages from the queues/topics If this situation continues, it can unduly overload the server, forcing it to stall message delivery and return errors to the producers; this in turn may cause

Measurement	Description	Measurement Unit	Interpretation
			critical messages to not reach consumers on time. In this case therefore, the bottleneck areas must be identified and resolved quickly.
Pending message size	Indicates the total size of all the pending messages on the server.	KB	Ideally, the value of this measure should be low. While sporadic spikes in the value are normal and may indicate the entry of one/more bulky messages, a steady increase in the value of this measure over time could be a cause for concern, as it may induce the excessive consumption of the storage capacity of the server. This space erosion can adversely impact the overall server performance. To avoid such adversities, you might want to consider fine-tuning the EMS server to handle pending messages better. For instance, you can use the maxBytes configuration parameter of a queue/topic. maxbytes defines the maximum size (in bytes) of all messages that can be waiting in a queue or waiting to be delivered to the durable subscribers of a topic. By reducing the maxBytes of a queue/topic, you can reduce the size of all messages held by that destination, thus significantly reducing the resource foot-print of the messages within. If this limit is violated, the server returns an error to the producer. Similarly, you can try enabling 'Flow control'. Flow control is a feature that controls the flow of messages to a destination. If this capability is enabled for the server, then, you can use the flowControl configuration parameter to

Measurement	Description	Measurement Unit	Interpretation
			configure each destination with a target maximum size for storing pending messages. If need be, you can reduce the flowControl value, so that the storage capacity utilized by messages waiting to be delivered to a destination is reduced. If this limit is violated, the server blocks producers from sending any more messages.
Free message memory	Indicates the amount of storage memory unused on the server.	KB	Adequate storage memory should always be available on the server for storing messages. Excessive memory usage over time can exhaust the storage capacity of the server - the lack of sufficient memory can slowdown and can even bring to a halt, all server operations. To avoid such adversities, you can regulate the space usage of the messages on the server. For instance, you can set the maxBytes threshold for a queue/topic. maxbytes defines the maximum size (in bytes) of all messages that can be waiting in a queue or waiting to be published to a topic. By reducing the maxBytes of a queue/topic, you can reduce the size of all messages held by that destination, thus significantly reducing the resource foot-print of the messages within. If this limit is violated, the server returns an error to the producer. Similarly, you can try enabling 'Flow control'. Flow control is a feature that controls the flow of messages to a destination. If this capability is enabled for the server, then, you can use the flowControl configuration parameter to

Measurement	Description	Measurement Unit	Interpretation
			configure each destination with a target maximum size for storing pending messages. If need be, you can reduce the flowControl value, so that the storage capacity utilized by messages waiting to be delivered to a destination is reduced. If this limit is violated, the server blocks producers from sending any more messages.
Used message memory	Indicates the amount of storage memory that is currently in use for storing messages on the server.	KB	
Total message memory	Indicates the total memory allocated for storing messages on the server.	KB	
Message memory used	Indicates the percentage of memory used for storing messages.	Percent	
Message memory pooled	Indicates the size of the pools of storage allocated for messages.	KB	
Synchronous storage	Indicates the size of the synchronous storage memory.	KB	A synchronous storage memory is the one which can store only one message at a time.
Asynchronous storage	Indicates the size of the asynchronous storage memory.	KB	A asynchronous storage memory is the one which allows multiple message storing at a time.
Fsync	A asynchronous storage memory is the one which allows multiple message storing at a time.	Number	Critical message transaction will be done successfully by enabling the "fail safe" configuration for a specific queue. By this the messages for a "fail safe" queue are synchronously written to the disk. By enabling this, the message transactions will not be lost even when the server goes down. In other words, this mode is decribed as "reliable mode". The value 0 for this measure indicates that the 'fail safe' mode is disabled and and the value 100 for this measure indicates the 'fail safe' is enabled.

Measurement	Description	Measurement Unit	Interpretation
Inbound messages	Indicates the number of incoming messages to the message memory.	Number	
Outbound messages	Indicates the number of outgoing messages from the message memory.	Number	

3.1.3 Tibco EMS Activity Test

This test provides a snapshot of the level of activity on the server by reporting the number of sessions on and applications connecting to the server. Besides revealing how busy the server is, this test helps isolate idle connections to the server so that such connections can be promptly removed.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Used connections	Indicates the number of virtual connections to the server that are actively used.	Number	A Connection object encapsulates a virtual connection with the server. A connection typically represents a communication link between the application and the messaging server. A connection is a fairly heavyweight object, and therefore most clients will use one connection for all messaging. You may create multiple connections, if needed by your application. The value of this measure therefore will indicate whether any applications require multiple connections to the EMS server.
Connections used percent	Indicates the percentage of total connections to the server that are in use.	Percent	Ideally, this value should be high. A low ratio of used connections indicates that too many connections to the server are currently unused - i.e., are idle. Unused open connections are eventually closed, but they do consume resources that could be used for other applications. Too many idle connections therefore can accelerate the resource drain on the EMS server, thereby compelling other applications to contend for limited resources. A low value for this measure is hence a cause for concern.
Sessions	Indicates the number of sessions on the server.	Number	A session represents a single-threaded context for sending and receiving messages. A session is single-threaded so that messages are serialized, meaning that messages are received one-by-one in the order sent. The benefit of a session is that it supports transactions. If the user

Measurement	Description	Measurement Unit	Interpretation
			selects transaction support, the session context holds a group of messages until the transaction is committed, then delivers the messages. Before committing the transaction, the user can cancel the messages using a rollback operation. A session allows users to create message producers to send messages, and message consumers to receive messages. This measure is generally an indicator of how busy the EMS server is.
Producers	Indicates the number of producers communicating with the server.	Number	Messages are structured data that one application sends to another. The creator of a message is known as a producer. Using the value reported by this measure, you can accurately figure out how many applications are sending messages to the queues on the server.
Consumers	Indicates the number of consumers communicating with the server.	Number	Messages are structured data that one application sends to another. The receiver of messages is known as a consumer. Using the value reported by this measure, you can accurately figure out how many applications are retrieving messages from the queues on the server.
Durables	Indicates the number of messages for the durable subscriptions.	Number	Each message consumer subscribes to a topic. When a message is published to that topic, all subscribed consumers receive the message. By default, subscribers only receive messages when they are active. If the messages are delivered when the subscriber is not available, the

Measurement	Description	Measurement Unit	Interpretation
			subscriber does not receive these messages. Java Message Service(JMS) specifies a way to remove the part of timing dependency by allowing subscribers to create durable subscriptions. Messages for durable subscriptions are stored on the server until the message expires or the storage limit is reached. Subscribers can receive messages from a durable subscription even if the subscriber was not available when the message was originally delivered.

3.1.4 Tibco EMS Server Test

This test reports the current state of the server, and also indicates the load on the server in terms of the number and type of topics and queues handled by it.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter.

Parameter	Description
	The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is active?	Indicates the current state of the EMS Server.	Number	The value 0 for this measure indicates that the server being monitored is inactive, and the value 2 indicates that it is active. The value 1 on the other hand indicates that the server is currently in the fault tolerant standby mode. The fault-tolerant state is operated by configuring a pair of servers - namely, primary and backup. The primary and backup servers act as a pair, with the primary server accepting the client connections and performing the work of handling messages, and the secondary server acting as a backup in case of failure. If the primary server fails, the backup server resumes operation in its place.
Total topics	Indicates the total number of topics available on the server.	Number	Topics is a distribution mechanism for publishing messages that are delivered to multiple subscribers.The publisher generally addresses messages to a topic. Many publishers can publish to the same topic, and a message from a single publisher can be received by many subscribers. Subscribers subcribe to topics and all messages published to the topic are received by all subscribers to the topic. The value reported by this measure includes dynamic, static and temporary

Measurement	Description	Measurement Unit	Interpretation
			topics, and thus indicates the total topic load on the server.
Dynamic topics	Indicates the total number of dynamic topics on the server.	Number	Destinations for message can be either queues or topics, and these destinations can either be created statically or dynamically. Dynamic topics do not appear in the configuration files and exist as long as there are messages or consumers on the destination. Since these topics are created dynamically by applications, this measure serves as a good indicator of the topic load imposed by applications on the EMS server.
Temporary topics	Indicates the number of temporary topics on the server.	Number	Temporary destinations (temporary queues or temporary topics) are proposed as a lightweight alternative in a scalable system architecture that could be used as unique destinations for replies. Such destinations have a scope limited to the connection that created it, and are removed on the server side as soon as the connection is closed. Owing to their short life span, these topics may not have a lasting effect on the load/overall performance of the EMS server. For the same reason and those listed below, the temporary topics are not ideal destinations for messages: a. A temporary destination can only be consumed by the connection that created it. b. When you close the connection that has a temporary destination,

Measurement	Description	Measurement Unit	Interpretation
			the destination is closed and its contents are lost. c. You cannot have durable subscriptions to a TemporaryTopic. d. Each temporary destination is unique and cannot be copied. e. Temporary destinations cannot be routed using an enterprise messaging service. Typically, the value of this measure will be low. However, if the value appears to significantly increase with time, it could indicate a bottleneck in the closure of connections; further investigation may be required in this regard.
Static topics	Indicates the number of static topics on the server..	Number	The destination for a message can be either created statically in the server configuration files or dynamically by a client application. Static topics are those topics for which configuration information resides within configuration files for the EMS server. These topics are typically administered by the server itself. The value of this measure therefore indicates the server workload in terms of the number of static topics it is currently administering.
Total queues	Indicates the total number of queues..	Number	The point-to-point style of messaging uses a queue to store messages until they are received. The message producer sends the message to the queue; the message consumer retrieves messages from the queue and sends acknowledgement that the

Measurement	Description	Measurement Unit	Interpretation
			message was received. Based on how it is created, a queue can be classified as dynamic, static, or temporary. Based on what its created for, a queue can be configured to be 'exclusive' or 'non-exclusive'. If a queue is exclusive, then all messages in this queue can only be retrieved by the first consumer specified for the queue. Exclusive queues are useful when you want only one application to receive messages for a specific queue. If the queue is not exclusive, any number of receivers can retrieve messages from the queue. Non-exclusive queues are useful for balancing the load of incoming messages across multiple receivers. The value of this measure includes queues of all the above-mentioned types, and is useful for determining the load on the server in terms of the number of queues it handles.
Dynamic queues	Indicates the total number dynamic queues that are created.	Number	Dynamic queues do not appear in the configuration files and exist as long as there are messages or consumers on the destination. Since these queues are created dynamically by applications, this measure serves as a good indicator of the queue load imposed by applications on the EMS server.
Temporary queues	Indicates the number of temporary queues on the server.	Number	Temporary destinations (temporary queues or temporary topics) are proposed as a lightweight alternative in a scalable system architecture that could be used as unique destinations

Measurement	Description	Measurement Unit	Interpretation
			for replies. Such destinations have a scope limited to the connection that created it, and are removed on the server side as soon as the connection is closed. Owing to their short life span, temporary queues may not have a lasting effect on the load/overall performance of the EMS server. For the same reason and those listed below, the temporary queues are not ideal destinations for messages: a. A temporary destination can only be consumed by the connection that created it. b. When you close the connection that has a temporary destination, the destination is closed and its contents are lost. You cannot have durable subscriptions to a TemporaryTopic. c. Each temporary destination is unique and cannot be copied. d. Temporary destinations cannot be routed using an enterprise messaging service. Typically, the value of this measure will be low. However, if the value appears to significantly increase with time, it could indicate a bottleneck in the closure of connections; further investigation may be required in this regard.
Static queues	Indicates the number of	Number	The destination for a message can be

Measurement	Description	Measurement Unit	Interpretation
	static queues on the server.		either created statically in the server configuration files or dynamically by a client application. Static queues are those queues for which configuration information resides within configuration files for the EMS server. These queues are typically administered by the server itself. The value of this measure therefore indicates the server workload in terms of the number of static queues it is currently managing.

3.2 The Tibco EMS Objects Layer

This layer zooms into the performance of each durable subscriber, queue, and topic on the server, and reports which subscriber/queue/topic has the maximum number of pending messages on the server, or is utilizing too much message memory.

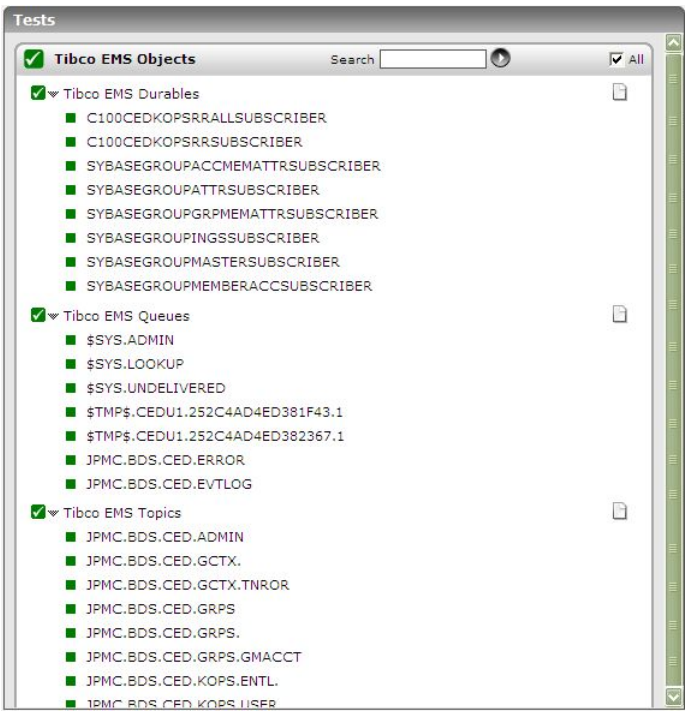

Figure 3.3: The tets mapped to the Tibco EMS Objects

3.2.1 Tibco EMS Durables Test

Generally, each message consumer subscribes to a topic. When a message is published to that topic, all subscribed consumers receive the message. Because of this there might be time dependency. By default, subscibers only receive messages when they are active. If the messages are delivered when the subscriber is not available, the subscriber does not receive these messages.

Java Message Service(JMS) specifies a way to remove the part of timing dependency by allowing subscribers to create durable subscriptions. Messages for durable subscriptions are stored on the server until the message expires or the storage limit is reached. Subscribers can receive messages from a durable subscription even if the subscriber was not available when the message was originally delivered.

When an application restarts and recreates a durable subscriber with the same ID, all messages stored on the server for that topic are published to the durable subscriber.

Until such time, the pending messages will remain on the EMS server, consuming machine resources. If the messages are allowed to grow limitlessly in size and number over time, they can drain the server of its resources, thereby choking critical server operations. Using this test, you can continuously monitor the status, activities, and pending messages of durable subscribers, and be proactively alerted if the pending messages appear to be growing significantly in size and/or number. Besides enabling you to control message growth, the metrics reported by this test also enable you to avert impending resource shortages on the server.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each durable subscriber on the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled

Parameter	Description
	with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is status offline?	Indicates the current state of this durable subscriber.	Boolean	If the value of this measure is 0, it indicates that the subscriber is offline, and if the value is 1, it indicates that the subscriber is online. Use the detailed diagnosis of this measure to know the topics to which the durable subscriber is subscribing.
Pending messages	Indicates the number of pending messages for this durable subscriber.	Number	If the value of this measure increases significantly over time, it could mean that more messages are still waiting to be delivered to subscribers. This could primarily be because, while the publishers are sending messages to

Measurement	Description	Measurement Unit	Interpretation
			the topic quickly, the subscriber is not receiving the messages quiet as quickly. One of the reasons for this slowdown could be a network connectivity issue between the subscriber and the EMS server - a poor network link could be delaying message delivery at one end, while on the other end, the topic may be experiencing a major influx of messages from the publisher. The pending message count can also increase, if the subscriber has been offline for a long time. During this period of inactivity, existing messages will remain undelivered on the server, and more messages may continue to be published to the topic by the publisher, thereby increasing the pending message count. However, regardless of the reason, the number of pending messages should be kept under control at all times, as they consume considerable server resources.
Message size	Indicates the total size of pending messages for this durable subscriber.	KB	If the size grows significantly over time, it could indicate that messages are getting published on the topic but are not being delivered to subscribers quiet as quickly. As already mentioned, the slowdown in delivery can be attributed to the durable subscriber being offline for a long time. Practical issues - such as bad network links - faced by subscribers while receiving messages from the topic can also delay/halt delivery. Regardless of the root cause, the message size

Measurement	Description	Measurement Unit	Interpretation
			should not be allowed to grow uncontrollably, as it may cause a serious resource contention on the server. To limit the growth of pending messages on a topic, use the maxBytes configuration for that topic. For topics, maxBytes limits the total size (in bytes) of all messages waiting for delivery to each durable subscriber on that topic. If this limit is violated, then messages will be go undelivered, thus causing the receivers to lose critical data.

Use the detailed diagnosis of the *Is status offline?* measure to know the topics to which the durable subscriber is subscribing.

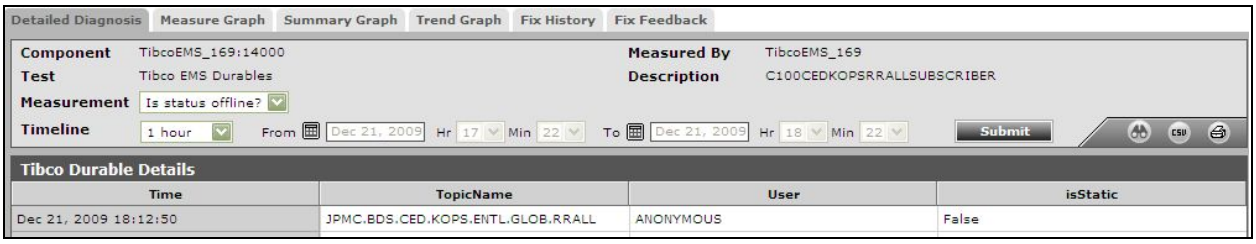

Figure 3.4: The detailed diagnosis of the Is status offline? measure

3.2.2 Tibco EMS Queues Test

The point-to-point style of messaging uses a queue to store messages until they are received. The message producer sends the message to the queue; the message consumer retrieves messages from the queue and sends acknowledgment that the message was received. More than one producer can send messages to the same queue, and more than one consumer can retrieve messages from the same queue. The queue can be configured to be exclusive, if desired. If the queue is exclusive, then all queue messages can only be retrieved by the first consumer specified for the queue. Exclusive queues are useful when you want only one application to receive messages for a specific queue. If the queue is not exclusive, any number of receivers can retrieve messages from the queue. Non-exclusive queues are useful for balancing the load of incoming messages across

multiple receivers. Regardless of whether the queue is exclusive or not, only one consumer can ever retrieve each message that is placed on the queue.

For every queue configured on the EMS server, this test enables you to track the length of the queue and its size (in bytes), so that you can be promptly alerted to any abnormal increase in queue length or to any load imbalances experienced by the queue. In addition, the test reports the queue type and the number of receivers to the messages in the queue.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each queue on the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.
Ignore Queues	Specify a comma-separated list of queues to be excluded from monitoring in the Ignore Queues text box. A * can be used in the queue name to indicate leading or trailing spaces. For instance, to ignore queues with names that embed the string *sys*, your Ignore Queues specification can be: **sys**

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Receivers	Indicates the number of	Number	Receivers can retrieve messages from

Measurement	Description	Measurement Unit	Interpretation
	currently active receivers in this queue.		the queue. By default receivers can receive multiple messages at the same instant. You can also configure the receivers to receive only one message at a time. The value of this measure will help you ascertain the workload of the queue in terms of the number of receivers it supports.
Pending messages	Indicates the number of pending messages in this queue.	Number	If the value of this measure increases significantly over time, it could mean that the messages are getting accumulated in the queue. This could primarily be because the producers are sending messages more rapidly than the speed of consumption by the consumers. One of the reasons for this slowdown could be that the consumer has been offline for a long time - this also allows more messages to be enqueued. Connectivity issues faced by the consumer while retrieving messages from the queue could also delay consumption. Regardless of the reason, the queue length should be kept under a check at all times, so as to conserve server resources and the loss of critical messages.
Message size	Indicates the total size of pending messages in this queue.	KB	If the size grows significantly over time, it could indicate that messages are getting added to the queue but are not being consumed quiet as quickly. As already mentioned, the slowdown in consumption can be attributed to the consumer being offline for a long time. Practical issues - such as bad network links - faced by consumers while retrieving messages from the queue can also delay/halt consumption.

Measurement	Description	Measurement Unit	Interpretation
			Regardless of the root cause, the message size should not be allowed to grow uncontrollably, as it may completley erode the storage resources of the server. To limit the growth of a queue, use the the maxBytes configuration for that queue. maxbytes defines the maximum size (in bytes) of all messages that can be waiting in a queue. If this limit is violated, an error is returned to the producers.
IsStatic	Indicates whether the queue is static or not.	Boolean	The value 0 for this measure indicates that the queue is static, and the value 1 indicates that it is non-static - in this case, the queue can be of type dynamic or temporary. A static queue is typically created statically in the server configuration files, and the dynamic queue is created on-the-fly by applications. The dynamic queues do not appear in the configuration files, and exist as long as there are messages or consumers on the destination. You can identify dynamic queues in the administration tool using the asterisk (*) that will pre-fix their names. Temporary destinations (temporary queues or temporary topics) are proposed as a lightweight alternative in a scalable system architecture that could be used as unique destinations for replies. Such destinations have a scope limited to the connection that created it, and are removed on the server side as soon as the connection is closed. Owing to their short life span, these queues may not have a lasting effect on the

Measurement	Description	Measurement Unit	Interpretation
			load/overall performance of the EMS server.

3.2.3 Tibco EMS Topics Test

Topics are a distribution mechanism for publishing messages that are delivered to multiple subscribers. A topic's properties are set when the destination is created. The publisher generally addresses messages to a topic. Many publishers can publish to the same topic, and a message from a single publisher can be received by many subscribers. Subscribers subcribe to topics and all messages published to the topic are received by all subscribers to the topic.

This test allows you to keep tabs on the number and size of pending messages to each topic, so that you can accurately identify topics to which too many messages are pending and investigate the reasons for the same. In addition, the test reports the topic type and the number of subscribers to a topic, so that you can easily determine how popular the topic is.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each topic on the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.

Parameter	Description
Ignore Topics	Specify a comma-separated list of topics to be excluded from monitoring in the Ignore Topics text box. A * can be used in the topic name to indicate leading or trailing spaces. For instance, to Ignore Topics with names that embed the string *sys*, your Ignore Topics specification can be: **sys**

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Subscribers	Indicates the number of current subscribers to this topic.	Number	The value reported by this measure includes the number of durable subscribers as well. This measure therefore is a good indicator of how popular a particular topic is.
Durables	Indicates the number of durable subscribers to this topic.	Number	
Pending messages	Indicates the number of messages on this topic waiting to be delivered to subscribers.	Number	If the value of this measure increases significantly over time, it could mean that more messages are still waiting to be delivered to subscribers. This could primarily be because, while the publishers are sending messages to the topic quickly, the subscriber is not receiving the messages quiet as quickly. One of the reasons for this slowdown could be a network connectivity issue between the subscriber and the EMS server - a poor network link could be delaying message delivery at one end, while on the other end, the topic may be experiencing a major influx of messages from the publisher. The pending message count can also increase, if a topic has durable subscribers, and one/more of these subscriber have been offline for a long time. During this period of inactivity,

Measurement	Description	Measurement Unit	Interpretation
			existing messages will remain undelivered on the server, and more messages may continue to be published to the topic by the publisher, thereby increasing the pending message count. However, regardless of the reason, the number of pending messages should be kept under control at all times, as they consume considerable server resources.
Messages size	Indicates the total size of pending messages.	KB	If the size grows significantly over time, it could indicate that messages are getting published on the topic but are not being delivered to subscribers quiet as quickly. As already mentioned, the slowdown in delivery can be attributed to the durable subscriber being offline for a long time. Practical issues - such as bad network links - faced by subscribers while receiving messages from the topic can also delay/halt delivery. Regardless of the root cause, the message size should not be allowed to grow uncontrollably, as it may cause a serious resource contention on the server. To limit the growth of pending messages on a topic, use the maxBytes configuration for that topic. For topics, maxBytes limits the total size (in bytes) of all messages waiting for delivery to each durable subscriber on that topic. If this limit is violated, then messages will be go undelivered, thus causing the receivers to lose critical data.
IsStatic	Indicates whether the topic is static or not.	Boolean	The value 0 for this measure indicates that the topic is in static state and the

Measurement	Description	Measurement Unit	Interpretation
			value 1 for this measure indicates that it is in a non-static (i.e., dynamic or temporary) state. A static topic is typically created statically in the server configuration files, and the dynamic topic is created on-the-fly by applications. The dynamic topics do not appear in the configuration files, and exist as long as there are messages or consumers on the destination. You can identify dynamic topics in the administration tool using the asterisk (*) that will pre-fix their names. Servers connected by routes exchange messages sent to temporary topics. As a result, temporary topics are ideal destinations for reply messages in request/reply interactions.

3.3 The Tibco EMS Service Layer

The test mapped to this layer focus on the user activity on the server, and reveals which user is having the maximum number of sessions and connections open on the server.

Figure 3.5: The tests mapped to the Tibco EMS Service layer

3.3.1 Tibco EMS Connections Test

This test monitors the user activity on the EMS server, and reports the number of connections and sessions initiated by each user on the server. The users with the maximum number of open sessions on the server can thus be quickly identified and their activities closely tracked.

Target of the test : A Tibco EMS Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each user connected to the Tibco EMS server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
CommandPath	Prior to monitoring the Tibco EMS server, you will have to build a .bat or .sh file (depending upon the operating system on which Tibco EMS is functioning) bundled with the commands that the eG agent needs to execute on the Tibco EMS server for collecting the required metrics. The commands to be invoked by the .bat or .sh file are provided in **How does eG Enterprise Monitor Tibco EMS?** chapter. The .bat/.sh file so created can be saved to any location on the Tibco EMS host. Then, while configuring this test, make sure you provide the full path to this .bat or .sh file in the CommandPath text box so that, the agent can execute the file, invoke the commands bundled into it, and extract the desired metrics from the server.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis

Parameter	Description
	measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total connections	Indicates the total number of connections on the server for this user.	Number	The detailed diagnosis of this measure, if enabled, will provide the details of each connection that this user has established with the EMS server. The host from which the user is connecting and the total uptime of the connection can be determined using this information.
Total sessions	Indicates the total number of sessions on the server for this user.	Number	

The detailed diagnosis of the *Total connections* measure, if enabled, will provide the details of each connection that this user has established with the EMS server. The host from which the user is connecting and the total uptime of the connection can be determined using this information.

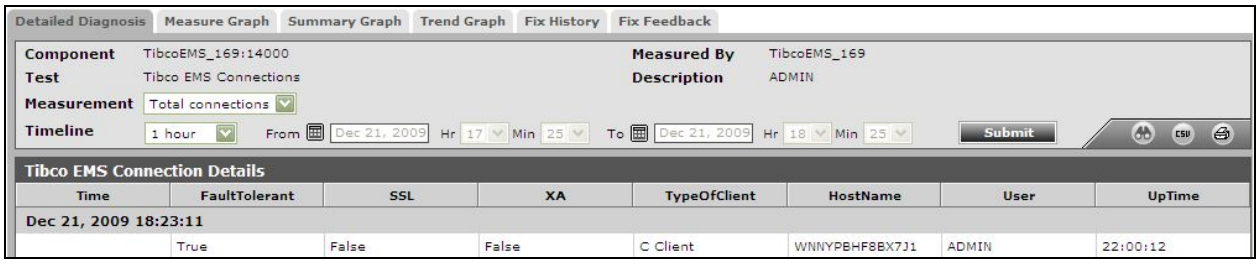

Figure 3.6: The detailed diagnosis of the Total connections measure

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.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.