Monitoring Microsoft SQL Cluster

eG Innovations Product Documentation

www.eginnovations.com

eG
Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

When two or more MS SQL servers exist in an environment, they can be grouped together to form a SQL cluster. Requests to a cluster are routed through a virtual cluster server that is assigned a cluster IP address and TCP port. Requests to this server can be handled by any of the individual nodes in the cluster at any given point in time, depending on which node is active at that time.

Since clusters are deployed in environments where 24*7 availability and responsiveness are critical, it is imperative that the performance of the clusters is monitored all the time. The eG Enterprise aids administrators in this regard!

Chapter 2: Administering the eG Manager to monitor the Microsoft SQL Cluster Server

1. Log into the eG administrative interface.

2. eG Enterprise cannot automatically discover Microsoft SQL Cluster Server. You need to manually add the server using the **COMPONENTS** page (see Figure 2.1) that appears when the Infrastructure -> Components -> Add/Modify menu sequence is followed. Remember that components manually added are managed automatically.

Figure 2.1: Adding a Microsoft SQL Cluster Server

3. Specify the **Host IP/Name** and the **Nick name** for the Microsoft SQL Cluster Server in Figure 2.1. The default port number appears in the **Port number** text box. The Microsoft SQL Cluster Server can only be monitored in an agentless manner. Therefore, the **Agentless** flag would be set to **Yes**, by default. Then, select **Windows** or **Linux** from the **OS** list box. If you select **Windows** as the **OS**, then select **Perfmon** as the **Mode**. If you select **Linux** as the **OS**, then select **SSH** as the **Mode**. Finally, click the **Add** button to register the changes.

The eG Enterprise also enables you to configure and monitor the Microsoft SQL Cluster Server that is installed as a 'named instance' in your environment. In such a case, you have to enter the host name or IP address of the SQL instance, which you specified while installing the SQL

instance, against the **Host IP/Name** text box in Figure 2.1. The IP address and host name of the SQL instance can be viewed in the Failover Cluster Manager console as shown in Figure 2.2.
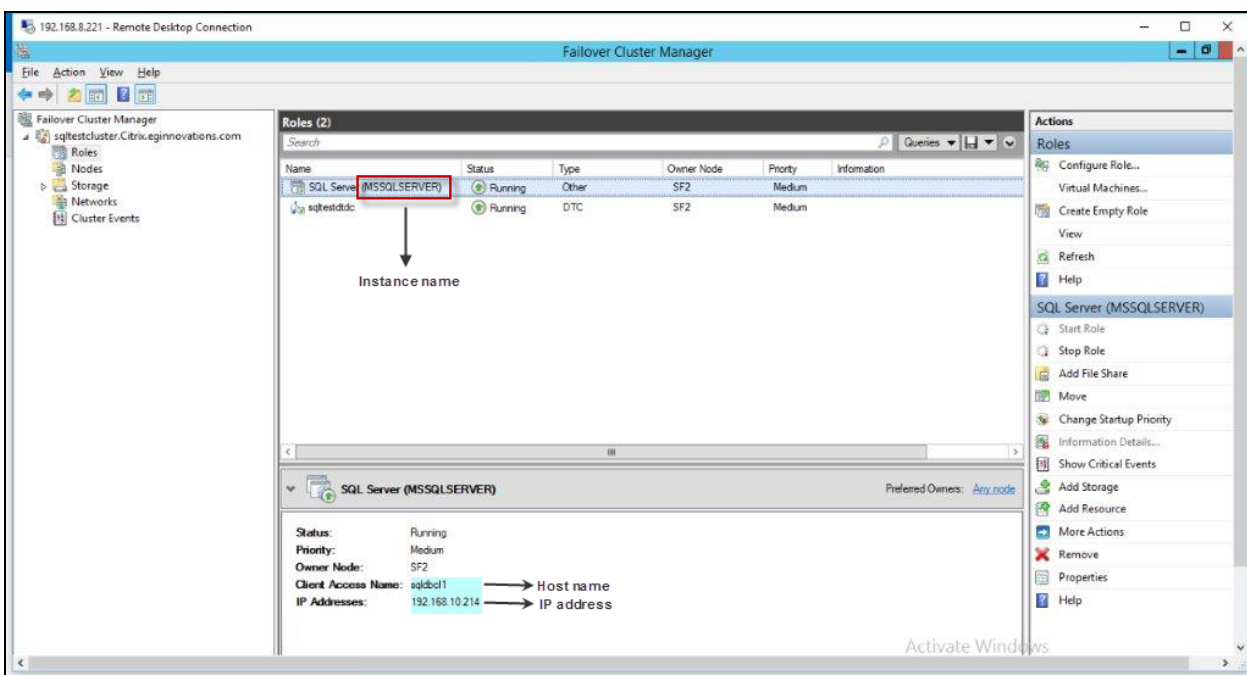
Figure 2.2: Failover Cluster Manager console

In addition, name of the SQL instance (see Figure 2.2) has to be specified against the **INSTANCE NAME** parameter while configuring the tests for monitoring the Microsoft SQL Cluster Server.

4. When you attempt to sign out, a list of unconfigured tests will appear as shown in Figure 2.3.

Figure 2.3: List of Unconfigured tests to be configured for the Microsoft SQL Cluster Server

5. Click on the **SQL Cluster Connection** test to configure it. To know how to configure this test, refer to Section **3.5.1**.

6. Once all the tests are configured, signout of the eG administrative interface.

Chapter 3: Monitoring Microsoft SQL Cluster Server

n

Figure 3.1 depicts the *Microsoft SQL Cluster Server* monitoring model.
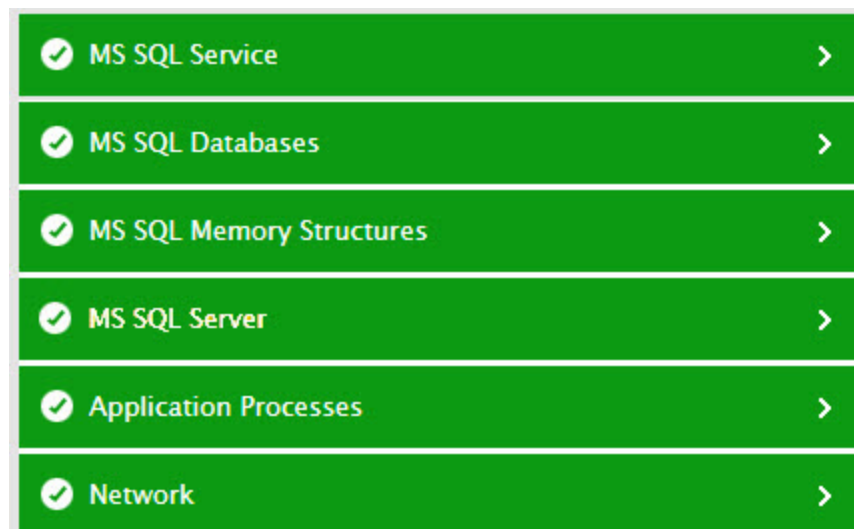
Figure 3.1: The layer model of a SQL cluster service

This section will be discussing the top 3 layers of Figure 3.1 only, as the **Network** layer has already been extensively discussed in the *Monitoring Unix and Windows Servers* document.

3.1 The Application Processes Layer

The tests mapped to this layer report on the availability of the cluster server port and the cluster process.
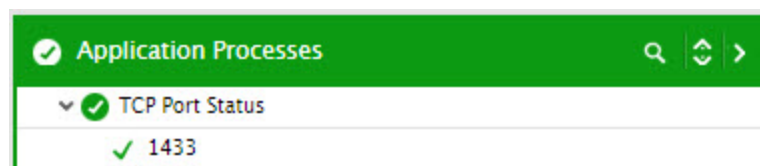
Figure 3.2: The tests mapped to the Application Processes layer

3.1.1 SQL Cluster Process Test

This test reports the current state and resource usage of the cluster process, and promptly alerts administrators if the cluster goes down or is up and consuming more resources than it should.

Note:

This test will report metrics only if the eG agent is installed on a remote Windows host. The eG agent should be installed using the credentials of a domain user possessing valid permission to log into the target Microsoft SQL server.

Target of the test : A Microsoft SQL Cluster

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the cluster being monitored

Configurable parameters for the test

1. **TEST PERIOD** – How often should the test be executed

2. **HOST** – The IP address of the SQL cluster

3. **PORT** – The port on which the cluster is listening.

4. **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
Service availability	Indicates the availability of the cluster service.	Percent	The availability is 100% when the cluster process is running. This measure will report the value 0, if the cluster process is not running.

Measurement	Description	Measurement Unit	Interpretation
Processes running	Indicates the number of instances of the cluster process that are currently running.	Number	The value 0 for this measure is indicative of the non-availability of the cluster service. The detailed diagnosis of this measure reveals the process ID of each running instance of the cluster process and the resource and I/O usage of each instance, so that you can quickly identify which instance is resource-intensive.
CPU utilization	Indicates the percentage of CPU resources used by the cluster process.	Percent	A value close to 100% indicates that the cluster process is hogging the CPU resources of the 'active' node in the cluster.
Handle count	Indicates the number of handles opened by the process.	Number	An increasing trend in this measure is indicative of a memory leak in the process.
Number of threads	Indicates the number of threads that are used by the process.	Number	
Virtual memory used	Indicates the amount of virtual memory that is being used by the process.	MB	
Private memory used	Indicates the amount of memory that this process has been allotted, that cannot be shared with other processes.	MB	
Pool paged	Indicates the amount of memory currently in the pool paged area of	MB	

Measurement	Description	Measurement Unit	Interpretation
	system memory.		
Pool non paged	Indicates the amount of memory currently in the pool non-paged area of system memory.	MB	
I/O data rate	Indicates the rate at which processes are reading and writing bytes in I/O operations.	Kbytes/Sec	This value counts all I/O activity generated by each process and includes file, network and device I/Os.
I/O data operations	Indicates the rate at which the process is issuing read and write data to file, network and device I/O operations.	Operations/Sec	
I/O read data rate	Indicates the rate at which the process is reading data from file, network and device I/O operations.	Kbytes/Sec	
I/O write data rate	Indicates the rate at which the process is writing data to file, network and device I/O operations.	Kbytes/Sec	
Page fault rate	Indicates the total rate at which page faults are occurring for the threads of all matching processes.	Faults/Sec	A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.

Measurement	Description	Measurement Unit	Interpretation
Memory working set	Indicates the current size of the working set of a process.	MB	The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed they will then be soft-faulted back into the Working Set before leaving main memory. The detailed diagnosis for this test provides details of the individual process instances and their individual working sets. Comparing the working set across process instances indicates which instances taking up excessive memory.

3.2 The Microsoft SQL Server Layer

The tests associated with this layer monitor the health of the SQL server engine and the number and type of system processes executing on the Microsoft SQL server. In addition, the test also reports the count of blocker processes executing on the Microsoft SQL server. These tests are been discussed in the *Monitoring Microsoft SQL* document.
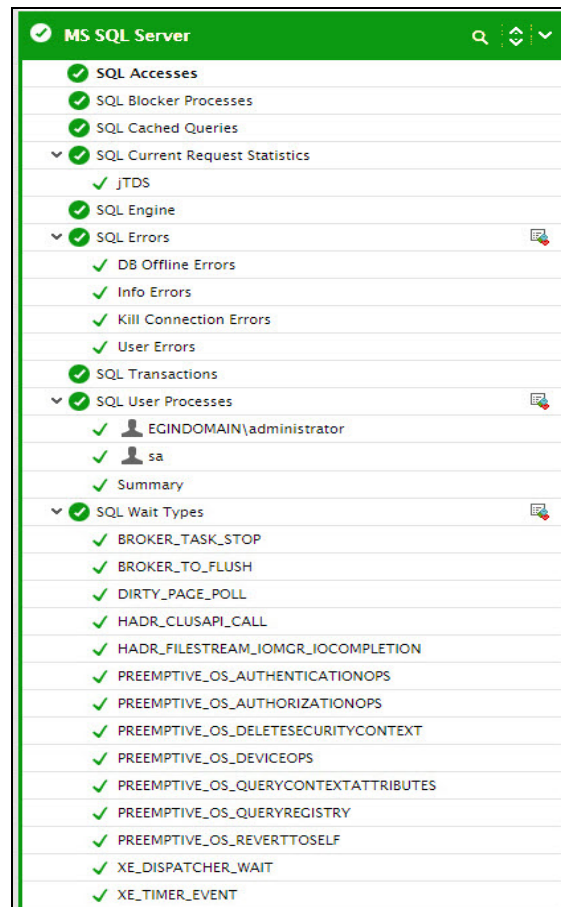
Figure 3.3: The Microsoft SQL Server layer

3.3 The Microsoft SQL Memory Structures Layer

This layer tracks the health of the memory and buffer structures of an Microsoft SQL server. The details of the tests are available *Monitoring Microsoft SQL* document.
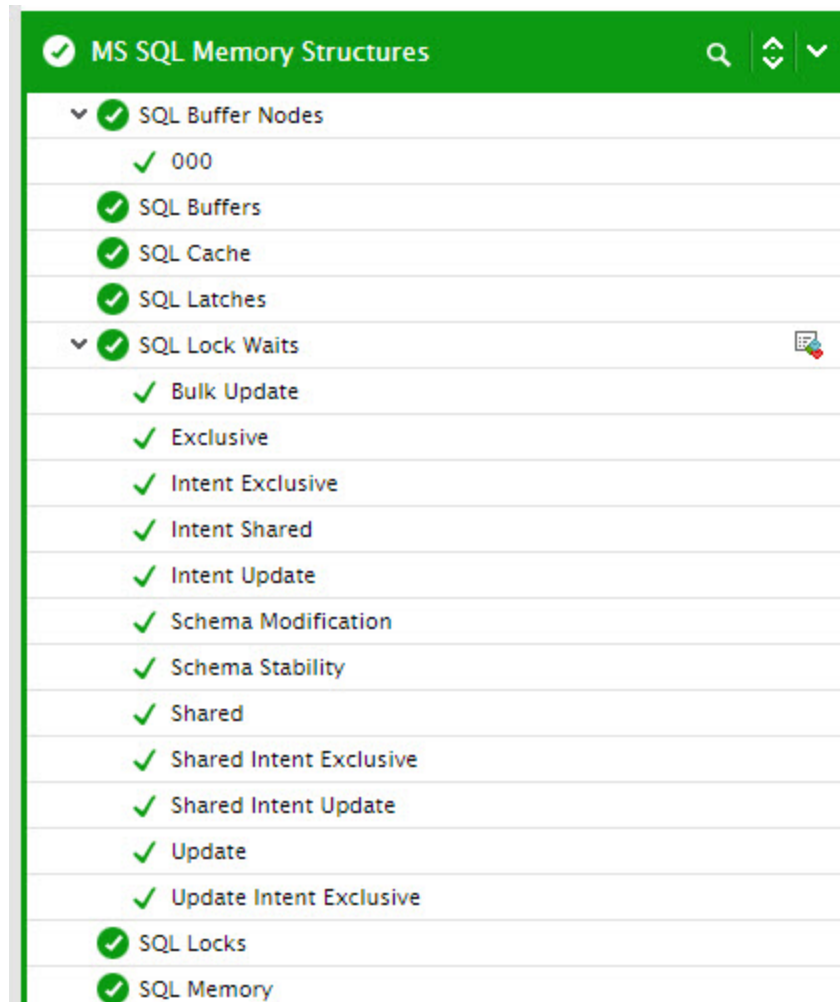
Figure 3.4: The tests mapped to the Microsoft SQL Memory Structures layer

3.4 The Microsoft SQL Databases Layer

The space usage on the Microsoft SQL server databases and the transaction log space usage can be tracked using the tests associated with this layer. These tests have been discussed in the *Monitoring Microsoft SQL Server* document.

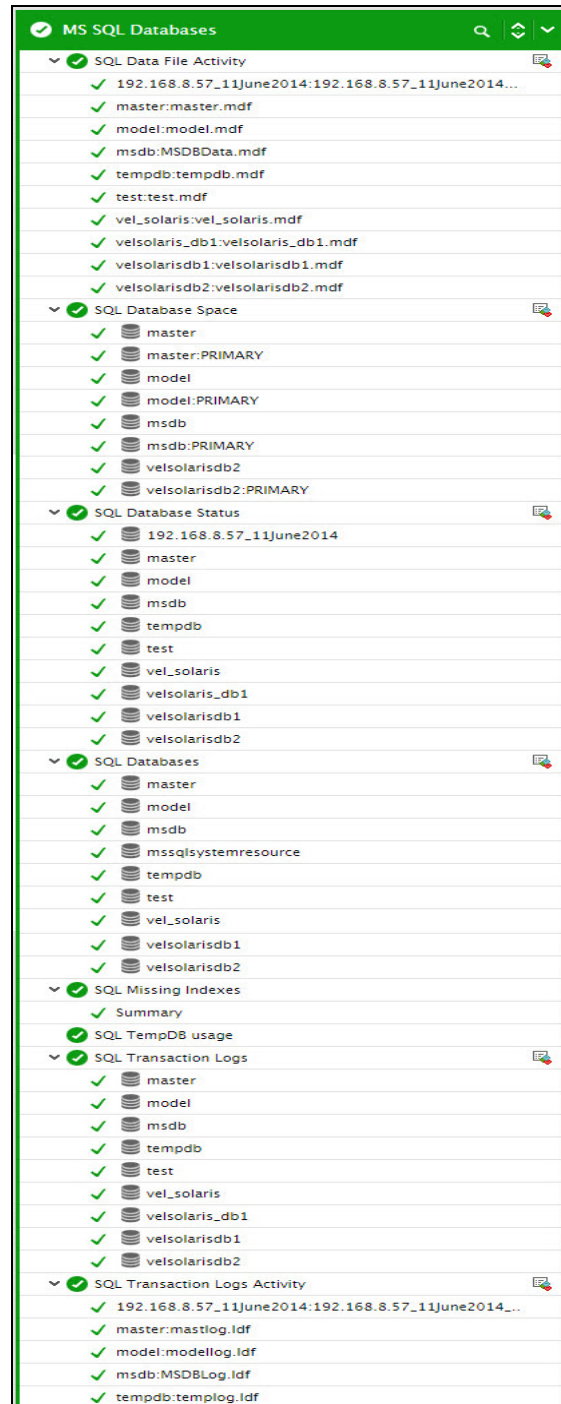Figure 3.5: The tests mapped to the Microsoft SQL Databases layer

3.5 The Microsoft SQL Service Layer

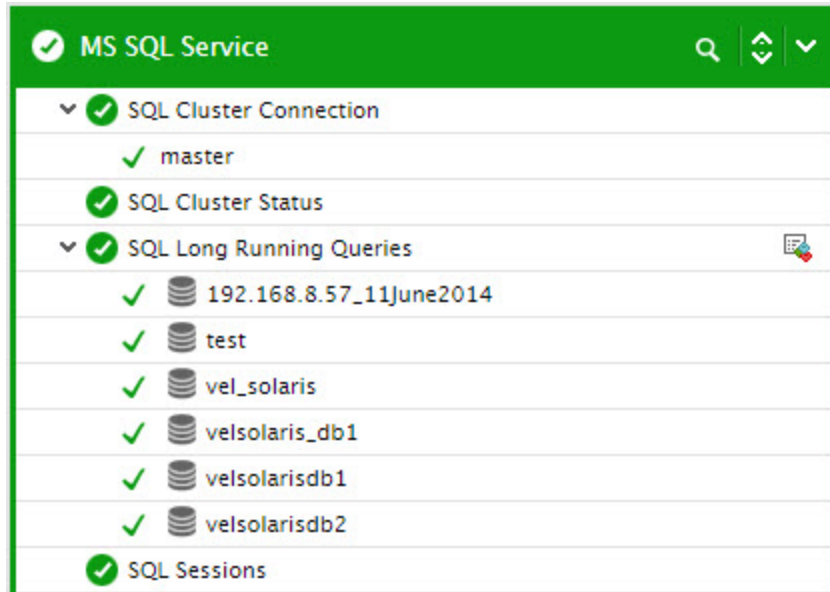Figure 3.6 depicts the tests associated with this layer.

Figure 3.6: The tests associated with the Microsoft SQL Service layer

The **SQL Sessions** and **SQL Long Running Queries** tests have been discussed in the *Monitoring Microsoft SQL* document. Therefore, the sub-sections that follow will elaborate on the cluster-specific tests mapped to this layer only.

3.5.1 SQL Cluster Connection Test

This test emulates a user executing a query on the cluster, and in the process, captures the availability of the cluster service and the responsiveness of the cluster.

Target of the test : A Microsoft SQL Cluster

Agent deploying the test : An external agent; if you are running this test using the external agent on the eG manager box, then make sure that this external agent is able to communicate with the port on which the virtual cluster server is listening. Alternatively, you can deploy the external agent that will be running this test on a host that can access the port on which the virtual cluster server is listening.

Outputs of the test : One set of results for the cluster being monitored

Configurable parameters for the test

1. **TEST PERIOD** – How often should the test be executed
2. **HOST** – The IP address of the SQL cluster
3. **PORT** – The port on which the cluster is listening

4. **USER** – A database user name.

5. **PASSWORD** - The password associated with the above user name (can be 'NULL'). Here, 'NULL' means that the user does not have any password.

6. **CONFIRM PASSWORD** – Confirm the **PASSWORD** (if any) by retyping it here.

7. **DATABASE** - The name of the database to connect to. The default is "master".

8. **QUERY** – The select query to execute. The default is "select * from master.dbo.spt_monitor".

9. **CASE** – Takes the value "upper" or "lower" depending upon the case-sensitivity of the SQL server installation.

10. **INSTANCE** – The name of a specific Microsoft SQL instance to be monitored. The default value of this parameter is "default". To monitor an Microsoft SQL instance named "CFS", enter this as the value of the "**INSTANCE**" parameter.

11. **DOMAIN** - By default, *none* is displayed in the **DOMAIN** text box. If the 'SQL server and Windows' authentication has been enabled for the server being monitored, then the **DOMAIN** can continue to be *none*. On the other hand, if 'Windows only' authentication has been enabled, then, in the **DOMAIN** text box, specify the Windows domain in which the managed Microsoft SQL server exists. Also, in such a case, the **USER** name and **PASSWORD** that you provide should be that of a user authorized to access the monitored SQL server.

12. **SSL** - If the Microsoft SQL server being monitored is an SSL-enabled server, then set the **SSL** flag to **Yes.** If not, then set the **SSL** flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
SQL availability:	Indicates the availability of the cluster service	Percent	The availability is 100% when the cluster is able to respond to a request. This can happen if any one server in the cluster is currently 'active' and is responding to client requests. This measure will report the value 0, if the cluster service is not up and running. Such an eventuality can be caused by the non-availability of active nodes in the cluster to handle

Measurement	Description	Measurement Unit	Interpretation
			the emulated query.
SQL response time:	Indicates the time taken by the cluster to respond to a user query	Seconds	A sudden increase in response time is indicative of a bottleneck in query processing on the 'active' server of the cluster.
Database connection availability:	Indicates whether the database connection is available or not.	Percent	If this measure reports the value 100 , it indicates that the database connection is available. The value 0 on the other hand indicates that the database connection is unavailable. A connection to the database may be unavailable if the database is down or if the database is listening on a port other than the one configured for it in the eG manager or owing to a poor network link. If the SQL availability measure reports the value 0, then, you can check the value of this measure to determine whether/not it is due to the unavailability of a connection to the server.
Query processor availability:	Indicates whether the database query is executed successfully or not.	Percent	If this measure reports the value 100, it indicates that the query executed successfully. The value 0 on the other hand indicates that the query failed. In the event that the SQL availability measure reports the value 0, check the value of this measure to figure out whether the failed query is the reason why that measure reported a server unavailability.
Database	Indicates the time taken	Secs	A high value could indicate a

Measurement	Description	Measurement Unit	Interpretation
connection time:	by the database connection.		connection bottleneck. Whenever the SQL response time of the measure soars, you may want to check the value of this measure to determine whether a connection latency is causing the poor responsiveness of the cluster.
Query execution time:	Indicates the time taken for query execution.	Secs	A high value could indicate that one/more queries to the cluster are taking too long to execute. Inefficient/badly designed queries often run for long periods. If the value of this measure is higher than that of the Connection time measure, you can be rest assured that long running queries are the ones causing the responsiveness of the cluster to suffer.
Records fetched:	Indicates the number of records fetched from the database.	Number	The value 0 indicates that no records are fetched from the database.

3.5.2 SQL Cluster Status Test

This test reports the current status (whether running or not) and uptime of the cluster service. This way, administrators can quickly find out if the cluster service was restarted recently. In addition, the test also indicates whether/not fail-over occurred recently. The IP address of the 'active' server in the cluster is also revealed as part of detailed diagnosis.

Note:

This test will report metrics only on Microsoft SQL Server 2008 (and above).

Target of the test : A Microsoft SQL Cluster

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for the cluster being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed

2. **HOST** - The host for which the test is to be configured.

3. **SSL** – If the Microsoft SQL server being monitored is an SSL-enabled server, then set the **SSL** flag to **Yes**. If not, then set the **SSL** flag to **No**.

4. **INSTANCE** - In this text box, enter the name of a specific Microsoft SQL instance that is to be monitored. The default value of this parameter is "default". To monitor an Microsoft SQL instance named "CFS", enter this as the value of the **INSTANCE** parameter.

5. **USER** – Provide the name of a SQL user with the **VIEW SERVER STATE** role.

6. **PASSWORD** - The password of the specified **USER**.

7. **CONFIRM PASSWORD** - Confirm the password by retyping it

8. **DOMAIN** - By default, *none* is displayed in the **DOMAIN** text box. If the 'SQL server and Windows' authentication has been enabled for the server being monitored, then the **DOMAIN** can continue to be *none*. On the other hand, if 'Windows only' authentication has been enabled, then, in the **DOMAIN** text box, specify the Windows domain in which the managed Microsoft SQL server exists. Also, in such a case, the **USER** name and **PASSWORD** that you provide should be that of a user authorized to access the monitored SQL server.

9. **ISNTLMV2** - In some Windows networks, NTLM (NT LAN Manager) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the **ISNTLMV2** flag is set to **No**, indicating that NTLMv2 is not enabled by default on the target Microsoft SQL host. Set this flag to **Yes** if NTLMv2 is enabled on the target host.

10. **ISPASSIVE** – If the value chosen is **YES**, then the Microsoft SQL server under consideration is a passive server in a SQL cluster. No alerts will be generated if the server is not running. Measures will be reported as "Not applicable' by the agent if the server is not up.

11. **REPORTMANAGERTIME** – By default, this flag is set to **Yes**, indicating that, by default, the detailed diagnosis of this test, if enabled, will report the shutdown and reboot times of the device in the manager's time zone. If this flag is set to **No**, then the shutdown and reboot times are shown in the time zone of the system where the agent is running(i.e., the system being managed for agent-based monitoring, and the system on which the remote agent is running - for agentless monitoring).

12. 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 cluster running ?:	Indicates whether/not the cluster is currently running.		The values that this measure can report and their corresponding numeric values are listed in the table below: | Measure Value	Numeric Value
---	---
Yes	1
No	0	 **Note:** By default, the test reports the **Measure Value**s in the table above to indicate whether/not the cluster is running currently. In the graph of this measure however, the same is represented using the numeric equivalents only. The detailed diagnosis of the Is cluster running? measure reveals the IP address of the currently 'active' node in the cluster and the

Measurement	Description	Measurement Unit	Interpretation
			date/time at which the active node was last started.
Has SQL cluster switched?:	Indicates whether/not fail-over occurred in the last measurement period.	Boolean	The values that this measure can report and their corresponding numeric values are listed in the table below: **Note:** By default, the test reports the **Measure Value**s in the table above to indicate whether/not fail- over occurred. In the graph of this measure however, the same is represented using the numeric equivalents only.
Uptime of the SQL cluster:	Indicates the total time the SQL cluster has been up since its last reboot.	Secs	Administrators may wish to be alerted if a cluster has been running without a reboot for a very long period. Setting a threshold for this metric allows administrators to determine such conditions.
SQL cluster uptime since last measure:	Indicates how long the SQL cluster has been up since the last measurement period.	Secs	If the cluster has not been rebooted during the last measurement period and the agent has been running continuously, this value will be equal to the measurement period. If the cluster was rebooted during the last measurement period, this value will be less than the measurement

Within the "Has SQL cluster switched?" interpretation cell:

Measure Value	Numeric Value
Yes	1
No	0

Measurement	Description	Measurement Unit	Interpretation
			period of the test. For example, if the measurement period is 300 secs, and if the cluster was rebooted 120 secs back, this metric will report a value of 120 seconds. The accuracy of this metric is dependent on the measurement period – the smaller the measurement period, greater the accuracy.
Has the SQL cluster been restarted?:	Indicates whether the SQL cluster server has been rebooted during the last measurement period or not.		The values that this measure can report and their corresponding numeric values are listed in the table below: **Note:** By default, the test reports the **Measure Value**s in the table above to indicate whether/not the cluster was restarted in the last measure period. In the graph of this measure however, the same is represented using the numeric equivalents only.

Measure Value	Numeric Value
Yes	1
No	0

The detailed diagnosis of the *Is cluster running?* measure reveals the IP address of the currently 'active' node in the cluster and the date/time at which the active node was last started.
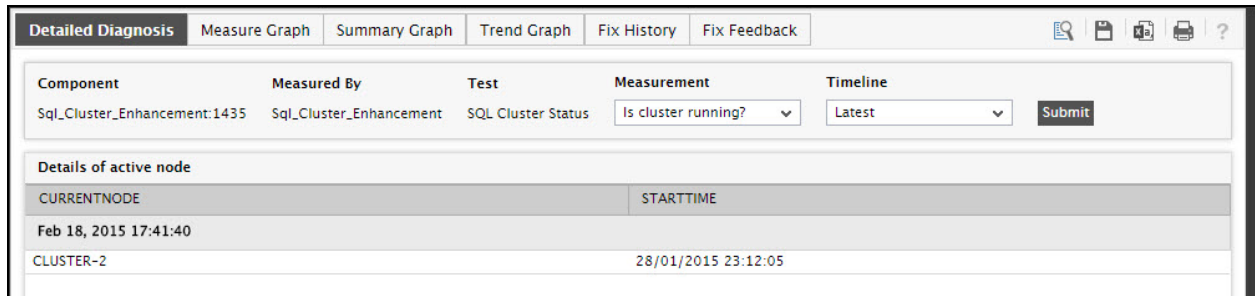
Figure 3.7: The detailed diagnosis of the Is cluster running? measure of the SQL Cluster Status test

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.