Monitoring Microsoft File Server

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

In the client/server model, a file server is a computer responsible for the central storage and management of data files so that other computers on the same network can access the files. A file server allows users to share information over a network without having to physically transfer files. Any computer can be configured to be a host and act as a file server. In its simplest form, a file server may be an ordinary PC that handles requests for files and sends them over the network. In a more sophisticated network, a file server might be a dedicated network-attached storage device that also serves as a remote hard disk drive for other computers, allowing anyone on the network to store files on it as if to their own hard drive.

The true indicator of the efficiency of a File server is the speed with which it serves concurrent file requests. If users are unable to access important files stored on the file server as and when they need due to a temperorary break in connection to the server or because of a long request queue, it might severely hamper the productivity of the users, and might unnecessarily delay critical operations. If such a problem situation is to be averted, the file server needs to be monitored, and administrators promptly warned about probable performance issues. eG Enterprise Suite helps administrators in this regard.

This document describes the monitoring model that eG Enterprise prescribes for the File server, and the performance metrics the model collects.

Chapter 2: How to Monitor Microsoft File Server Using eG Enterprise?

eG Enterprise adopts an agent-based approach to monitoring the Microsoft Exchange 2007 /2010 server with Mailbox Server role

To do the above, do the following:

1. Log into the eG administrative interface.
2. If an Microsoft File server is already discovered, then directly proceed towards managing it using the **COMPONENTS - MANAGE/UNMANAGE** page (Infrastructure -> Components -> Manage/Unmanage). However, if it is yet to be discovered, then run discovery (Infrastructure -> Components -> Discover) to get it discovered or add the MS File server manually using the **COMPONENTS** page (Infrastructure -> Components -> Add/Modify). Remember that components manually added are managed automatically. Discovered components, however, are managed using the **COMPONENTS - MANAGE/UNMANAGE** page. Figure 2.1 and Figure 2.2 clearly illustrate the process of managing a Microsoft File server.

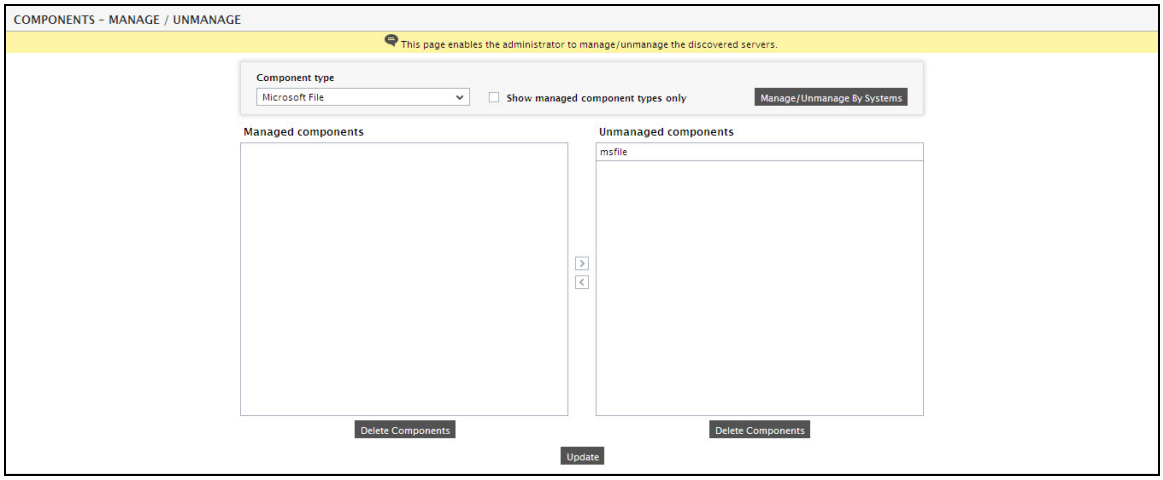

Figure 2.1: Viewing the list of unmanaged Microsoft File servers

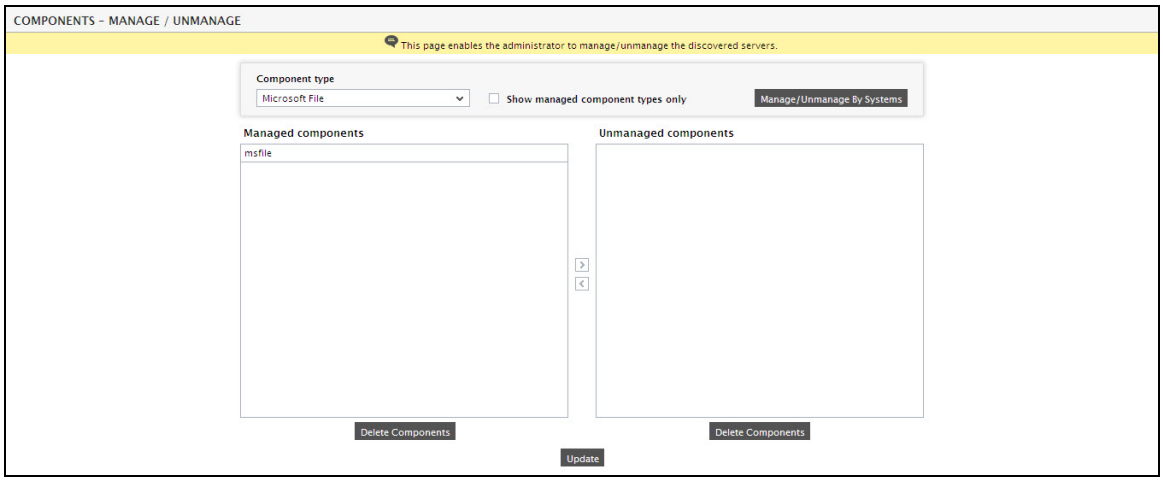

Figure 2.2: Managing an Microsoft File server

3. Next, sign out of the eG administrative interface.

Chapter 3: Monitoring Microsoft File Servers

eG Enterprise provides out-of-the-box a specialized *Microsoft File* server model (see Figure 3.1) that periodically runs diagnostic tests on the file server to ensure that it performs to peak capacity at all times.

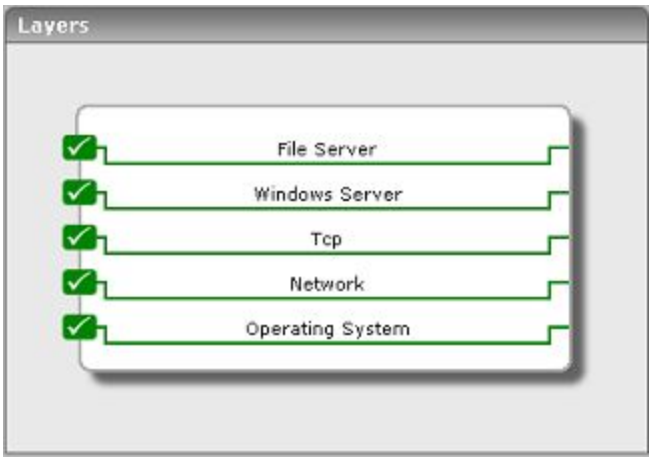

Figure 3.1: Layer model of a Microsoft File server

The sections to come discuss the top 2 layers of Figure 3.1, since the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

3.1 The Windows Server Layer

Using the tests associated with the Windows Server layer, administrators can closely observe the user logins to and session behavior on the MS File server.

Figure 3.2: Tests associated with the Windows Server layer

3.1.1 Windows Access Test

This test monitors the accesses to the Microsoft File server.

Target of the test : A Microsoft File server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the file server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the File server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Blocking request rejects	The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items	Reqs/sec	If the number of blocking request rejects is high, you may need to adjust the *MaxWorkItem* or *MinFreeWorkItems* server parameters
Permission errors	The number of times opens on behalf of clients have failed with *STATUS_ACCESS_ DENIED* in the last measurement period	Number	Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected.
File access denied errors	The number of times accesses to files opened successfully were denied in the last measurement period	Number	This number indicates attempts to access files without proper access authorization.
Internal server errors	This value indicates the	Number	Unexpected errors usually indicate a

Measurement	Description	Measurement Unit	Interpretation
	number of times an internal server error was detected in the last measurement period.		problem with the server.
Data received	The rate at which the server has received data from the network	Kbytes/sec	This metric indicates how busy the server is.
Data transmitted	The rate at which the server has sent data over the network	Kbytes/sec	This metric indicates how busy the server is.
Resource shortage errors	The number of times *STATUS_DATA_NOT_ACCEPTED* was returned to clients in the last measurement period	Number	A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the *InitWorkItems* or *MaxWorkItems* server parameters might need to be adjusted.
Avg response time	Average time taken by the server to respond to client requests	Secs	This is a critical measure of server health.

3.1.2 Windows Sessions Test

This test reports various session-related statistics for a Microsoft File server.

Target of the test : A Microsoft File server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the file server that is being monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.

Parameters	Description
Port	Refers to the port used by the File server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Logons	Rate of logons to the server	Reqs/sec	This measure reports the rate of all interactive, network, and service logons to an MS File server. The measure includes both successful and failed logons.
Logon errors	Number of logons in the last measurement period that had errors	Number	This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server.
Current sessions	The number of sessions currently active in a server	Number	This measure is one of the indicators of current server activity.
Sessions with errors	The number of sessions in the last measurement period that were closed to unexpected error conditions	Number	Sessions can be closed with errors if the session duration reaches the autodisconnect timeout.
Sessions forced off	The number of sessions in the last measurement period that have been forced to logoff	Number	This value indicates how many sessions were forced to logoff due to logon time constraints.
Sessions logged off	The number of sessions in the last measurement period that were terminated normally	Number	Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low.
Sessions timed out	The number of sessions that have been closed in the last measurement	Number	The number of session timed out gives an indication of whether the **AutoDisconnect** setting is helping to

Measurement	Description	Measurement Unit	Interpretation
	period due to their idle time exceeding the **AutoDisconnect** parameter for the server		conserve server resources

3.2 The File Server Layer

With the help of the tests associated with this layer, administrators can:

- Accurately determine the current work load on the server in terms of the number of files currently accessed on the server and the current user traffic to the server
- Quickly identify locked files and the users who have acquired a lock on those files

Figure 3.3: Tests associated with the File server layer

3.2.1 MS File Stats Test

This test tracks various statistics pertaining to open file connections at the host.

Target of the test : A Microsoft File server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the file server that is being monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
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
File locks count	The number of files locked at the host	Number	A high value can indicate too many files being opened at the host. The detailed diagnosis of this measure, if enabled, lists the files that have been locked, the user who holds the lock, and the number of locks on the file.
Unique users count	A unique count of users who have opened files at this host	Number	A high value can indicate too many users connected to the host.

3.2.2 Windows Usage Test

This test tracks various statistics pertaining to sessions open at the host.

Target of the test : A Microsoft File server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the file server that is being monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
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
Open files	The number of files opened over the network by users connecting to the file server	Number	This measurement is an indicator of the workload on the file server. The detailed diagnosis of this measure, if enabled, provides the number of open sessions for every user, and the time for which the sessions have been idle. If the idle time displayed here is very high, then measures for closing the inactive open sessions can be initiated.
Unique users	A unique count of users who have opened sessions at this host	Number	A high value can indicate too many users connected to the host.

3.2.3 HBA Port Stats Test

By periodically checking the target port status and measuring the I/O processing capability of the ports, you can identify overloaded ports, and thus proactively detect potential/existing load-balancing irregularities and/or processing bottlenecks on the host bus adapter. The **HBA Port Stats** test facilitates this port check.

For every port configured on the Host Bus Adapter, this test reports the port state, the I/O processing ability of the ports, and the errors encountered by each port. In the process, the test not only points administrators to overloaded ports, but also puts a finger on ports that are slow in processing I/O requests and the ports that are erroneous.

Target of the test : A Microsoft File server

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each HBA: port combination connecting to the file server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the File server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
HBA port state	Indicates the current state of this port.		The numeric values that correspond to the states are as follows: State: Unknown — Numeric Value: 1 State: Operational — Numeric Value: 2

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>State</th><th>Numeric Value</th></tr>
<tr><td>User Offline</td><td>3</td></tr>
<tr><td>Bypassed</td><td>4</td></tr>
<tr><td>In diagnostics mode</td><td>5</td></tr>
<tr><td>Link Down</td><td>6</td></tr>
<tr><td>Port Error</td><td>7</td></tr>
<tr><td>Loopback</td><td>8</td></tr>
</table>
Note:
By default, the test reports the States listed in the table above to indicate the current state of the port. In the graph of this measure however, the state is indicated using the corresponding numeric equivalents only.
</td></tr>
<tr><td>HBA port type</td><td>Indicates the current type of this port.</td><td></td><td>
The values that this measure can take and their corresponding numeric values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Unknown</td><td>1</td></tr>
<tr><td>Other</td><td>2</td></tr>
<tr><td>Not present</td><td>3</td></tr>
<tr><td>Fabric</td><td>5</td></tr>
<tr><td>Public Loop</td><td>6</td></tr>
<tr><td>Fabric on a loop</td><td>7</td></tr>
<tr><td>Fabric Port</td><td>8</td></tr>
<tr><td>Fabric expansion port</td><td>9</td></tr>
<tr><td>Generic Fabric Port</td><td>10</td></tr>
<tr><td>Private Loop</td><td>20</td></tr>
<tr><td>Point to Point</td><td>21</td></tr>
</table>
Note:
</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			By default, this measure reports one of the **Measure Value**s listed in the table above. In the graph of this measure however, the port type is indicated by the corresponding numeric equivalents only.
HBA port operational speed	Indicates the current operational speed of this port.	Gbps	
Received frames	Indicates the rate at which frames were received by this port during the last measurement period.	Frames/sec	Compare the value of these measures across the ports to identify the slowest port in terms of receiving / transferring frames.
Transfer frames	Indicates the rate at which frames were transferred from this port during the last measurement period.	Frames/sec	
Received words	Indicates the rate at which words were received by this port during the last measurement period.	Words/sec	Compare the value of these measures across the ports to identify the slowest port in terms of receiving / transferring words.
Transfer words	Indicates the rate at which words were transferred from this port during the last measurement period.	Words/sec	
Error frames	Indicates the rate at which frames were received / transferred with errors from this port during the last measurement period.	Frames/sec	Compare the value of this measure across the ports to identify the port that is more error prone.
Dumped frames	Indicates the number of frames dumped by this port per second	Frames/sec	Buffer credits, also called buffer-to-buffer credits (BBC) are used as a flow control method by Fiber Channel technology and

Measurement	Description	Measurement Unit	Interpretation
	due to lack of buffer credit during the last measurement period.		represent the number of frames a port can store. Each time a port transmits a frame that port's BB Credit is decremented by one; for each R_RDY received, that port's BB Credit is incremented by one. Transmission of an R_RDY indicates that the port has processed a frame, freed a receive buffer, and is ready for one more. If the BB Credit is zero, the corresponding node cannot transmit until an R_RDY is received back. A high value for this measure therefore indicates that an R_RDY was not received by the FC port for a long time. This is a cause for concern, as until the R_RDY is received, the FC port will not resume communication. The solution for this problem is to allocate optimal buffer credits to the FC port. The optimal number of buffer credits is determined by the distance (frame delivery time), the processing time at the receiving port, the link signaling rate, and the size of the frames being transmitted. As the link speed increases, the frame delivery time is reduced and the number of buffer credits must be increased to obtain full link utilization, even in a short-distance environment. Smaller frame sizes need more buffer credits.
Loop initialization primitive events	Indicates the number of Loop initialization primitive events on this port during the last measurement period.	Number	
Non-operational state primitive events	Indicates the number of Non-operational state (NOS) primitive	Number	

Measurement	Description	Measurement Unit	Interpretation
	events on this port during the last measurement period.		
Link failure count	Indicates the number of link failures experienced per second by this target during the last measurement period.	Number	Ideally, the value of this measure should be zero. A non-zero value indicates that Fiber Channel connectivity with this target was "broken" that many times. This is likely an indicator for a faulty connector or cable. These are also caused when the device connected to this target is restarted, replaced or being serviced when the Fiber Channel cable connected to this target is temporarily disconnected.
Loss of synchronization count	Indicates the number of times this port failed to synchronize during the last measurement period.	Number	Ideally, the value of this measure should be zero. A non-zero value for this measure indicates that port went into the "loss of synchronization" state, where it encountered continuous Disparity errors. This is likely an indicator for a faulty connector or cable. These are also caused when the device connected to the port is restarted, replaced or being serviced when the Fiber Channel cable connected to the port is temporarily disconnected. If the port is in the "loss of synchronization" state for longer than a specific period, the port will get into the link failure state which could degrade the performance of the Fiber Channel link.
Loss of signal count	Indicates number of signals lost on this port during the last measurement period.	Number	Ideally, the value of this measure should be zero. A non-zero value for this measure indicates that the port detected a loss of the electrical or optical signal used to transfer data on the port. This is likely an indicator for a faulty connector or cable. These are also caused when the device connected to the port is

Measurement	Description	Measurement Unit	Interpretation
			restarted, replaced or being serviced when the Fiber Channel cable connected to the port is temporarily disconnected. If the port is in the "loss of signal" state for longer than a specific period, the port will get into the link failure state which could degrade the performance of the Fiber Channel link.
Primitive sequence protocol errors	Indicates the number of Primitive sequence protocol errors that occurred on this port during the last measurement period.	Number	Ideally, the value of this measure should be zero.
Invalid word transmissions	Indicates the number of invalid word transmissions detected on this port per second during the last measurement period.	Words/sec	A low value is desired for this measure.
Cyclic redundancy checksum errors	Indicates the number of invalid Cyclic Redundancy Checksums that occurred on this port during the last measurement period.	Number	Ideally, the value of this measure should be zero. A high value for this measure indicates poor health of the target port. These are usually recoverable errors and will not degrade system performance unless their occurrence is sustained when the data cannot be relayed after retransmissions.

3.2.4 HBA Status Test

A host bus adapter (HBA) is a circuit board or integrated circuit adapter that provides input/output (I/O) processing and physical connectivity between a host system, or server, and a storage or network device. Failure of the host bus adapter may abruptly stop the I/O processing between the host system and the storage/network device. It is therefore necessary to keep a vigil on the status of the host bus adapter round the clock. The **HBA Status** test helps administrators in this regard.

This test auto-discovers the HBAs in the target environment and for each HBA, reports the status and the count of ports connecting to the target file server. Administrators can also figure out the HBAs that are currently inactive. Using this test, administrators can easily identify the HBAs that are error-prone and are frequently inactive.

Target of the test : A Microsoft File server

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each HBA connecting to the file server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the File server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is HBA active?	Indicates whether/not this HBA is active.		If the HBA is active, then this measure will return the value *Yes*. If the HBA is inactive or not available, then this measure will report a value *No*. The numeric values that correspond to these measure values are as follows: **Measure Value** / **Numeric Value**: No = 0; Yes = 1 **Note:** By default, the test reports the **Measure Value**s listed in the table above to indicate whether/not this HBA is active. In the graph of this measure however, the same is

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>indicated using the corresponding numeric equivalents only.</td></tr>
<tr><td>HBA status</td><td>Indicates the current status of this HBA.</td><td></td><td>The values that this measure can take and their corresponding numeric values are as follows:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Ok</td><td>0</td></tr><tr><td>Error</td><td>1</td></tr><tr><td>Not Supported</td><td>2</td></tr><tr><td>Invalid Handle</td><td>3</td></tr><tr><td>Bad Argument</td><td>4</td></tr><tr><td>WWN Not Recognized</td><td>5</td></tr><tr><td>Index Not Recognized</td><td>6</td></tr><tr><td>Larger Buffer Required</td><td>7</td></tr><tr><td>Information Has Changed</td><td>8</td></tr><tr><td>SCSI Check Condition Reported</td><td>9</td></tr><tr><td>Adapter Busy/Reserved</td><td>10</td></tr><tr><td>Request Timed out</td><td>11</td></tr><tr><td>HBA Removed/Deactivated</td><td>12</td></tr><tr><td>ELS Rejected</td><td>13</td></tr><tr><td>Invalid LUN</td><td>14</td></tr><tr><td>Incompatible</td><td>15</td></tr><tr><td>Ambiguous WWN</td><td>16</td></tr><tr><td>Bad Local SCSI Bus Number</td><td>17</td></tr><tr><td>Bad Local SCSI Target Number</td><td>18</td></tr><tr><td>Bad Local SCSI LUN</td><td>19</td></tr><tr><td>Local SCSI ID</td><td>20</td></tr></table></td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Invalid FCP Target FCID</td><td>21</td></tr>
<tr><td>Bad FCP Target Node WWN</td><td>22</td></tr>
<tr><td>Bad FCP Target Port WWN</td><td>23</td></tr>
<tr><td>FCP LUN Not Regonize by Target</td><td>24</td></tr>
<tr><td>Inaccessible LUID by Target</td><td>25</td></tr>
<tr><td>No Such Binding</td><td>26</td></tr>
<tr><td>Not a Target</td><td>27</td></tr>
<tr><td>Unsupported FC-4 protocol</td><td>28</td></tr>
<tr><td>Incapable Port</td><td>29</td></tr>
<tr><td>SCSI Target Busy</td><td>30</td></tr>
<tr><td>HBA Free Library</td><td>31</td></tr>
<tr><td>HBA Load Library</td><td>32</td></tr>
<tr><td>HBA Send RNIDV2</td><td>33</td></tr>
</table>
Note:
By default, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the status of the HBA is indicated by the corresponding numeric equivalents only.
</td></tr>
<tr><td>HBA ports</td><td>Indicates the number of ports in this HBA.</td><td>Number</td><td></td></tr>
</table>

3.3 The DFS Replication Service Layer

The tests mapped to this layer monitor the DFS replication service and reports abnormalities (if any).

3.3.1 DFS Replication Folders Test

DFS Replication is an efficient, multiple-master replication engine that you can use to keep folders synchronized between servers across limited bandwidth network connections.

DFS Replication uses a compression algorithm known as remote differential compression (RDC). RDC detects changes to the data in a file and enables DFS Replication to replicate only the changed file blocks instead of the entire file.

To use DFS Replication, you must create replication groups and add replicated folders to the groups. Replication groups, replicated folders, and members are illustrated in the following figure.

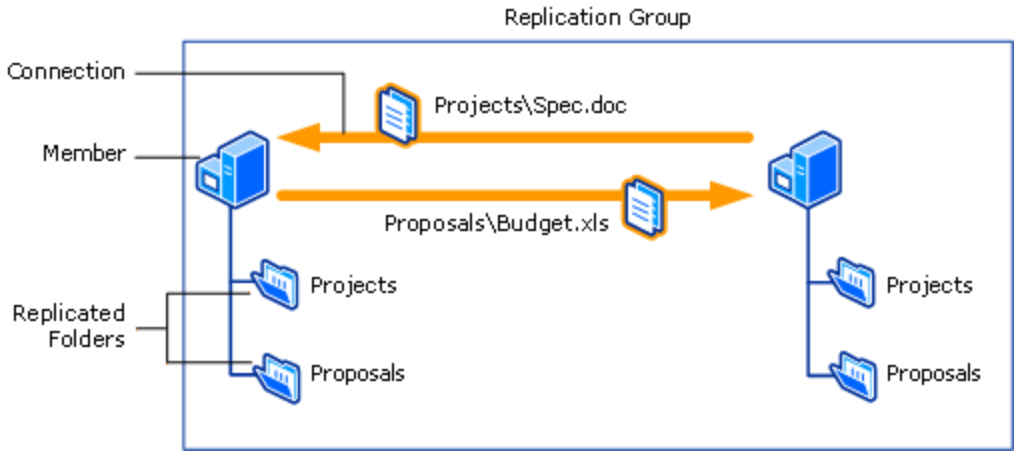

Figure 3.4: How DFS replication works

This figure shows that a replication group is a set of servers, known as members , which participates in the replication of one or more replicated folders. A replicated folder is a folder that stays synchronized on each member. In the figure, there are two replicated folders: Projects and Proposals. As the data changes in each replicated folder, the changes are replicated across connections between the members of the replication group. The connections between all members form the replication topology.

DFS Replication uses staging folders to act as caches for new and changed files to be replicated from sending members to receiving members. The sending member begins staging a file when it receives a request from the receiving member. The process involves reading the file from the replicated folder and building a compressed representation of the file in the staging folder. This is the *staged file*. After being constructed, the staged file is sent to the receiving member; if remote differential compression [RDC] is used, only a fraction of the staging file might be replicated. The receiving member downloads the data and builds the file in its staging folder. After the file has completed downloading on the receiving member, DFS Replication decompresses the file and installs it into the replicated folder. Each replicated folder has its own staging folder, which by default is located under the local path of the replicated folder in the DfsrPrivate\Staging folder.

DFS Replication uses a "last-writer wins" method for determining which version of a file to keep when a file is modified on two or more members. The losing file is stored in the Conflict and Deleted

folder on the member that resolves the conflict. This member might not be the member where the changes originated. Each replicated folder has its own Conflict and Deleted folder, which is located under the local path of the replicated folder in the DfsrPrivate\ConflictandDeleted folder.

Slow replication can cause replicated folders to remain out-of-sync across members for long time periods. Because of this, users may end up receiving an inconsistent/incomplete view of data. Typically, replication may slow down owing to the following:

- The lack of adequate bandwidth resources;
- Staging and conflict and deleted folders that have not been sized to deal with high levels of replication activity;

By closely monitoring these factors, administrators can detect bottlenecks to replication early and take pre-emptive measures. This is where the **DFS Replication Folders** test helps. This test auto-discovers replicated folders and for each folder reports the bandwith saved during replication and tracks the growth in size of the associated staging folders and config and deleted items folders. This way, administrators can proactively detect probable slowdowns in replication. In addition, they can pinpoint what is causing the slowdown and which replication folders will be affected by this.

Target of the test : A server that hosts the DFS namespace (this can even be a server that contains the Dfs root or a replica of it)

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each replicated folder

Configurable parameters for the test

Parameters	Description
Test period	This indicates 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 to.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Bandwidth savings using DFS replication	Indicates the percentage of bandwidth that was saved by the DFS	Percent	A high degree of compression results in high bandwidth savings and faster/inexpensive replication. A high

Measurement	Description	Measurement Unit	Interpretation
	Replication service for the replicated folder. .		value is therefore desired for this measure. A low value is often an outcome of a poor compression algorithm.
Conflict files generated	Indicates the number of files and folders in this replicated folder that were moved to the conflict and deleted folder during the last measurement period.	Number	The value of this measure is a good indicator of how often the files in a replication folder are changed.
Conflict files cleanups	Indicates the number of conflict loser files and folders that were deleted from the conflict and deleted folder of this replicated folder during the last measurement period.	Number	The Conflict and Deleted folder has high and low watermarks (90 percent and 60 percent of Conflict and Deleted folder quota, respectively) that govern cleanup and excessive usage of the folder. This implies that if the size of the Conflict and Deleted folder is over 90% of the quota configured for it, DFS will start deleting files from the folder automatically and will keep at it until the size of the folder touches its low watermark – i.e., until size becomes 60% of the configured quota. If the value of this measure is consistently high for a replicated folder, it can only mean that the Conflict and Deleted folder is continuously growing in size, thus forcing DFS to cleanup old files from it on a regular basis. This could also indicate that files in the corresponding replicated folder are changing very often, thus adding to the contents and the size of the Conflict and Deleted folder. Another reason for a large number of cleanups could be a

Measurement	Description	Measurement Unit	Interpretation
			low quota configured for the Conflict and Deleted folder.
Deleted files generated	Indicates the number of replicated deleted files and folders that were moved from this replicated folder to the conflict and deleted folder during the last measurement period.	Number	The Conflict and Deleted folder can also be used to store files that are deleted from replicated folders.
Deleted files cleanups	Indicates the number of replicated deleted files and folders that were deleted from the conflict and deleted folder of this replicated folder during the last measurement period. .	Number	Deleted files are treated like conflict files in that they are purged when the Conflict and Deleted folder reaches 90 percent of the configured quota. If the size of the Conflict and Deleted folder grows above 90% of its configured quota, DFS will start deleting files from the folder automatically and will keep at it until the size of the folder touches its low watermark – i.e., until size becomes 60% of the configured quota. If the value of this measure is consistently high for a replicated folder, it can only mean that the Conflict and Deleted folder is continuously growing in size, thus forcing DFS to cleanup old files from it on a regular basis. This could also indicate that files in the corresponding replicated folder are either changing or are being deleted very often, thus adding to the contents and the size of the Conflict and Deleted folder.
Staging files generated	Indicates the number of times files in this replicated folder were	Number	The value of this measure is a good indicator of the number of times files were added to or changed in a

Measurement	Description	Measurement Unit	Interpretation
	staged during the last measurement period.		replication folder. A high value is therefore indicative of a high level of replication activity. A consistent and significant increase in the value of this measure over time could also indicate that staging files are not being sent to the receiving member as quickly as they are created by the sending member – this in turn could indicate a bottleneck in replication.
Staging files cleanups	Indicates the number of files and folders that were cleaned up from the staging folder of this replication folder by the DFS Replication service during the last measurement period.	Number	The DFS Replication service stages files and folders in the staging folder before they are replicated, and automatically cleans up the staging folder when it exceeds a pre-configured threshold of the quota. If the value of this measure remains high consistently, it could indicate that the staging folder is rapidly exhausting its configured quota size, causing DFS to keep deleting old staging files from the folder. While this can happen if the level of replication activity is high, it can also happen if the staging folder is configured with a low quota size. If a staging folder quota is configured to be too small, DFS Replication might consume additional CPU and disk resources to regenerate the staged files. Replication might also slow down because the lack of staging space can effectively limit the number of concurrent transfers with partners.
Conflict space in use	Indicates the total size of conflict loser files and folders in the conflict and	MB	The quota size of the Conflict and Deleted folder is 660 MB by default;

Measurement	Description	Measurement Unit	Interpretation
	deleted folder of this replicated folder during the last measurement period.		however, this can be changed. Also, the Conflict and Deleted folder is configured with a high watermark of 90% and a low watermark of 60% of the quota. If the value of this measure is over 90% of the default/configured (as the case may be) quota size, then DFS will start purging files from the conflict and deleted folder until its size falls to 60% of the quota. Compare the value of this measure across replication folders to know which replication folder's conflict and deleted folder is growing close to its configured quota size.
Deleted space in use	Indicates the total size of replicated deleted files and folders currently in the conflict and deleted folder of this replicated folder.	MB	Deleted files are treated like conflict files in that they are purged when the Conflict and Deleted folder reaches 90 percent of the configured quota.
Staging space in use	Indicates the total size of the files and folders in the staging folder of this replication folder during the last measurement period.	MB	The default size of each staging folder is 4,096 MB. This is not a hard limit, however. It is only a quota that is used to govern cleanup and excessive usage based on high and low watermarks (90 percent and 60 percent of staging folder size, respectively). This quota can be changed. If the value of this measure reaches 90% of the configured quota, the oldest staged files are purged until the staging folder reaches 60 percent of the configured quota. For good operational performance, increasing the quota size of a staging folder is recommended when you have

Measurement	Description	Measurement Unit	Interpretation
			multiple large files that change frequently. Microsoft also recommends that you increase the staging folder quota on hub members that have many replication partners. If a staging folder quota is configured to be too small, DFS Replication might consume additional CPU and disk resources to regenerate the staged files. Replication might also slow down because the lack of staging space can effectively limit the number of concurrent transfers with partners.

3.3.2 DFS Namespace Referrals Test

A **DFS Namespace** is a virtual view of shared folders on different servers as provided by DFS. A DFS namespace consists of a root and many links and targets. The namespace starts with a root that maps to one or more root targets. Below the root are links that map to their own targets.

A **Referral** is a list of targets, transparent to the user, which a DFS client receives from DFS when the user is accessing a root or a link in the DFS namespace.

User experience with DFS is often measured by how quickly DFS processes referral requests. Prolonged delays in processing combined with frequent request failures can leave users frustrated. To prevent this, administrators must track the referral requests to each namespace managed by DFS, keep an eye out for request failures, monitor how DFS handles these requests, and promptly capture abnormalities. This is exactly what the **DFS Namespace Referrals** test does. For each namespace, this test reports the count of referral requests to the namespace, the percentage of requests that failed, and how swiftly the requests were processed. Besides proactively alerting administrators to processing bottlenecks, the test also leads them to the exact namespace where the problem lies.

Target of the test : A server that hosts the DFS namespace (this can even be a server that contains the Dfs root or a replica of it)

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each DFS namespace to which the namespace server being monitored belongs.

Configurable parameters for the test

Parameters	Description
Test period	This indicates 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 to.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Avg response time	Indicates the average time taken by the DFS namespace service to respond to referral requests to this namespace during the last measurement period.	Secs	Ideally, the value of this measure should be low. A consistent increase in the value of this measure is indicative of poor responsiveness of the DFS namespace service. Compare the value of this measure across namespaces to know which namespace is responding slowly to referral requests.
Requests processed	Indicates the number of referral requests processed by the DFS namespace service for this namespace during the last measurement period.	Number	
Requests failed	Indicates the number of requests for this namespace that failed.	Number	Ideally, the value of this measure should be 0.
Percentage of failed requests	Indicatesthe percentage of referral requests to this namespace that failed.	Percent	The value 0 is desired for this measure. A high value is a cause for concern. Compare the value of this measure across namespaces to know requests to which namespace are failing often.

Measurement	Description	Measurement Unit	Interpretation
Request processing rate	Indicates the rate at which the referral requests to this namespace are processed by the DFS namespace service.	Requests/Sec	A high value is desired for this measure. A steady dip in this value is indicative of a processing bottleneck. Compare the value of this measure across namespaces to know the requests that are being serviced very slowly.

3.3.3 DFS Replication Volumes Test

In DFS, replication folders are hosted on volumes. The DFS replication service maintains one ESE (Extensible Storage Engine) database per volume. This database is used to store metadata about each file and folder in the replicated folder.

How quickly the DFSR performs replication depends upon how well the volume (i.e., the ESE database in the volume) is sized. Because volumes are to be sized on the basis of the transaction load that replication imposes on them, administrators may first want to ascertain the level of database activity that replication typically generates per volume and then plan volume capacity accordingly. The DFS Replication Volumes test helps administrators with this exercise.

For every volume that hosts replication folders, this test reports the count of database lookups and commits performed on the volume since the last measurement period. By tracking variations to these measures over time, administrators can easily understand the current load, forecast the future load, and use these results to determine/fine-tune the volume size.

Target of the test : A server that hosts the DFS namespace (this can even be a server that contains the Dfs root or a replica of it)

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each volume hosting replication folders.

Configurable parameters for the test

Parameters	Description
Test period	This indicates 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 to.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Database lookups	Indicates the number of database search operations performed by the DFS replication service on this volume since the last measurement period.	Number	These measures are good indicators of the level of database activity that replication generates on a volume. Compare the value of these measures across volumes to know which volume is used the most. For this volume, closely track variations in the value of these measures over time, understand load changes, and accordingly right-size (if required) the volume.
Database commits	Indicates the number of database commit operations performed by the DFS replication service on this volume since the last measurement period.	Number	

3.3.4 DFS Namespace API Requests Test

The management (eg., DFS Management UI, DFS BPA, DFS Namespaces Windows PowerShell cmdlets, File Server Resource Manager) used by administrators to manage DFS typically employ NetDfs API functions to perform critical management tasks – eg., creating/removing namespaces, adding roots and links, viewing/configuring information about roots and links, and so on. In the process, these tools impose a significant amount of operational load on the namespace server. If the server is not sized right to handle this load, API requests may either fail frequently or may end-up being processed slowly by the server, thus increasing the DFS management challenges of administrators. This in turn will impact the user experience with DFS. If such situations are to be avoided, administrators should use the **DFS Namespace API Requests** test to track the number and type of API requests to the namespace server, measure how quickly the server processes requests of each type, identify request types that failed often, and accordingly, decide whether/not the server should be resized to improve its processing ability.

Target of the test : A server that hosts the DFS namespace (this can even be a server that contains the DFS root or a replica of it)

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each type of API request to the namespace server being monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates 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 to.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Avg response time	Indicates the average time taken by the DFS namespace server to respond to API requests of this type.	Secs	Ideally, the value of this measure should be low. A consistent increase in the value of this measure is indicative of poor responsiveness of the DFS namespace service. Compare the value of this measure across request types to know which API requests are being processed slowly by the server.
Requests processed	Indicates the number of API requests of this type processed by the server since the last measurement period.	Number	
Requests failed	Indicates the number of API requests of this type that failed during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Percentage of failed requests	Indicates the percentage of API requests of this type that failed.	Percent	The value 0 is desired for this measure. A high value is a cause for concern. Compare the value of this measure across types to know which type of requests is prone to failures.
Request processing rate	Indicates the rate at which the API requests of this type are processed.	Requests/Sec	A high value is desired for this measure. A steady dip in this value is indicative of a processing bottleneck. Compare the value of this measure

Measurement	Description	Measurement Unit	Interpretation
			across types to know the requests that are being serviced very slowly.

3.3.5 DFS Namespace API Queue Test

To manage DFS, administrators often rely on management tools and functions that Windows offers - for eg., DFS Management UI, DFS BPA, DFS Namespaces Windows PowerShell cmdlets, File Server Resource Manager, and many more. All these management tools use NetDfs API functions to perform a wide variety of management tasks such as:

- Adding a DFS link to a DFS root;
- Creating or removing stand-alone and domain-based DFS namespaces;
- Adding targets to an existing DFS link;
- Removing a DFS link from a DFS root;
- Removing a target from a DFS link;
- Viewing and configuring information about DFS roots and links

In the process, these tools can generate a significant amount of management traffic on the namespace server. The true test of strength of the namespace server lies in how well it processes this traffic. To determine whether/not the namespace server is able to handle this traffic, administrators can use the **DFS Namespace API Queues** test. This test continuously tracks the length of the API queues on the server to figure out how quickly the server processes the API requests in the queue. This sheds light on the load imposed by API requests on the server and the ability of the server to respond to these requests.

Target of the test : A server that hosts the DFS namespace (this can even be a server that contains the DFS root or a replica of it)

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the DFS namespace server being monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates 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 to.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
API queue length	Indicates the number of requests (made using the NetDfs API) currently in the queue for the DFS Namespace service to process.	Number	If the value of this measure keeps increasing while at the peak level of activity, it could indicate a bottleneck in processing API requests. One of the common reasons why a namespace server may be unable to process API requests quickly is improper server sizing. A server that is sized right should be able to crank through its work queues and be responsive. A server with insufficient resources on the other hand will not be able to handle this load, and may hence end-up processing requests slowly; this in turn will increase the length of the API request processing queues. In such cases, add more processing power to the server and see if it helps reduce queue length.

3.3.6 DFS Replication Connections Test

The replication topology consists of the logical connections that DFS Replication uses to replicate files among servers. Members in a replication topology communicate via two one-way connections. These two connections allow data to flow in both directions.

Where replication activities consume considerable bandwidth, administrators may want to quickly identify the bandwidth-intensive replication connections, determine how much bandwidth is currently saved by these connections by compressing the traffic, and decipher how the compression algorithm can be tweaked to conserve more bandwidth. This is exactly what the **DFS Replication Connections** test does. This test monitors each replication connection to a namespace server and

reports the throughput and bandwidth savings of each connection, so that administrators can identify those connections that could be candidates for additional traffic compression.

Target of the test : A server that hosts the DFS namespace (this can even be a server that contains the Dfs root or a replica of it)

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each replication connection to the namespace server monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates 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 to.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total files received	Indicates the number of files received on this connection during the last measurement period.	Number	Compare the value of this measure across connections to know which connection receives the maximum number of files.
Total data received	Indicates the total number of bytes received on this connection during the last measurement period.	MB	Compare the value of this measure across connections to know which connection receives the maximum data.
RDC data received	Indicates the bytes that were received on this connection while replicating files using remote differential compression during the last measurement period.	MB	This is the actual bytes received over the network without the networking protocol overhead. By comparing the value of this measure with that of the Total data received measure for a connection, administrators can determine whether/not data compression has improved throughtput and reduced

Measurement	Description	Measurement Unit	Interpretation
			bandwidth consumption of that connection.
Bandwidth savings using DFS replication	Indicates the percentage of bandwidth that was saved by the DFS Replication service.	Percent	This measure indicates how much bandwidth was saved for this connection using a combination of remote differential compression (RDC) and other compression technologies that minimize network bandwidth use. For example, a value of 20 indicates that the DFS Replication service used 20% less bandwidth than it would have used if it had transmitted the entire files uncompressed over the network. By comparing the value of this measure across connections, administrators can identify that connection which has saved the least bandwidth. For such connections, the compression technologies employed may have to be tweaked further to ensure optimal bandwidth usage.

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.