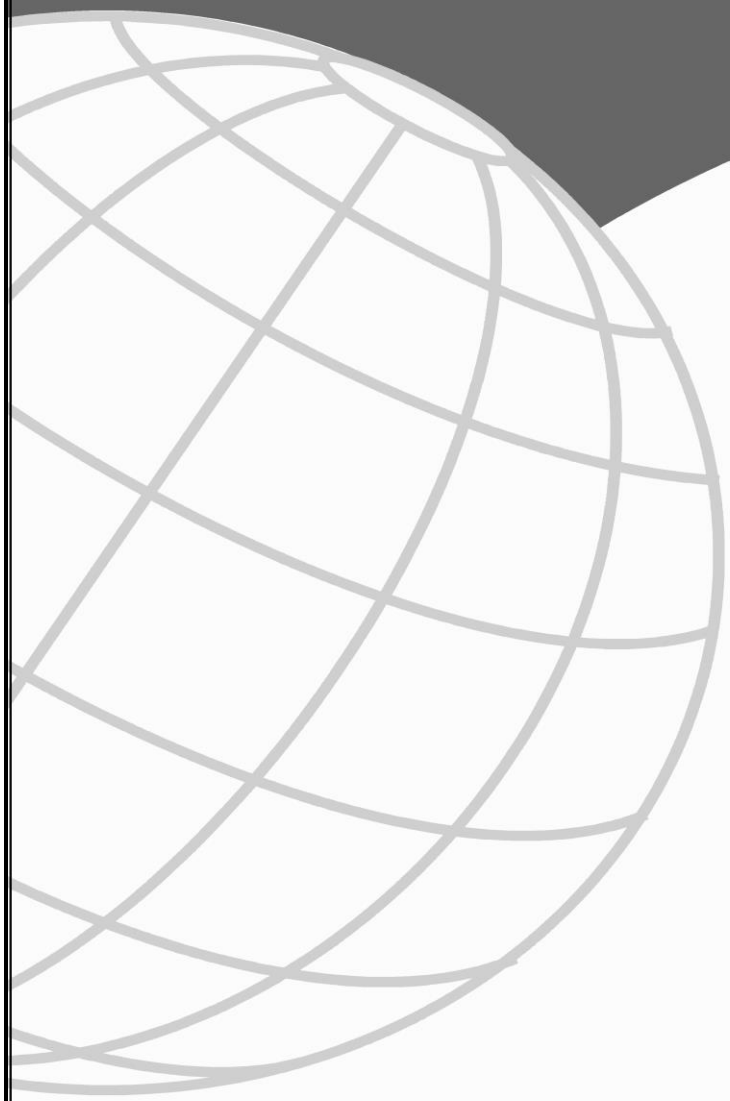

Monitoring the NetApp USD

Table of Contents

MONITORING THE NETAPP UNIFIED STORAGE 1

1.1 How does eG Enterprise Monitor the NetApp Unified Storage? 2

1.1.1 Configuring the eG Agent to Receive SNMP Traps from the NetApp Unified Storage 3

1.1.2 Configuring the eG Agent to Poll the SNMP MIB of the NetApp Unified Storage to Pull Out the Metrics 3

1.1.3 Configuring the eG Agent to Use the NetApp Manageability SDK 3

1.1.4 Creating a New User with the Privileges Required for Monitoring the NetApp Unified Storage 5

1.2 The Hardware Layer 6

1.2.1 Failure Traps Test 7

1.2.2 Shutdown Traps Test 10

1.2.3 Warning Traps Test 13

1.2.4 NetApp System Components Test 16

1.2.5 NetApp Environment Test 21

1.3 The Network Layer 25

1.4 The Physical Storage Layer 25

1.4.1 NetApp Disks Test 26

1.4.2 Ungrouped Disks 31

1.4.3 NetApp Aggregates Test 34

1.4.4 Raid Groups Test 47

1.4.5 Disk Health Monitor Events 51

1.5 The NetApp OS Layer 54

1.5.1 Consistency points Test 54

1.5.2 NetApp File Layouts (WAFL) Test 56

1.6 The NetApp Access Layer 70

1.6.1 NetApp Block I/O Protocol 70

1.6.2 NetApp iSCSI Connections Test 75

1.6.3 NetApp iSCSI Protocol Test 77

1.7 The File Access Protocols Layer 84

1.7.1 CIFS Test 84

1.7.2 NetApp IGroup Config Mismatches Test 88

1.7.3 NetApp NFS I/O Test 93

1.8 The Logical Storage Layer 98

1.8.1 NetApp Volume Details Test 99

1.8.2 Busy Snapshots Test 108
1.8.3 NetApp High Utilization Quotas Test 113
1.8.4 NetApp Clone Operations Test 117
1.8.5 NetApp LUN Config Errors Test 120
1.8.6 NetApp LUNs Test 122
1.9 The NetApp System Layer 128
1.9.1 Virus Scanner Stats Test 128
1.9.2 System Status Test 131
1.9.3 NetApp Fiber Channel Adapters Test 136
1.9.4 NetApp Initiator Config Mismatches 144
1.9.5 NetApp Syslog Test 147
1.9.6 NetApp System Performance Test 150
CONCLUSION 154

Chapter 1

Monitoring the NetApp Unified Storage

NetApp storage systems are hardware- and software-based data storage and retrieval systems. They respond to network requests from clients and fulfill them by writing data to or retrieving data from the disk arrays. They provide a modular hardware architecture running the Data ONTAP operating system and WAFL (Write Anywhere File Layout) software.

The NetApp storage system consists of the following components:

- The storage system main unit, or chassis, is also known as the storage engine. It is the hardware device that receives and sends data. This unit also houses the storage system components and detects and gathers information about the hardware and the hardware configuration, the storage system components, operational status, hardware failures, and error conditions.
- The disk shelves are the containers, or device carriers, that hold disks and associated hardware (such as power supplies, connectivity, and cabling) that are connected to the main unit of the storage systems.

More specifically, the NetApp storage system includes:

- Internal components such as the system board, system memory, NVRAM, boot device, LCD and LEDs, environmental adapters, etc.
- Slots and ports
- Disk shelves and disks

Owing to their high availability and efficient load distribution features, the NetApp storage system is very popular in large, mission-critical IT infrastructures, which require ready and reliable storage services. In such environments, the non-availability of the storage system or any of its core components, rapid erosion of storage space provided by the storage system, and inconsistencies in I/O load-balancing across disks/LUNs/RAIDs can result in short/prolonged delays in the delivery of storage services, which will ultimately slowdown the dependent end-user services. To avoid this, it is imperative to watch out for issues in the operations and usage of the storage system on a regular basis.

eG Enterprise provides out-of-the-box monitoring for the NetApp storage system. The comprehensive *NetApp Unified Storage* monitoring model offered by the eG Enterprise Suite monitors various aspects of the performance of the

NetApp storage system and promptly alerts storage administrators to potential I/O processing bottlenecks or space crunches.

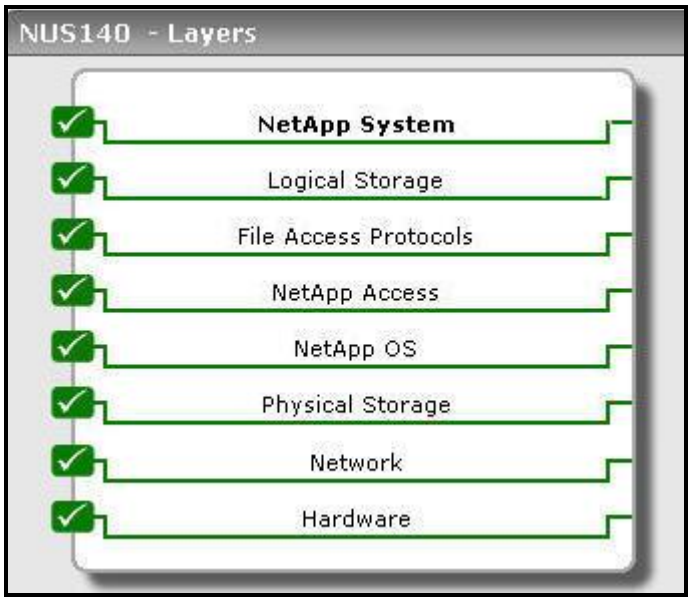

Figure 1: The layer model of the NetApp Unified Storage

The metrics so collected report on the following:

- The performance of the USD hardware
- The health of the network components that interface with (and depend on) the hardware
- The current status and space usage of physical storage entities (disks, raid groups etc)
- The status of Data ONTAP, the NetApp OS which runs on these physical entities
- The efficiency of the access framework and protocols that are used to access the USD - eg., FCP and iSCSI
- The NAS aspect of the USD; this includes the file access protocols such as CIFS and NFS
- Abnormalities related to the usage of logical storage entities (such as volumes, LUNs, Qtrees etc.,) which are accessed using their framework by the end users;
- The overall USD status and performance which depends on all its underlying components

1.1 How does eG Enterprise Monitor the NetApp Unified Storage?

In order to monitor a NetApp Unified Storage, eG uses best-of-both-worlds combination of SNMP and NetApp Manageability SDK. While a few tests intercept SNMP traps to obtain information of interest, a few others directly poll the SNMP MIB of the device to pull out the metrics. Most tests however run the NetApp Manageability SDK commands for metrics collection.

The metrics so collected report on the following:

- The performance of the USD hardware
- The health of the network components that interface with (and depend on) the hardware
- The current status and space usage of physical storage entities (disks, raid groups etc)
- The status of Data ONTAP, the NetApp OS which runs on these physical entities
- The efficiency of the access framework and protocols that are used to access the USD - eg., FCP and iSCSI
- The NAS aspect of the USD; this includes the file access protocols such as CIFS and NFS
- Abnormalities related to the usage of logical storage entities (such as volumes, LUNs, Qtrees etc.,) which are accessed using their framework by the end users;
- The overall USD status and performance which depends on all its underlying components

To know how to configure each of these monitoring mechanisms, refer to the sections below.

1.1.1 Configuring the eG Agent to Receive SNMP Traps from the NetApp Unified Storage

Whenever an SNMP agent detects an error in an SNMP-enabled network device / application, it sends SNMP traps with the error information to a daemon process known as the SNMP Trap Receiver (Snmptrapd). In the eG Enterprise system, the external agent includes an optional SNMP trap receiver that can log traps it receives into a log file which can be parsed/interpreted by the external agent. Therefore, to enable the eG external agent which externally monitors the NetApp Unified Storage to intercept SNMP traps sent out by that device, you need to **setup Snmptrapd on the external agent host**. The procedure for setting up Snmptrapd differs according to the operating system of the external agent host. For detailed setup procedures per operating system, refer to the *Handling SNMP Traps using eG Enterprise* document. Once the Snmptrapd is setup and started, you need to configure the following tests to integrate with Snmptrapd for pulling out the desired metrics:

a. Failure Traps test

b. Shutdown Traps test

c. Warning Traps test

To know how to configure these tests, refer to Section 1.2.1, Section 1.2.2, and Section 1.2.3 of this document.

1.1.2 Configuring the eG Agent to Poll the SNMP MIB of the NetApp Unified Storage to Pull Out the Metrics

You can configure tests to periodically poll the SNMP MIB of the NetApp Unified Storage for collecting metrics of interest. For this, **you have to SNMP-enable the NetApp Unified Storage**.

1.1.3 Configuring the eG Agent to Use the NetApp Manageability SDK

The NetApp Manageability SDK (NMSDK) provides resources to develop applications that monitor and manage NetApp storage systems.

Many tests that execute on the NetApp Unified Storage run API commands provided by this SDK to extract the performance metrics.

To run these commands, the following pre-requisites need to be fulfilled:

- An **eG remote agent** should be installed on a remote Windows/Unix host in the environment. This remote agent should be assigned to the target storage device when managing that device using the eG administrative interface.
- The eG remote agent should be able to access the target storage device.
- The NMSDK should be available on the eG remote agent host. To achieve this, follow the steps discussed below:
    - Download the NMSDK from the following URL to any location on the remote agent host:

      http://support.netapp.com/NOW/cgi-bin/software

      To download the NMSDK, you will have to create a NOW login; to achieve this, go to the following URL:

      http://support.netapp.com

      Note:

      While you download the NMSDK, you should select the platform for which the download is applicable. For an eG agent to collect metrics, you should select **All platforms** option as shown in Figure 1.2.

Figure 1.2: Selecting the All Platforms option to download the NMSDK

    - The NMSDK will be downloaded as a zip file named **netapp-manageability-sdk-<SDK_version>.zip**. Extract the contents of the zip file to any location on the eG remote agent host.
    - Next, copy the **netapp-manageability-sdk-<SDK_version>\netapp-manageability-sdk-<SDK_version>\lib\java\classes\manageontap.jar** file from the extracted contents to the **<EG_AGENT_INSTALL_DIR>\lib** directory (on a Windows host; on Unix, this will be the **/opt/egurkha/lib** directory). Sometimes, the name of the jar file may be suffixed by the NMSDK version number. For instance, instead of **manageontap.jar**, you might find **manageontap-5.2.jar** in **\java\classes**. In such a case, first, rename the jar file to **manageontap.jar**, and then copy the jar file to the **<EG_AGENT_INSTALL_DIR>\lib** directory.
    - Then, start the eG agent.
- To invoke the API commands, the eG agent has to be configured with the credentials of a NetApp user with the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**

For this purpose, you can either grant the aforesaid privileges to an existing user, or create a new user for this purpose. The new user creation process has been detailed in Section 1.1.3.1 below. Once such a user is created, make sure that you configure the eG tests with the credentials of such a user.

- Manage the storage device as a *NetApp Cluster* in eG using its **Cluster Management IP address**. Before that, make sure that the target is indeed a *NetApp Cluster* device and not a stand-alone *NetApp Unified Storage* device. For that, check the full version string for the Data ONTAP version in the NetApp device. If the version string contains the word "c-mode" or the word "cDOT", then it means that the target NeApp device is part of a cluster.

- Manage the storage device as a *NetApp Unified Storage* in eG using the device's IP address. Before that, make sure that the target is indeed a stand-alone *NetApp Unified Storage* device and not a device in a *NetApp Cluster.* For that, check the full version string for the Data ONTAP version in the NetApp device. If the version string contains "c-mode" or the word "cDOT", then it means that the target NetApp device is part of a cluster. If this word is not part of the version string, then it means that the target device is a stand-alone NetApp Unified Storage device.

1.1.4 Creating a New User with the Privileges Required for Monitoring the NetApp Unified Storage

As mentioned earlier, to run the API commands provided by the NMSDK and collect metrics, the eG agent requires the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**

To create a new user with the aforesaid privileges, do the following:

1. Login to the system hosting the remote agent.
2. Connect to the storage controller's console via SSH (say, using **puTTy.exe**).
3. Run the following command at the console to create a new role:

   useradmin role add <Name_of_new_role> -c "<A_brief_description_of_new_role>" -a <Comma-separated_list_of_privileges_to_be_granted_to_the_new_role>

   For instance, to create a role named **eG_role** with all the privileges required for monitoring NetApp Unified Storage, the command will be as follows:

   useradmin role add eG_Role -c "role for eG user" -a login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter*

4. Once the role is created successfully, proceed to create a new user group and assign the newly created role to it. The command for this will be:

useradmin group add <Name_of_new_group> -c "<A_brief_description_of_new_group>" -r <Name_of_new_role>

5. For instance, to create a group named **eG_Group** and to assign the **eG_Role** to it, the command will be as follows:

   useradmin group add eG_Group -c "Group for eG user" -r eG_Role

6. Then, create a new user and add that user to the newly created group. The command for the same is as follows:

   useradmin user add Mname_of_new_user> -c "<A_brief_description_of_new_user>" -g <Name_of_new_group>

   For instance, to create a user named **eG_User** and to add that user to the **eG_Group** that you created previously, the command will be as follows:

   useradmin user add eG_User -c "User for eG to monitor NetApp" -g eG_Group

7. This command, upon execution, will request for the password of the new user. The password is case-sensitive, and should be atleast 8 characters long. **It must contain atleast 2 alphabets and 1 digit**.

   ```
   New password:
   Reype new password:
   ```

   Then, confirm the new user's password by retyping it.

1.2 The Hardware Layer

The tests mapped to this layer report on the overall health of the hardware supporting the NetApp storage system.

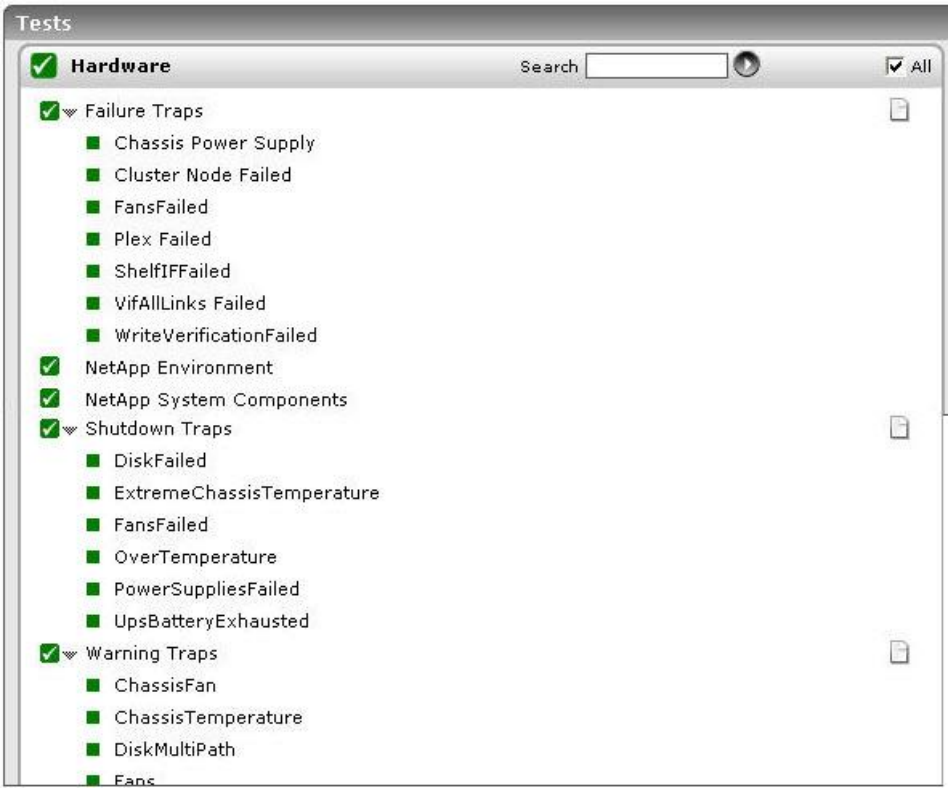

Figure 3: The tests mapped to the Hardware layer

1.2.1 Failure Traps Test

Hardware errors/failures, if not promptly detected and resolved, can prove to be fatal to the availability and overall health of a storage system. This test intercepts the traps sent by the storage system, extracts information related to hardware errors/failures from the traps, and reports the count and detailed description of these trap messages to the eG manager. This information enables administrators to detect current and potential hardware failures, understand the nature of these failures, and accordingly decide on the remedial measures.

Purpose	Intercepts the traps sent by the storage system, extracts information related to hardware errors/failures from the traps, and reports the count and detailed description of these trap messages to the eG manager
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test

1. **TESTPERIOD** - How often should the test be executed
2. **HOST -** The host for which the test is to be configured.
3. **SOURCEADDRESS** - Specify a comma-separated list of IP addresses or address patterns of the hosts from which traps are considered in this test. For example, *10.0.0.1,192.168.10.**. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.
4. **OIDVALUE** - Provide a comma-separated list of OID and value pairs returned by the traps. The values are to be expressed in the form, *DisplayName:OID-OIDValue*. For example, assume that the following OIDs are to be considered by this test: *.1.3.6.1.4.1.9156.1.1.2* and *.1.3.6.1.4.1.9156.1.1.3*. The values of these OIDs are as given hereunder:

OID	Value
.1.3.6.1.4.1.9156.1.1.2	Host_system
.1.3.6.1.4.1.9156.1.1.3	NETWORK

In this case the **OIDVALUE** parameter can be configured as *Trap1:.1.3.6.1.4.1.9156.1.1.2-Host_system,Trap2:.1.3.6.1.4.1.9156.1.1.3-Network*, where *Trap1* and *Trap2* are the display names that appear as descriptors of this test in the monitor interface.

An * can be used in the OID/value patterns to denote any number of leading or trailing characters (as the case may be). For example, to monitor all the OIDs that return values which begin with the letter 'F', set this parameter to *Failed:*-F**.

Typically, if a valid value is specified for an OID in the *OID-value* pair configured, then the test considers the configured OID for monitoring only when the actual value of the OID matches with its configured value. For instance, in the example above, if the value of OID *.1.3.6.1.4.1.9156.1.1.2* is found to be *HOST* and not *Host*_system, then the test ignores OID *.1.3.6.1.4.1.9156.1.1.2* while monitoring. In some cases however, an OID might not be associated with a separate value – instead, the OID itself might represent a value. While configuring such OIDs for monitoring, your **OIDVALUE** specification should be: *DisplayName:OID-any*. For instance, to ensure that the test monitors the OID *.1.3.6.1.4.1.9156.1.1.5*, which in itself, say represents a failure condition, then your specification would be:

Trap5: .1.3.6.1.4.1.9156.1.1.5-any.

<table>
<tr><td></td><td colspan="3">5. SHOWOID – Specifying true against SHOWOID will ensure that the detailed diagnosis of this test shows the OID strings along with their corresponding values. If you enter false, then the values alone will appear in the detailed diagnosis page, and not the OIDs.

6. TRAPOIDS – By default, this parameter is set to all, indicating that the eG agent considers all the traps received from the specified SOURCEADDRESSes. To make sure that the agent considers only specific traps received from the SOURCEADDRESS, then provide a comma-separated list of OIDs in the TRAPOIDS text box. A series of OID patterns can also be specified here, so that the test considers only those OIDs that match the specified pattern(s). For instance, *94.2*,*.1.3.6.1.4.25*, where * indicates leading and/or trailing spaces.

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

8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of failure event that occurred on the hardware of the target storage system</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Number of messages:
Indicates the number of events of this type that were captured during the last measurement period.</td><td>Number</td><td>The failure events may be generated due to the failure of hardware units like fans, chassis power supply etc., or failure of the cluster node, the shell interface module failure etc. If the failure events are not rectified within a certain pre-defined timeperiod, the storage system will be shutdown automatically.

Ideally, the value of this measure should be zero. A high value is an indication of performance degradation of the storage system.

The detailed diagnosis capability, if enabled provides you with a more detailed information about the failure events that were captured by this measure.</td></tr>
</table>

1.2.2 Shutdown Traps Test

This test provides administrators with a heads up on those failure events that have caused/could cause the storage system to come to a standstill!

Purpose	Provides administrators with a heads up on those failure events that have caused/could cause the storage system to come to a standstill
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test

1. **TESTPERIOD** - How often should the test be executed
2. **HOST -** The host for which the test is to be configured.
3. **SOURCEADDRESS** - Specify a comma-separated list of IP addresses or address patterns of the hosts from which traps are considered in this test. For example, *10.0.0.1,192.168.10.**. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.
4. **OIDVALUE** - Provide a comma-separated list of OID and value pairs returned by the traps. The values are to be expressed in the form, *DisplayName:OID-OIDValue*. For example, assume that the following OIDs are to be considered by this test: *.1.3.6.1.4.1.9156.1.1.2* and *.1.3.6.1.4.1.9156.1.1.3*. The values of these OIDs are as given hereunder:

OID	Value
.1.3.6.1.4.1.9156.1.1.2	Host_system
.1.3.6.1.4.1.9156.1.1.3	NETWORK

In this case the **OIDVALUE** parameter can be configured as *Trap1:.1.3.6.1.4.1.9156.1.1.2-Host_system,Trap2:.1.3.6.1.4.1.9156.1.1.3-Network*, where *Trap1* and *Trap2* are the display names that appear as descriptors of this test in the monitor interface.

An * can be used in the OID/value patterns to denote any number of leading or trailing characters (as the case may be). For example, to monitor all the OIDs that return values which begin with the letter 'F', set this parameter to *Failed:*-F**.

Typically, if a valid value is specified for an OID in the *OID-value* pair configured, then the test considers the configured OID for monitoring only when the actual value of the OID matches with its configured value. For instance, in the example above, if the value of OID *.1.3.6.1.4.1.9156.1.1.2* is found to be *HOST* and not *Host_*system, then the test ignores OID *.1.3.6.1.4.1.9156.1.1.2* while monitoring. In some cases however, an OID might not be associated with a separate value – instead, the OID itself might represent a value. While configuring such OIDs for monitoring, your **OIDVALUE** specification should be: *DisplayName:OID-any*. For instance, to ensure that the test monitors the OID *.1.3.6.1.4.1.9156.1.1.5*, which in itself, say represents a failure condition, then your specification would be:

Trap5: .1.3.6.1.4.1.9156.1.1.5-any.

<table>
<tr><td></td><td colspan="3">5. SHOWOID – Specifying true against SHOWOID will ensure that the detailed diagnosis of this test shows the OID strings along with their corresponding values. If you enter false, then the values alone will appear in the detailed diagnosis page, and not the OIDs.

6. TRAPOIDS – By default, this parameter is set to all, indicating that the eG agent considers all the traps received from the specified SOURCEADDRESSes. To make sure that the agent considers only specific traps received from the SOURCEADDRESS, then provide a comma-separated list of OIDs in the TRAPOIDS text box. A series of OID patterns can also be specified here, so that the test considers only those OIDs that match the specified pattern(s). For instance, *94.2*,*.1.3.6.1.4.25*, where * indicates leading and/or trailing spaces.

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

8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of failure event that occurred on the target storage system</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Number of messages:
Indicates the number of failure events that caused system shutdown during the last measurement period.</td><td>Number</td><td>The failure events may be generated due to the failure of hardware units like fans, chassis power supply etc, failure of the cluster node, the shell Interface module failure etc. When such failure events are generated, the storage system will shutdown automatically and the storage sytem must be restarted only after rectifying the failure.

Ideally, the value of this measure should be zero. A high value is an indication of performance degradation of the storage system.

The detailed diagnosis capability, if enabled provides you with a more detailed information about the failure events that were captured by this measure.</td></tr>
</table>

1.2.3 Warning Traps Test

SNMP traps carrying warning messages serve as early indicators of 'probable' failures/errors that can occur on the storage system. By intercepting and reading the warning traps sent by the storage system, this test proactively alerts administrators to potential issues in the performance of the storage system.

Purpose	Intercepts and reads the warning traps sent by the storage system, and proactively alerts administrators to potential issues in the performance of the storage system
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test

1. **TESTPERIOD** - How often should the test be executed
2. **HOST -** The host for which the test is to be configured.
3. **SOURCEADDRESS** - Specify a comma-separated list of IP addresses or address patterns of the hosts from which traps are considered in this test. For example, *10.0.0.1,192.168.10.**. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.
4. **OIDVALUE** - Provide a comma-separated list of OID and value pairs returned by the traps. The values are to be expressed in the form, *DisplayName:OID-OIDValue*. For example, assume that the following OIDs are to be considered by this test: *.1.3.6.1.4.1.9156.1.1.2* and *.1.3.6.1.4.1.9156.1.1.3*. The values of these OIDs are as given hereunder:

OID	Value
.1.3.6.1.4.1.9156.1.1.2	Host_system
.1.3.6.1.4.1.9156.1.1.3	NETWORK

In this case the **OIDVALUE** parameter can be configured as *Trap1:.1.3.6.1.4.1.9156.1.1.2-Host_system,Trap2:.1.3.6.1.4.1.9156.1.1.3-Network*, where *Trap1* and *Trap2* are the display names that appear as descriptors of this test in the monitor interface.

An * can be used in the OID/value patterns to denote any number of leading or trailing characters (as the case may be). For example, to monitor all the OIDs that return values which begin with the letter 'F', set this parameter to *Failed:*-F**.

Typically, if a valid value is specified for an OID in the *OID-value* pair configured, then the test considers the configured OID for monitoring only when the actual value of the OID matches with its configured value. For instance, in the example above, if the value of OID *.1.3.6.1.4.1.9156.1.1.2* is found to be *HOST* and not *Host_*system, then the test ignores OID *.1.3.6.1.4.1.9156.1.1.2* while monitoring. In some cases however, an OID might not be associated with a separate value – instead, the OID itself might represent a value. While configuring such OIDs for monitoring, your **OIDVALUE** specification should be: *DisplayName:OID-any*. For instance, to ensure that the test monitors the OID *.1.3.6.1.4.1.9156.1.1.5*, which in itself, say represents a failure condition, then your specification would be:

Trap5: .1.3.6.1.4.1.9156.1.1.5-any.

<table>
<tr><td></td><td colspan="3">5. SHOWOID – Specifying true against SHOWOID will ensure that the detailed diagnosis of this test shows the OID strings along with their corresponding values. If you enter false, then the values alone will appear in the detailed diagnosis page, and not the OIDs.
6. TRAPOIDS – By default, this parameter is set to all, indicating that the eG agent considers all the traps received from the specified SOURCEADDRESSes. To make sure that the agent considers only specific traps received from the SOURCEADDRESS, then provide a comma-separated list of OIDs in the TRAPOIDS text box. A series of OID patterns can also be specified here, so that the test considers only those OIDs that match the specified pattern(s). For instance, *94.2*,*.1.3.6.1.4.25*, where * indicates leading and/or trailing spaces.
7. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of warning event that occurred on the target storage system</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Number of messages:
Indicates the number of virus scanner related events that were captured during the last measurement period.</td><td>Number</td><td>The warning events may be generated due to the abnormal behavior of the fan/power supply etc, abnormal chassis temperature, an UPS drain, remote system warning, a directory whih is almost full, configuration errors etc. Such events are an indication for an administrator to take remedial steps to rectify the issue as soon as possible.
Ideally, the value of this measure should be zero. A high value is an indication of performance degradation of the storage system.
The detailed diagnosis capability, if enabled provides you with a more detailed information about the warning events that were captured by this measure.</td></tr>
</table>

1.2.4 NetApp System Components Test

This test periodically monitors the processors, spare disks, Vfilers, and the DMA channels used by the storage system, and proactively alerts you to abnormalities such as the following:

- Excessive CPU usage by the storage system;
- Over-utilization of processors supported by the storage system;
- Write latencies experienced by the NVRAM DMA transactions;
- Unavailability of spare disks;

Purpose	Periodically monitors the processors, spare disks, Vfilers, and the DMA channels used by the storage system, and proactively alerts you to abnormalities related to the same
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAMENAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAMENAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored.</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>CPU busy:
Indicates the percentage of time for which the CPU time was busy performing system-level processing.</td><td>Percent</td><td>A high value indicates that the storage system is utilizing CPU resources excessively. A consistent increase in this value could indicate a potential CPU contention on the storage system.</td></tr>
<tr><td>Avg processor busy:
Indicates what percentage of time, on an average, a processor is busy processing requests.</td><td>Percent</td><td>A high value indicates that processors have been over-utilized in more than one instance. This is a cause for concern, as it reveals load-balancing irregularities and the need for additional processors to handle the load.</td></tr>
</table>

	Total processor busy: Indicates the total percentage of time all the processors were actively serving requests.	Percent	A high value indicates that processors have been over-utilized in more than one instance. This is a cause for concern, as it reveals load-balancing irregularities and the need for additional processors to handle the load.
	NVRAM DMA write latency: Indicates the NVRAM DMA wait time per transaction in this storage system.	Milliseconds	When CP (consistency point) is triggered, Data ONTAP reads the journal of write requests from the NVRAM, and uses DMA (Direct Memory Access) to update the disk with the data. Direct memory access (DMA) is a feature that allows hardware subsystems to access system memory independently of the central processing unit (CPU). Any latencies experienced by the DMA channel can slowdown writes to the disk, consequently degrading the storage system's write performance. This is why, a low value is desired for this measure.
	NVRAM DMA transaction rate: Indicates the rate at which NVRAM DMA transactions were performed in this storage system.	Ops/sec	A consistent decrease in the value of this measure could indicate latencies. Any latencies experienced by the DMA channel can slowdown writes to the disk, consequently degrading the storage system's write performance.
	Are sufficient spare disks available? Indicates whether/not sufficient spare disks are available.		A hot spare disk is a disk that is assigned to a storage system but is not in use by a RAID group. It does not yet hold data but is ready for use. If a disk failure occurs within a RAID group, Data ONTAP automatically assigns hot spare disks to RAID groups to replace the failed disks. At a minimum, you should have at least one matching or appropriate hot spare available for each kind of disk installed in your storage system. However, having two available hot spares for all disks provides the best protection against disk failure. This measure indicates the value *Yes* if sufficient spare disks are available, and the value *No* if no spare disk are available in the storage system.

<table>
<tr>
<td></td>
<td></td>
<td></td>
<td>The numeric values that correspond to the above-mentioned measure values are as follows:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
By default, Data ONTAP issues warnings to the console and logs if you have fewer than one hot spare disk that matches the attributes of each disk in your storage system. You can change the threshold value for these warning messages by using the raid.min_spare_count option.
To make sure that you always have two hot spares for every disk (a best practice), you can set the raid.min_spare_count option to 2.
Setting the raid.min_spare_count option to 0 disables low spare warnings. You might want to do this if you do not have enough disks to provide hot spares (for example if your storage system does not support external disk shelves). You can disable the warnings only if the following requirements are met:
➢ Your system has 16 or fewer disks.
➢ You have no RAID groups that use RAID4.
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether sufficient spare disks are available in this storage system. However, in the graph of this measure, spare disk availability will be represented using the corresponding numeric equivalents i.e., 0 or 1.</td>
</tr>
<tr>
<td></td>
<td>Num offline/inconsistent vFiler resources:
Indicates the number of offline/inconsistent storage resources available across all vfilers in this storage system.</td>
<td>Number</td>
<td>MultiStore is also known as vFiler. A Unified Storage System's storage space could be divided into vFiler units. Each vFiler unit is run by a separate administrator, and is available on a separate network interface. One vFiler cannot view the storage space owned by other vFiler units (except for the special vFiler units "vFiler zero", which is the actual physical machine).</td>
</tr>
</table>

1.2.5 NetApp Environment Test

This test monitors the NetApp storage system's support environment - which includes its hardware, the fans, the power supply units, the battery, and the buffer cache - and promptly alerts you to current/potential issues in the health of this environment. These issues can range from abnormal hardware temperature to batteries fast-approaching their end-of-life to power rail failures and more!

Purpose	Monitors the NetApp storage system's support environment - which includes its hardware, the fans, the power supply units, the battery, and the buffer cache - and promptly alerts you to current/potential issues in the health of this environment
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the Cisco Router. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **SNMPPORT** - The port number through which the Cisco router exposes its SNMP MIB. The default value is 161. 5. **SNMPVERSION** – By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the **SNMPVERSION** list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list. 6. **SNMPCOMMUNITY** – The SNMP community name that the test uses to communicate with the Cisco router. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the **SNMPVERSION** chosen is **v3**, then this parameter will not appear. 7. **USERNAME** – This parameter appears only when **v3** is selected as the **SNMPVERSION**. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against the **USERNAME** parameter. 8. **AUTHPASS** – Specify the password that corresponds to the above-mentioned **USERNAME**. This parameter once again appears only if the **SNMPVERSION** selected is **v3**. 9. **CONFIRM PASSWORD** – Confirm the **AUTHPASS** by retyping it here. 10. **AUTHTYPE** – This parameter too appears only if **v3** is selected as the **SNMPVERSION**. From the **AUTHTYPE** list box, choose the authentication algorithm using which SNMP v3 converts the specified **USERNAME** and **PASSWORD** into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: ➢ **MD5** – Message Digest Algorithm ➢ **SHA** – Secure Hash Algorithm 11. **ENCRYPTFLAG** – This flag appears only when **v3** is selected as the **SNMPVERSION**. By default, the eG agent does not encrypt SNMP requests. Accordingly, the **ENCRYPTFLAG** is set to **NO** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **YES** option. 12. **ENCRYPTTYPE** – If the **ENCRYPTFLAG** is set to **YES**, then you will have to mention the encryption type by selecting an option from the **ENCRYPTTYPE** list. SNMP v3 supports the following encryption types: ➢ **DES** – Data Encryption Standard ➢ **AES** – Advanced Encryption Standard 13. **ENCRYPTPASSWORD** – Specify the encryption password here. 14. **CONFIRM PASSWORD** – Confirm the encryption password by retyping it here.

<table>
<tr><td></td><td colspan="3">15. TIMEOUT - Specify the duration (in seconds) within which the SNMP query executed by this test should time out in the TIMEOUT text box. The default is 10 seconds.

16. DD FREQUENCY - The DD FREQUENCY refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 2:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.

17. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG system 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 against DETAILED DIAGNOSIS. 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 bad and normal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Temperature status:

Indicates whether/not the hardware temperature is normal.</td><td></td><td>This measure reports the value Normal if the hardware is operating at a normal temperature and the value High if the hardware is operating at a temperature higher than the normal.

The values reported by this measure and their numeric equivalents are available in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Normal</td><td>1</td></tr><tr><td>High</td><td>2</td></tr></table>
Note:

This measure reports the Measure Values listed in the table above to indicate the current hardware temperature. However, in the graph of this measure, the temperature is indicated using only the Numeric Values listed in the above table.</td></tr>
</table>

<table>
<tr>
<td></td>
<td>Failed fans:
Indicates the number of main unit backplane fans that failed during the last measurement period.</td>
<td>Number</td>
<td>The detailed diagnosis capability, if enabled for this test, will list the fans that failed and the reason for their failure.</td>
</tr>
<tr>
<td></td>
<td>Failed power supplies:
Indicates the number of power supplies and the power rails that failed during the last measurement period.</td>
<td>Number</td>
<td>The detailed diagnosis capability, if enabled for this test, will list the power rails that have failed and the reason for their failure.</td>
</tr>
<tr>
<td></td>
<td>Battery status:
Indicates the current status of the NVRAM battery.</td>
<td></td>
<td>The values reported by this measure and their numeric equivalents are available in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Ok</td><td>1</td></tr>
<tr><td>Partially discharged</td><td>2</td></tr>
<tr><td>Fully discharged</td><td>3</td></tr>
<tr><td>Not present</td><td>4</td></tr>
<tr><td>Near EndOfLife</td><td>5</td></tr>
<tr><td>At EndOfLife</td><td>6</td></tr>
<tr><td>Unknown</td><td>7</td></tr>
<tr><td>Over charged</td><td>8</td></tr>
<tr><td>Fully charged</td><td>9</td></tr>
</table>
Note:
This measure reports the Measure Values listed in the table above to indicate the current battery status. However, in the graph of this measure, the same will be represented using only the Numeric Values listed in the above table.</td>
</tr>
</table>

	Cache age: Indicates the age of the oldest read only block in the buffer cache.	Hours	The value of this measure indicates how fast the read operations are cycling through the system memory. When the appliance is reading very large files (i.e., the files that are larger than the machine's memory size), buffer cache age will be very low.

1.3 The Network Layer

Use the tests mapped to this layer to determine whether the storage device is available over the network or not, and to identify speedy and bandwidth-intensive nework interfaces.

Figure 4: The tests mapped to the Network layer

Since these tests have been dealt with extensively in the *Monitoring Network Elements* document, let us proceed to the next layer.

1.4 The Physical Storage Layer

The tests associated with this layer reveal abnormalities related to the core physical storage components such as LUNs, disks, aggregates, and RAIDs.

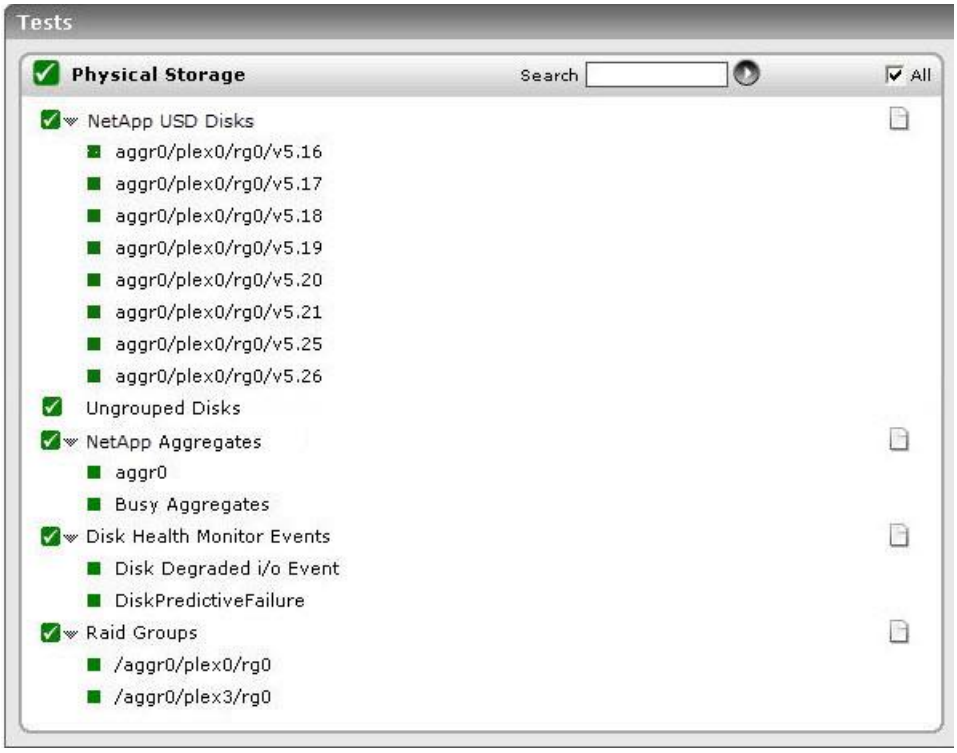

Figure 5: The tests mapped to the Physical Storage layer

1.4.1 NetApp Disks Test

Disks form the basic storage device in the NetApp storage systems. ATA disks, Fibre Channel disks, SCSI disks, SAS disks or SATA disks are used, depending on the storage system model.

Data ONTAP assigns and makes use of four different disk categories to support data storage, parity protection, and disk replacement. The disk category can be one of the following types: **Data disk** - Holds data stored on behalf of clients within RAID groups (and any system management data) **Global hot spare disk** - Does not hold usable data, but is available to be added to a RAID group in an aggregate. Any functioning disk that is not assigned to an aggregate functions acts as a hot spare disk. **Parity disk** - Stores information required for data reconstruction within RAID groups. **Double-parity disk** - Stores double-parity information within RAID groups, if RAIDDP is used.

Administrators should closely monitor the space usage and the level of I/O activity of each of these disks, so that they can proactively detect a space crunch or an I/O latency and receive early warnings of inconsistencies in load-balancing across disks. The **NetApp Unified Storage Disks** test aids administrators in this endeavor. This test auto-discovers the disks used by the storage system and reports how well every disk uses the available space and processes I/O requests. This way, potential space contentions and I/O latencies can be isolated, and slow disks and those that are running short of space can be identified. In addition, the test also reports the current state of each disk and how busy each disk is, thus pointing administrators to broken disks and over-used disks. In the process, the test turns the spotlight on irregularities in load-balancing.

Purpose	Auto-discovers the disks used by the storage system and reports how well every disk uses the available space and processes I/O requests
Target of the test	A NetApp Unified Storage
Agent	An external/remote agent

deploying the test	
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

	10. **VFILERNAMENAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAMENAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box. 11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds. 12. **DISK BUSY THRESHOLD** - A disk is termed as *Busy* if there is atleast one outstanding request that is awaiting a response. Alternately, you can set a threshold value in terms of percentage of time to classify the disk as a *Busy* disk. Specify such a threshold value in the **DISK BUSY THRESHOLD** text box. By default, this value is set to 70 (percent). **This parameter has been deprecated in v5.6.5 (and above).** 13. **READ LATENCY THRESHOLD** - Sometimes, the read operations by users on a disk may take too long to complete. In such a case, specify a threshold value in the **READ LATENCY THRESHOLD** text box, above which you can classify the disk as a Slow disk(read) i.e., you can term this disk as a slow disk (read) when the read operation by the user violates the threshold value mentioned in this text box. By default, this value is set to 20 (milliseconds). **This parameter has been deprecated in v5.6.5 (and above).** 14. **WRITE LATENCY THRESHOLD** - Sometimes, the write operations by users on a disk may take too long to complete. In such a case, specify a threshold value in the **WRITE LATENCY THRESHOLD** text box, above which you can classify the disk as a Slow disk(write) i.e., you can term this disk as a slow disk (write) when the write operation by the user violates the threshold value mentioned in this text box. By default, this value is set to 20 (milliseconds). **This parameter has been deprecated in v5.6.5 (and above).** 15. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**. 16. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.
Outputs of the test	One set of results for each disk on the NetApp storage system being monitored.

<table>
<tr><th>Measurements made by the test</th><th>Measurement</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td>Number of disks:
Indicates the total number of disks in this disk group.</td><td>Number</td><td>This measure is applicable only for disk groups and not individual disks. This measure has been deprecated in v5.6.5 (and above).</td></tr>
<tr><td></td><td>Raid state:
Indicates the curent RAID status of this disk in this Storage system.</td><td></td><td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>partner</td><td>1</td></tr><tr><td>Present</td><td>2</td></tr><tr><td>Zeroing</td><td>3</td></tr><tr><td>Spare</td><td>4</td></tr><tr><td>Copy</td><td>5</td></tr><tr><td>Pending</td><td>6</td></tr><tr><td>Reconstructing</td><td>7</td></tr><tr><td>Broken</td><td>8</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating the current RAID status of this disk in this Storage system. However, in the graph of this measure, status will be represented using the corresponding numeric equivalents i.e., 1 to 8.</td></tr>
<tr><td></td><td>Free space:
Indicates the amount of free space that is currently available for use in this disk of this Storage system.</td><td>MB</td><td>A high value is desired for this measure.</td></tr>
<tr><td></td><td>Physical space:
Indicates the total amount of space available in this disk of this Storage system.</td><td>MB</td><td></td></tr>
</table>

	Used space: Indicates the amount of space that is already utilized in this disk of this Storage system.	MB	A consistent increase in the value of these measures could indicate that the disk space is getting slowly but steadily eroded. Compare the value of these measures across all disks to identify the disks that are utilizing disk space excessively.
	Used space percentage: Indicates the percentage of space that has been already utilized in this disk.	Percent	
	Transfers: Indicates the rate at which data transfer is being initiated from this disk.	Ops/Sec	
	User reads: Indicates the rate at which data or metadata associated with user requests is being retrieved from this disk.	Ops/Sec	A consistent decrease in the value of this measure is indicative of a gradual slowdown in a user's ability to read from the disk. Compare the value of this measure across disks to know which disks service read requests slowly.
	User writes: Indicates the rate at which data or metadata associated with user requests is being stored in this disk.	Ops/Sec	A consistent decrease in the value of this measure is indicative of a gradual slowdown in a user's ability to write to a disk. Compare the value of this measure across disks to know which disks are servicing write requests slowly.
	User read latency: Indicates the time taken for retrieving data or metadata associated with user requests from this disk during the last measurement period.	Msecs	Very high values for these measures are indicative of the existence of road-blocks to rapid reading/writing by the Storage system. By observing the variations in these measures over time, you can understand whether the latencies are sporadic or consistent. Consistent delays in reading/writing could indicate that there are persistent bottlenecks (if any) in the Storage device to speedy I/O processing.
	User write latency: Indicates the time taken for a write operation on this disk during the last measurement period.	Msecs	

	Disk busy: Indicates the percentage of time when there is atleast one outstanding request (i.e., read or write) to this disk.	Percent	Comparing the percentage of time that the different disks are busy, an administrator can determine whether the application load is properly balanced across the different disks.

1.4.2 Ungrouped Disks

This test monitors the disks such as spare disks that do not belong to any RAID group in the NetApp Unified Storage system and reports the following:

- The number of disks that are currently zeroing
- The number of disks that are offline and the number of broken disks
- How well media scrubbing has been completed in those disks?

Purpose	Monitors the disks such as spare disks that do not belong to any RAID group in the NetApp Unified Storage system
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAMENAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAMENAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Zeroing disks:
Indicates the number of disks that are currently zeroing in this storage system.</td><td>Number</td><td>Disk zeroing is usually a time consuming background operation that is used to initialize the spare disks before they can be used.
Disk zeroing is the process of formatting the disk by filling zeroes i.e., overwriting the files with zeroes before being used.</td></tr>
<tr><td>Offline disks:
Indicates the number of disks that are currently offline in this storage system.</td><td>Number</td><td>Unresponsive or semi-responsive disks are taken offline by the operating system and its data is reconstructed from the associated parity disks. This puts a strain on the performance of the associated RAID group. Irrecoverable offline disks will be failed.</td></tr>
</table>

	Broken disks: Indicates the number of disks whose RAID status is *Broken* in this storage system.	Number	The disks may be broken due to disk failure, labeling issues or intentional setting to physical removal. Broken disks affect constituent raid group performance and put the system at risk of losing data if spares are unavailable.
	Average media scrub percentage: Indicates the average percentage of media scrubbing that is currently completed across all spare disks in this storage system.	Percent	Media scrubbing is a continuous background process. The purpose of the continuous media scrub is to detect and correct media errors in order to minimize the chance of storage system disruption due to a media error while a storage system is in degraded or reconstruction mode. By default, Data ONTAP runs continuous background media scrubbing for media errors on all storage system disks. If a media error is found, Data ONTAP uses RAID to reconstruct the data and repairs the error. Due to media scrubbing process, the disk LEDs may blink on an apparently idle storage system and some CPU activity may occur even when no user workload is present.

1.4.3 NetApp Aggregates Test

To support the differing security, backup, performance, and data sharing needs of your users, you group the physical data storage resources on your storage system into one or more aggregates. These aggregates provide storage to the volume or volumes that they contain. Each aggregate has its own RAID configuration, plex structure, and set of assigned disks or array LUNs.

Periodically, you must monitor the state, I/O activity, processing power, and space usage of each of the aggregates configured on your storage system, so that probable space contentions and I/O overloads can be rapidly detected, and failed/inconsistent/busy aggregates can be easily identified. Also, to be able to accurately point to failed checksum storage, problematic RAID groups, or issues in plex resynchronization in an aggregate, the key components of each aggregate - such as, RAID groups, plex structures and checksum disks - should also be monitored from time to time. The **NetApp Aggregates** test provides all these performance insights. This test auto-discovers the aggregates configured on a storage system, and periodically reports the following:

- What is the current state of each aggregate?
- Which are the busy aggregates?
- Is any aggregate running short of storage space?
- Is I/O load uniformly distributed across all aggregates, or is any aggregate overloaded with read-write requests?
- What is the current status of the checksum storage in each aggregate?
- What is the current status of the plex structures in each aggregate?
- Are the RAID groups in an aggregate in a normal state?
- Did any aggregate experience issues during plex resynchronization?

Purpose	Auto-discovers the aggregates configured on a storage system, and periodically monitors the state, I/O activity, processing power, and space usage of each of the aggregates, so that probable space contentions and I/O overloads can be rapidly detected, and failed/inconsistent/busy aggregates can be easily identified. Also, to be able to accurately point to failed checksum storage, problematic RAID groups, or issues in plex synchronization in an aggregate, the key components of each aggregate - such as, RAID groups, plex structures and checksum disks - are also monitored from time to time
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">
10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAMENAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. TRANSFERS THRESHOLD - You can set a threshold value for the rate at which the transfers are serviced by an aggregate. Specifying such a value in the TRANSFERS THRESHOLD text box implies that the aggregates violating this threshold value will be termed as Busy aggregates. The default value is 15 (Transfers/Sec). This parameter is deprecated in v5.6.5 (and above).
13. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
14. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each aggregate on the NetApp storage system being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>NetApp aggregates:
Indicates the number of busy aggregates in the storage system.</td><td>Number</td><td>This measure is applicable only to the Busy Aggregates descriptor.
The detailed diagnosis capability of this measure, if enabled, lists out the name of the aggregate and the Transfer rate of each aggregate i.e., the rate at which data transfer is serviced by an aggregate.
This measure is deprecated in v5.6.5 (and above).</td></tr>
</table>

<table>
<tr>
<td></td>
<td>State:
Indicates the current state of this aggregate.</td>
<td></td>
<td>The values that this measure can report and their corresponding numeric values have been listed in the table below. A brief description for each Measure Value is also provided:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr>
<tr><td>Creating</td><td>1</td><td></td></tr>
<tr><td>Online</td><td>2</td><td>Read and write access to volumes hosted on this aggregate is allowed.</td></tr>
<tr><td>Restricted</td><td>3</td><td>Some operations, such as parity reconstruction, are allowed, but data access is not allowed.</td></tr>
<tr><td>Iron Restricted</td><td>4</td><td>A WAFL consistency check is being performed on the aggregate.</td></tr>
<tr><td>Partial</td><td>5</td><td>At least one disk was found for the aggregate, but two or moredisks are missing.</td></tr>
<tr><td>Offline</td><td>6</td><td>No access to the aggregate is allowed.</td></tr>
<tr><td>Failed</td><td>7</td><td></td></tr>
<tr><td>Unknown</td><td>8</td><td></td></tr>
</table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating the current status of an aggregate. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., 1 to 8.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Is aggregate inconsistent?
Indicates whether/not this aggregate is inconsistent.</td>
<td></td>
<td>One of the reasons why an aggregate is marked as inconsistent or corrupted, is when the Lost write protection feature detects an issue. Lost write protection is a feature of Data ONTAP that occurs on each WAFL read. Data is checked against block checksum information (WAFL context) and RAID parity data. If an issue is detected, there are two possible outcomes:
d. The drive containing the data is failed.
e. The aggregate containing the data is marked inconsistent.
If an aggregate is marked inconsistent, it will require the use of WAFL iron to be able to return the aggregate to a consistent state.
This measure indicates a value of Yes if the aggregate is inconsistent and the value No if the aggregate is not inconsistent. The numeric values that correspond to the above-mentioned values are detailed in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Yes</td><td>1</td></tr>
<tr><td>No</td><td>2</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether/not this aggregate is inconsistent. However, in the graph of this measure, the inconsistent state of an aggregate will be represented using the corresponding numeric equivalents i.e., 1 or 2.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Mirror status:
Indicates the current mirror status of this aggregate.</td>
<td></td>
<td>The values that this measure can report and their corresponding numeric values have been listed in the table below. A brief description for a few Measure Values is also provided:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr>
<tr><td>Unmirrored</td><td>1</td><td>The aggregate is not mirrored. Unmirrored aggregates have only one plex (copy of their data), which contains all of the RAID groups belonging to that aggregate.</td></tr>
<tr><td>Mirrored</td><td>2</td><td>The aggregate is mirrored. Mirrored aggregates have two plexes (copies of their data), which use the SyncMirror functionality to duplicate the data to provide redundancy</td></tr>
<tr><td>Mirror Resynchronizing</td><td>3</td><td>One of the mirrored aggregate's plexes is being resynchronized</td></tr>
<tr><td>Un Initialized</td><td>4</td><td></td></tr>
<tr><td>CP Count Check In Progress</td><td>5</td><td>WAFL consistency check is in progress</td></tr>
</table>
</td>
</tr>
</table>

Needs CP Count Check	6	WAFL consistency check needs to be performed on the aggregate
Mirror Degraded	7	The aggregate is mirrored and one of its plexes is offline or resynchronizing
Invalid	8	The aggregate contains no volumes and none can be added. Typically this happens only after an aborted `aggr copy` operation.
Failed	9	
Limbo	10	

Note:

By default, this measure reports the above-mentioned **Measure Value**s while indicating the current mirror status of this aggregate in this storage system. However, in the graph of this measure, the mirror status will be represented using the corresponding numeric equivalents - i.e., *1 to 10*.

<table>
<tr>
<td></td>
<td>Is Raid state abnormal?
Indicates whether/not the RAID of this aggregate is in an abnormal state currently.</td>
<td></td>
<td>This measure indicates a value of Yes if the RAID of this aggregate is in an abnormal state and the value No if the RAID of this aggregate is normal. The numeric values that correspond to the above-mentioned values are detailed in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Yes</td><td>1</td></tr>
<tr><td>No</td><td>2</td></tr>
</table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether the RAID of this aggregate is in an abnormal state. However, in the graph of this measure, the RAID states will be represented using the corresponding numeric equivalents i.e., 1 or 2.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Checksum status:
Indicates the current checksum status of this aggregate.</td>
<td></td>
<td>The values that this measure can report and their corresponding numeric values have been listed in the table below.
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Active</td><td>1</td></tr>
<tr><td>Off</td><td>2</td></tr>
<tr><td>Reverting</td><td>3</td></tr>
<tr><td>None</td><td>4</td></tr>
<tr><td>Unknown</td><td>5</td></tr>
<tr><td>Initializing</td><td>6</td></tr>
<tr><td>Reinitializing</td><td>7</td></tr>
<tr><td>Reinitialized</td><td>8</td></tr>
<tr><td>Upgrading Phase1</td><td>9</td></tr>
<tr><td>Upgrading Phase2</td><td>10</td></tr>
</table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating the current checksum status of this aggregate. However, the graph of this measure will be represented using the corresponding numeric equivalents i.e., 1 to 10.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Are plexes offline?:
Indicates whether/not the plexes in this aggregate are currently offline.</td>
<td></td>
<td>A plex is a collection of one or more RAID groups that together provide the storage for one or more WAFL® file system volumes. Data ONTAP uses plexes as the unit of RAID-level mirroring when the SyncMirror® feature is enabled. All RAID groups in one plex are of the same level, but may have a different number of disks.

This measure reports the value Yes if the plexes in this aggregate are currently offline and the value No if the plexes are not offline. The numeric values that correspond to the above-mentioned values are detailed in the table below:

<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>No</td><td>1</td></tr>
<tr><td>Yes</td><td>2</td></tr>
</table>

Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether the plexes in this aggregate are currently offline or not. However, in the graph of this measure, the state of the plexes will be represented using the corresponding numeric equivalents i.e., 1 or 2.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Are plexes resyncing?
Indicates whether/not the plexes of this aggregate are currently being resynchronized.</td>
<td></td>
<td>Plex resynchronization is a process that ensures two plexes of a mirrored aggregate have exactly the same data. When plexes are unsynchronized, one plex contains data that is more up to date than that of the other plex. Plex resynchronization updates the out-of-date plex so that both plexes are identical.

Data ONTAP resynchronizes the two plexes of a mirrored aggregate if one of the following situations occurs:

f. One of the plexes was taken offline and then brought online later.

g. You add a plex to an unmirrored aggregate.

This measure reports the value Yes if the plexes in this aggregate are currently resyncing and the value No if the plexes are not resyncing. The numeric values that correspond to the above-mentioned values are detailed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>No</td><td>1</td></tr><tr><td>Yes</td><td>2</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether the plexes in this aggregate are currently offline or not. However, in the graph of this measure, the state of the plexes will be represented using the corresponding numeric equivalents i.e., 1 or 2.</td>
</tr>
<tr>
<td></td>
<td>Total size:
Indicates the total usable size of this aggregate.</td>
<td>MB</td>
<td>The size of this aggregate excludes the WAFL reserve and the aggregate snapshot reserve. This measure will report a value of 0 if the aggregate is restricted or offline.</td>
</tr>
<tr>
<td></td>
<td>Aggregate used size:
Indicates the amount of space that is currently used in this aggregate.</td>
<td>MB</td>
<td>This measure will report a value 0 if the aggregate is not usable i.e., offline.</td>
</tr>
</table>

	Percentage size used: Indicates the percentage of space that is currently used in this aggregate.	Percent	A value close to 100% is an indication of space constraint in the aggregate.
	Total files: Indicates the total number of files in this aggregate.	Number	
	Used files: Indicates the total number of files that are currently stored in this aggregate.	Number	
	Transfers: Indicates the rate at which the transfers are serviced by this aggregate.	Ops/Sec	Compare the value of this measure across aggregates to identify the busy aggregates.
	User reads: Indicates the rate at which the read request from the user is serviced by this aggregate.	Ops/Sec	A consistent decrease in the value of this measure could indicate a bottleneck when processing read requests. Compare the value of this measure across aggregates to know which aggregates service read requests slowly.
	User writes: Indicates the rate at which the write request from the user is serviced in this aggregate.	Ops/Sec	A consistent decrease in the value of this measure could indicate a bottleneck when processing write requests. Compare the value of this measure across aggregates to know which aggregates are servicing write requests slowly.
	CP reads: Indicates the rate at which the read request from the user is serviced during a Consistency Point (CP) operation in this aggregate.	Ops/Sec	A consistent decrease in the value of this measure could indicate that CP operations are slowing down the processing of read requests.
	Block read rate: Indicates the rate at which the blocks are read from this aggregate upon a user request.	Ops/Sec	A consistent decrease in the value of this measure could indicate a bottleneck when processing read requests. Compare the value of this measure across aggregates to know which aggregates service block read requests slowly.

	Block write rate: Indicates the rate at which the blocks are written to this aggregate upon a user request.	Ops/Sec	A consistent decrease in the value of this measure could indicate a bottleneck when processing write requests. Compare the value of this measure across aggregates to know which aggregates are servicing block write requests slowly.
	Block read rate during CP: Indicates the rate at which the blocks are read from this aggregate during a Consistency point (CP) operation.	Ops/Sec	A consistent decrease in the value of this measure could indicate that CP operations are slowing down the processing of read requests.

1.4.4 Raid Groups Test

Data ONTAP organizes disks into RAID groups, which are collections of data and parity disks to provide parity protection. For Data ONTAP 6.5 onwards the following RAID types are supported for NetApp storage systems:

- **RAID4 technology**: In this RAID, within each RAID group, a single disk is assigned for holding parity data, which ensures against data loss due to a single disk failure within a group.
- **RAID-DP™ technology (DP for double-parity)**: RAID-DP provides a higher level of RAID protection for Data ONTAP aggregates. Within its RAID groups, it allots one disk for holding parity data and one disk for holding double-parity data. Double-parity protection ensures against data loss due to a double disk failure within a group.

For native storage, Data ONTAP uses RAID-DP or RAID4 groups to provide parity protection. For third-party storage, Data ONTAP uses RAID0 groups to optimize performance and storage utilization. The storage arrays provide the parity protection for third-party storage. Data ONTAP RAID groups are organized into plexes, and plexes are organized into aggregates.

This test auto discovers the RAID groups in the storage system and helps the administrator figure out the following:

- How many disks are in abnormal state i.e., prefailed and replacing?
- What is the total size of this RAID group? Is any RAID group facing/is about to encounter a space crunch?
- The percentage of media scrubbing andd parity scrubbing that has been completed in this RAID group.

Purpose	Auto discovers the RAID groups in the storage system and helps the administrator figure out the following: • How many disks are in abnormal state i.e., prefailed and replacing? • What is the total size of this RAID group? Is any RAID group facing/is about to encounter a space crunch? • The percentage of media scrubbing andd parity scrubbing that has been completed in this RAID group.
Target of the	A NetApp Unified Storage

test	
Agent deploying the test	An external/remote agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">
10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAMENAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each aggregate on the NetApp storage system being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Prefailed disks:
Indicates the number of prefailed disks in this RAID group.</td><td>Number</td><td>The disks that are manually failed due to excessive error logging are termed as Prefailed disks. The contents of these disks are copied into suitable replacement disks i.e., the spare disks available in the storage system.
Ideally, the value of this measure should be 0.</td></tr>
</table>

	Replacing disks: Indicates the number of replacing disks in this RAID group.	Number	Mismatched disks that are part of an aggregate can be replaced with a more suitable spare disk without disrupting the data service. This process uses the Rapid RAID Recovery process to copy the data from the disk being replaced to a specified spare disk. Frequently replacing the disks will lead to the system degradation. Therefore, the frequent replacement of the disks needs to be avoided by proper initial configuration.
	Total physical space: Indicates the total size of this RAID group.	MB	
	Used space: Indicates the total amount of space used by all disks in this RAID group.	MB	Ideally, the value of this measure should be low. If this value grows close to that of the *Total physical space* measure, then you may want to consider adding more disks to the storage system, or free space in the disks by deleting unnecessary data.
	Used space percentage: Indicates the percent of space that is utilized across all disks in this RAID group.	Percent	A low value is desired for this measure. A value close to 100% indicates excessive disk space usage by a RAID group.
	Media scrub percentage: Indicates the percentage of media scrubbing that is currently completed in this RAID group.	Percent	Media scrubbing is a continuous background process. The purpose of the continuous media scrub is to detect and correct media errors in order to minimize the chance of storage system disruption due to a media error while a storage system is in degraded or reconstruction mode. By default, Data ONTAP runs continuous background media scrubbing for media errors on all storage system disks. If a media error is found, Data ONTAP uses RAID to reconstruct the data and repairs the error. Due to media scrubbing process, the disk LEDs may blink on an apparently idle storage system and some CPU activity may occur even when no user workload is present.
	Parity scrub percentage: Indicates the percentage of parity scrubbing that is currently completed in this RAID group.	Percent	The purpose of the parity scrub is to detect and correct errors in the parity disk of the RAID group. A consistent parity is required for disk reconstruction.

1.4.5 Disk Health Monitor Events

This test reports the number and nature of error events that occurred on the disks of this storage system. This way, administrators can promptly detect disk failures/related performance issues.

Purpose	Reports the number and nature of error events that occurred on the disks of this storage system
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST -** The host for which the test is to be configured.
3. **SOURCEADDRESS** - Specify a comma-separated list of IP addresses or address patterns of the hosts from which traps are considered in this test. For example, *10.0.0.1,192.168.10.**. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.
4. **OIDVALUE** - Provide a comma-separated list of OID and value pairs returned by the traps. The values are to be expressed in the form, *DisplayName:OID-OIDValue*. For example, assume that the following OIDs are to be considered by this test: *.1.3.6.1.4.1.9156.1.1.2* and *.1.3.6.1.4.1.9156.1.1.3*. The values of these OIDs are as given hereunder:

OID	Value
.1.3.6.1.4.1.9156.1.1.2	Host_system
.1.3.6.1.4.1.9156.1.1.3	NETWORK

In this case the **OIDVALUE** parameter can be configured as *Trap1:.1.3.6.1.4.1.9156.1.1.2-Host_system,Trap2:.1.3.6.1.4.1.9156.1.1.3-Network*, where *Trap1* and *Trap2* are the display names that appear as descriptors of this test in the monitor interface.

An * can be used in the OID/value patterns to denote any number of leading or trailing characters (as the case may be). For example, to monitor all the OIDs that return values which begin with the letter 'F', set this parameter to *Failed:*-F**.

Typically, if a valid value is specified for an OID in the *OID-value* pair configured, then the test considers the configured OID for monitoring only when the actual value of the OID matches with its configured value. For instance, in the example above, if the value of OID *.1.3.6.1.4.1.9156.1.1.2* is found to be *HOST* and not *Host_*system, then the test ignores OID *.1.3.6.1.4.1.9156.1.1.2* while monitoring. In some cases however, an OID might not be associated with a separate value – instead, the OID itself might represent a value. While configuring such OIDs for monitoring, your **OIDVALUE** specification should be: *DisplayName:OID-any*. For instance, to ensure that the test monitors the OID *.1.3.6.1.4.1.9156.1.1.5*, which in itself, say represents a failure condition, then your specification would be:

Trap5: .1.3.6.1.4.1.9156.1.1.5-any.

<table>
<tr><td></td><td colspan="3">5. SHOWOID – Specifying true against SHOWOID will ensure that the detailed diagnosis of this test shows the OID strings along with their corresponding values. If you enter false, then the values alone will appear in the detailed diagnosis page, and not the OIDs.

6. TRAPOIDS – By default, this parameter is set to all, indicating that the eG agent considers all the traps received from the specified SOURCEADDRESSes. To make sure that the agent considers only specific traps received from the SOURCEADDRESS, then provide a comma-separated list of OIDs in the TRAPOIDS text box. A series of OID patterns can also be specified here, so that the test considers only those OIDs that match the specified pattern(s). For instance, *94.2*,*.1.3.6.1.4.25*, where * indicates leading and/or trailing spaces.

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

8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of error event that occurred on the disks of the target storage system</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Number of messages:
Indicates the number of events of this type that were captured during the last measurement period.</td><td>Number</td><td>The event type may either be predictive failure of the disk or the degraded I/O event of the disk. When such types of events are generated, the Operating system automatically recovers the maximum amount of data from the affected disks and stores in a spare disk.

Ideally, the value of this measure should be zero. A high value is an indication of performance degradation of the disk.

The detailed diagnosis capability, if enabled provides you with a more detailed information about the events that were captured by this measure.</td></tr>
</table>

1.5 The NetApp OS Layer

With the help of the tests mapped to this layer, you can monitor Consistency Points (CP) and NetApp's Write Anywhere File Layout (WAFL).

Figure 6: The tests mapped to the NetApp OS layer

1.5.1 Consistency points Test

Consistency points (CP) are periodic tasks performed by Data ONTAP wherein unwritten data that is temporarily stored in the non- volatile RAM (NVRAM), is copied over (committed) to the disks thereby maintaining system consistency.

Typically, a CP occurs when the NVRAM journal is half full or when 10 seconds have passed since the most recent CP, whichever comes first. By carefully studying the number and frequency of CPs, you can accurately determine the level of write activity on your storage system. This test serves as a good indicator of the write request load on your storage system, as it reports the number of CPs that occurred and when it occurred.

Purpose	Reports the number of CPs that occurred and when it occurred
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the device. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **SNMPPORT** - The port number through which the device exposes its SNMP MIB. The default value is 161. 5. **SNMPVERSION** – By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the **SNMPVERSION** list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list. 6. **SNMPCOMMUNITY** – The SNMP community name that the test uses to communicate with the Cisco router. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the **SNMPVERSION** chosen is **v3**, then this parameter will not appear. 7. **USERNAME** – This parameter appears only when **v3** is selected as the **SNMPVERSION**. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against the **USERNAME** parameter. 8. **AUTHPASS** – Specify the password that corresponds to the above-mentioned **USERNAME**. This parameter once again appears only if the **SNMPVERSION** selected is **v3**. 9. **CONFIRM PASSWORD** – Confirm the **AUTHPASS** by retyping it here. 10. **AUTHTYPE** – This parameter too appears only if **v3** is selected as the **SNMPVERSION**. From the **AUTHTYPE** list box, choose the authentication algorithm using which SNMP v3 converts the specified **USERNAME** and **PASSWORD** into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: ➢ **MD5** – Message Digest Algorithm ➢ **SHA** – Secure Hash Algorithm 11. **ENCRYPTFLAG** – This flag appears only when **v3** is selected as the **SNMPVERSION**. By default, the eG agent does not encrypt SNMP requests. Accordingly, the **ENCRYPTFLAG** is set to **NO** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **YES** option. 12. **ENCRYPTTYPE** – If the **ENCRYPTFLAG** is set to **YES**, then you will have to mention the encryption type by selecting an option from the **ENCRYPTTYPE** list. SNMP v3 supports the following encryption types: ➢ **DES** – Data Encryption Standard ➢ **AES** – Advanced Encryption Standard 13. **ENCRYPTPASSWORD** – Specify the encryption password here. 14. **CONFIRM PASSWORD** – Confirm the encryption password by retyping it here.

	15. **TIMEOUT** - Specify the duration (in seconds) within which the SNMP query executed by this test should time out in the **TIMEOUT** text box. The default is 10 seconds. 16. **DATA OVER TCP –**By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the equalizer over TCP (and not UDP). For this, set the **DATA OVER TCP** flag to **Yes**. By default, this flag is set to **No**.		
Outputs of the test	One set of results for the NetApp storage system being monitored.		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	CP due to log-full operations: Indicates the number of consistency point operations that occurred due to the cache i.e., the NVRAM log being full during the last measurement period.	Number	The storage system automatically triggers a consistency point when the NVRAM log is 50% full and writes the data available in the NVRAM log to the disk. By doing so, the write latency of the disk is maintained along with a smooth transition of data to the disk from the NVRAM log.
	Back to back CP operations: Indicates the number of back to back consistency point operations that occurred during the last measurement period.	Number	The back to back consistency point operations indicate that the storage system is highly loaded and the write rate on the disk is more than the consistency point rate.
	Number of CP operations: Indicates the total number of consistency point operations that occurred during the last measurement period.	Number	This is a good indicator of the level of write activity on the storage system.

1.5.2 NetApp File Layouts (WAFL) Test

WAFL is the NetApp® Write Anywhere File Layout, which defines how NetApp lays out data on disk. The WAFL buffer cache is a read cache maintained by WAFL in system memory. On a storage system, if you attempt to read data that is not in the WAFL buffer cache, it results in a direct disk read. Disk reads are expensive operations that increase the processing overheads of your storage system. A well-tuned, right-sized buffer cache can alone help in keeping disk reads minimal. By closely tracking the requests to the storage system and how the WAFL buffer cache services these requests, the **NetApp File Layouts** test points you to the ineffective usage (if any) of the buffer cache, which can be

attributed to insufficient cache memory. Based on the findings of this test, you can then proceed to increase the cache memory (if required).

Purpose	By closely tracking the requests to the storage system and how the WAFL buffer cache services these requests, the **NetApp File Layouts** test points you to the ineffective usage (if any) of the buffer cache, which can be attributed to insufficient cache memory
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

	10. **VFILERNAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAMENAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box. 11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.		
Outputs of the test	One set of results the NetApp storage system being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Name cache hits:** Indicates the rate at which the name cache buffer was successfully queried for an entry during the last measurement period.	Hits/Sec	While a high value is desired for the *Name cache hits* measure, a low value is ideal for *Name cache misses*. A large number of cache misses and very few cache hits indicate that adequate entries are not available in the cache to service requests to the storage system. This in turn will force direct disk reads, thereby increasing the processing overheads of the storage system. To minimize disk reads and maximize cache reads, you can increase the WAFL cache memory using WAFL extended cache and the Performance Acceleration Module (PAM) family. WAFL extended cache is a software component of Data ONTAP and requires a license. WAFL extended cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. Once the extended cache is enabled, you can cache the following in it: h. **Cache normal user data blocks**: If you cache normal user data blocks, the WAFL extended cache interprets this setting as the buffer cache policy of *keep* and saves normal user data blocks in the extended cache. i. **Caching low-priority user data blocks**: You can cache low-priority user data blocks that are not normally stored by WAFL extended cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through a network-attached storage (NAS) protocol such as Network File System (NFS). Caching low-priority user data blocks is useful if you have workloads that fit within WAFL extended cache memory and if the workloads consist of either write followed by read or large sequential reads.

	Name cache misses: Indicates the rate at which the user query for an entry failed in the name cache buffer during the last measurement period.	Misses/Sec	j. **Caching only system metadata:** If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL extended cache memory by turning off both normal user data block caching and low-priority user data block caching. k. **Integrating FlexShare buffer cache policies with WAFL extended cache:** For additional cache control, you can integrate FlexShare buffer cache policies with the WAFL extended cache options. Doing so allows you to set caching policies on specific volumes. You can choose to enable only the FlexShare buffer cache policies without enabling all other FlexShare options.

	Directory find hits: Indicates the rate at which the user request successfully found a directory using the WAFL buffer during the last measurement period.	Hits/Sec	A large number of cache misses and very few cache hits indicate that adequate entries are not available in the cache to service requests to the storage system. This in turn will force direct disk reads, thereby increasing the processing overheads of the storage system. To minimize disk reads and maximize cache reads, you can increase the WAFL cache memory using WAFL extended cache and the Performance Acceleration Module (PAM) family. WAFL extended cache is a software component of Data ONTAP and requires a license. WAFL extended cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. Once the extended cache is enabled, you can cache the following in it: • **Cache normal user data blocks**: If you cache normal user data blocks, the WAFL extended cache interprets this setting as the buffer cache policy of *keep* and saves normal user data blocks in the extended cache. • **Caching low-priority user data blocks**: You can cache low-priority user data blocks that are not normally stored by WAFL extended cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through a network-attached storage (NAS) protocol such as Network File System (NFS). Caching low-priority user data blocks is useful if you have workloads that fit within WAFL extended cache memory and if the workloads consist of either write followed by read or large sequential reads.

	Directory find misses: Indicates the rate at which the user request failed to find a directory using the WAFL buffer during the last measurement period.	Misses/Sec	• **Caching only system metadata:** If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL extended cache memory by turning off both normal user data block caching and low-priority user data block caching. • **Integrating FlexShare buffer cache policies with WAFL extended cache:** For additional cache control, you can integrate FlexShare buffer cache policies with the WAFL extended cache options. Doing so allows you to set caching policies on specific volumes. You can choose to enable only the FlexShare buffer cache policies without enabling all other FlexShare options.

	Buffer hash hits: Indicates the rate at which the hash queue of the WAFL buffer was successfully queried for an entry during the last measurement period.	Hits/Sec	A large number of cache misses and very few cache hits indicate that adequate entries are not available in the cache to service requests to the storage system. This in turn will force direct disk reads, thereby increasing the processing overheads of the storage system. To minimize disk reads and maximize cache reads, you can increase the WAFL cache memory using WAFL extended cache and the Performance Acceleration Module (PAM) family. WAFL extended cache is a software component of Data ONTAP and requires a license. WAFL extended cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. Once the extended cache is enabled, you can cache the following in it: • **Cache normal user data blocks**: If you cache normal user data blocks, the WAFL extended cache interprets this setting as the buffer cache policy of *keep* and saves normal user data blocks in the extended cache. • **Caching low-priority user data blocks**: You can cache low-priority user data blocks that are not normally stored by WAFL extended cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through a network-attached storage (NAS) protocol such as Network File System (NFS). Caching low-priority user data blocks is useful if you have workloads that fit within WAFL extended cache memory and if the workloads consist of either write followed by read or large sequential reads.

	Buffer hash misses: Indicates the rate at which the user request failed to find an entry in the hash queue of the WAFL buffer during the last measurement period.	Misses/sec	• **Caching only system metadata:** If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL extended cache memory by turning off both normal user data block caching and low-priority user data block caching. • **Integrating FlexShare buffer cache policies with WAFL extended cache:** For additional cache control, you can integrate FlexShare buffer cache policies with the WAFL extended cache options. Doing so allows you to set caching policies on specific volumes. You can choose to enable only the FlexShare buffer cache policies without enabling all other FlexShare options.

	Inode cache hits: Indicates the rate at which the inode information of a file was successfully found using the WAFL buffer during the last measurement period.	Hits/Sec	A large number of cache misses and very few cache hits indicate that adequate entries are not available in the cache to service requests to the storage system. This in turn will force direct disk reads, thereby increasing the processing overheads of the storage system. To minimize disk reads and maximize cache reads, you can increase the WAFL cache memory using WAFL extended cache and the Performance Acceleration Module (PAM) family. WAFL extended cache is a software component of Data ONTAP and requires a license. WAFL extended cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. Once the extended cache is enabled, you can cache the following in it: • **Cache normal user data blocks**: If you cache normal user data blocks, the WAFL extended cache interprets this setting as the buffer cache policy of *keep* and saves normal user data blocks in the extended cache. • **Caching low-priority user data blocks**: You can cache low-priority user data blocks that are not normally stored by WAFL extended cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through a network-attached storage (NAS) protocol such as Network File System (NFS). Caching low-priority user data blocks is useful if you have workloads that fit within WAFL extended cache memory and if the workloads consist of either write followed by read or large sequential reads.

	Inode cache misses: Indicates the rate at which the inode information of a file was not found in the WAFL buffer during the last measurement period.	Misses/sec	• **Caching only system metadata:** If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL extended cache memory by turning off both normal user data block caching and low-priority user data block caching. • **Integrating FlexShare buffer cache policies with WAFL extended cache:** For additional cache control, you can integrate FlexShare buffer cache policies with the WAFL extended cache options. Doing so allows you to set caching policies on specific volumes. You can choose to enable only the FlexShare buffer cache policies without enabling all other FlexShare options.

	Buffer cache hits: Indicates the rate at which the WAFL buffer cache was successfully queried during the last measurement period.	Hits/Sec	A large number of cache misses and very few cache hits indicate that adequate entries are not available in the cache to service requests to the storage system. This in turn will force direct disk reads, thereby increasing the processing overheads of the storage system. To minimize disk reads and maximize cache reads, you can increase the WAFL cache memory using WAFL extended cache and the Performance Acceleration Module (PAM) family. WAFL extended cache is a software component of Data ONTAP and requires a license. WAFL extended cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. Once the extended cache is enabled, you can cache the following in it: • **Cache normal user data blocks**: If you cache normal user data blocks, the WAFL extended cache interprets this setting as the buffer cache policy of *keep* and saves normal user data blocks in the extended cache. • **Caching low-priority user data blocks**: You can cache low-priority user data blocks that are not normally stored by WAFL extended cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through a network-attached storage (NAS) protocol such as Network File System (NFS). Caching low-priority user data blocks is useful if you have workloads that fit within WAFL extended cache memory and if the workloads consist of either write followed by read or large sequential reads.

	Buffer cache misses: Indicates the rate at which an entry was not found in the the WAFL buffer cache upon a user query during the last measurement period.	Misses/sec	• **Caching only system metadata:** If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL extended cache memory by turning off both normal user data block caching and low-priority user data block caching. • **Integrating FlexShare buffer cache policies with WAFL extended cache:** For additional cache control, you can integrate FlexShare buffer cache policies with the WAFL extended cache options. Doing so allows you to set caching policies on specific volumes. You can choose to enable only the FlexShare buffer cache policies without enabling all other FlexShare options.
	Total number of buffers: Indicates the total number of buffers in this storage system.	Number	
	Number of available buffers: Indicates the number of available buffers in this storage system.	Number	A high value is desired for this measure.
	Total blocks read: Indicates the total number of blocks read from the WAFL buffer cache.	Number	
	Total blocks written: Indicates the total number of blocks written to the WAFL buffer cache.	Number	
	WAFL message rate: Indicates the total number of WAFL messages in this storage system.	Number	

	Average message latency: Indicates the average time taken for the execution of the WAFL messages during the last measurement period.	Milliseconds	Ideally, the value of this measure should be low. A high value indicates a slowdown indicating a processing bottleneck.
	Failures allocating extent messages: Indicates the total number of times the WAFL buffer failed to allocate the extent messages.	Number	Ideally, the value of this measure should be *0*. Too many failures may result in processing bottlenecks thus leading to the slowdown of the storage system.

1.6 The NetApp Access Layer

To monitor the load imposed by block access protocols and iSCSI connections to the storage system and to understand how well/poorly the NetApp system handles this load, use the tests mapped to this layer.

Figure 7: The tests mapped to the NetApp Access layer

1.6.1 NetApp Block I/O Protocol

Volumes are data containers. Clients can access the data in volumes through the access protocols supported by Data ONTAP. These protocols include Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Web-based Distributed Authoring and Versioning (WebDAV), Fibre Channel Protocol (FCP), and Internet SCSI (iSCSI).

Obviously, if one/more of these protocols are suddenly rendered unavailable, then clients will not be able to access critical data through these protocols. Moreover, whenever request processing delays are noticed, it becomes necessary for administrators to determine which protocol took the longest to perform read/write operations, so that slow protocol services can be identified. The **NetApp Block I/O Protocol** test provides these protocol-centric insights. For every protocol used for accessing data volumes, this test reports the availability of the protocol service, the rate of I/O operations performed through each protocol, and the time taken by each protocol to process read-write requests, so that problem-prone protocols can be accurately identified.

Purpose	For every protocol, this test reports the availability of the protocol service, the rate of I/O operations performed through each protocol, and the time taken by each protocol to process read-write requests, so that problem-prone protocols can be accurately identified
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each protocol that is active on the NetApp storage system being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Is service available?:
Indicates whether this protocol service is currently available.</td><td></td><td>This measure reports the value Yes if this protocol service is currently available and the value No if this protocol service is not available.
The values reported by this measure and their numeric equivalents are available in the table below:<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>Note:
This measure reports the Measure Values listed in the table above while indicating whether this protocol service is currently available or not. However, in the graph of this measure, the state is indicated using only the Numeric Values listed in the above table.</td></tr>
<tr><td></td><td>Operations rate:
Indicates the rate at which read/write operations were performed by users through this block protocol.</td><td>Ops/Sec</td><td></td></tr>
</table>

	Latency: Indicates the average time taken for performing the operations through this protocol.	Millisecs	A low value is desired for this measure. When users complaint of slowdowns when accessing data volumes, you can compare the value of this measure across protocols to know which protocol took the longest to perform the read-write operations.
	Read operations rate: Indicates the rate at which the read operations are performed across all LUNs of this storage system through this protocol.	Ops/Sec	Very high values for these measures are indicative of the existence of road-blocks to rapid reading/writing by the storage device. By observing the variations in these measures over time, you can understand whether the latencies are sporadic or consistent. Consistent delays in reading/writing could indicate that there are persistent bottlenecks (if any) in the storage device to speedy I/O processing.
	Read latency: Indicates the average time taken to perform read operations across all LUNs through this protocol.	Millisecs	
	Data read: Indicates the rate at which data is read from this storage system through this protocol.	Bytes/Sec	
	Write operations rate: Indicates the rate at which the write operations were performed across all LUNs of this storage system through this protocol.	Ops/Sec	
	Write latency: Indicates the average time taken to perform write operations across all LUNs through this protocol.	Millisecs	
	Data written: Indicates the rate at which data is written to this storage system through this protocol.	Bytes/Sec	

	Partner read latency: Indicates the average time taken to perform read operations across all the LUNs of the partner system (i.e., either the master/slave in a cluster setup of this storage system) through this protocol.	Millisecs	Very high values for these measures are indicative of the existence of road-blocks to rapid reading/writing by the storage device. By observing the variations in these measures over time, you can understand whether the latencies are sporadic or consistent. Consistent delays in reading/writing could indicate that there are persistent bottlenecks (if any) in the storage device to speedy I/O processing.
	Partner write latency: Indicates the average time taken to perform write operations on the LUNs of the partner system (i.e., either the master/slave in a cluster setup of this storage system) through this protocol.	Millisecs	

1.6.2 NetApp iSCSI Connections Test

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. The iSCSI protocol standard is defined by RFC 3720. In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). A host with an iSCSI host bus adapter (HBA), or running iSCSI initiator software, uses the iSCSI protocol to access LUNs on a storage system. The iSCSI protocol is implemented over the storage system's standard gigabit Ethernet interfaces using a software driver. The connection between the initiator and target uses a standard TCP/IP network. No special network configuration is needed to support iSCSI traffic. The network can be a dedicated TCP/IP network, or it can be your regular public network. The storage system listens for iSCSI connections on TCP port 3260.

This test monitors the iSCSI connections to the storage system, reports the load imposed by these connections on the storage device, and reveals the nature of these connections - i.e., the number of new connections, the number of connections used for data transfer, the number of connections used for discovery, and more.

Purpose	Monitors the iSCSI connections to the storage system, reports the load imposed by these connections on the storage device, and reveals the nature of these connections - i.e., the number of new connections, the number of connections used for data transfer, the number of connections used for discovery, and more
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

	10. **VFILERNAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box. 11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.		
Outputs of the test	One set of results for the NetApp storage system being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Number of initiators logged in: Indicates the number of initiators that were currently logged in.	Number	This measure reports the number of hosts with the initiator software that are currently accessing the LUNs on the storage device. This measure is a good indicator of the current workload of the device.
	Number of existing connections: Indicates the number of connections that were already established.	Number	This measure is a good indicator of the current workload of the device.
	Number of new connections: Indicates the number of connections that are not yet part of a session.	Number	
	Number of discovery sessions: Indicates the number of iSCSI sessions that are used to obtain information about iSCSI targets.	Number	

1.6.3 NetApp iSCSI Protocol Test

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. The iSCSI protocol standard is defined by RFC 3720. In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). A host with an iSCSI host bus adapter (HBA), or running iSCSI initiator software, uses the iSCSI protocol to access LUNs on a storage system. The iSCSI protocol is implemented over the storage system's standard gigabit Ethernet interfaces using a software driver. The connection between the initiator and target uses a standard TCP/IP network. No special

network configuration is needed to support iSCSI traffic. The network can be a dedicated TCP/IP network, or it can be your regular public network. The storage system listens for iSCSI connections on TCP port 3260.

This test monitors the active and attempted iSCSI sessions on the storage system, and promptly captures the processing ability, login failures, failed tasks, and errors encountered by these sessions.

Purpose	Monitors the active and attempted iSCSI sessions on the storage system, and promptly captures the processing ability, login failures, failed tasks, and errors encountered by these sessions
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

	10. **VFILERNAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box. 11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.		
Outputs of the test	One set of results for the NetApp storage system being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Command Descriptor Blocks CDB processed:** Indicates the total number of Command Descriptor Blocks that were processed by the inititator during the last measurement unit.	Number	The SCSI Command Descriptor Block (CDB) is a block of information that describes the command. Commands are sent from SCSI Initiators, which are contained in host computers, to SCSI Targets, which are controllers of some type of storage device (hard disk, tape drive, etc.). Almost every CDB contains 3 parts: l. a "What" field, m. a "Where" field, and n. a "How Much" field. For some commands, these fields are implied or not required. The "What" field is called the Operation Code (or OpCode) and tells the target what the command is supposed to do. A couple of examples would be READ or WRITE. The READ command moves data from the storage device to the host system, while the WRITE command moves data to the storage device for later access. The "Where" field tells the target where to begin the operation and is expressed as a Logical Block Address, or LBA. This address ranges from zero (0) to the maximum address of the device. Some commands, such as INQUIRY, do not require this field. The "How Much" field tells the target how many blocks (or bytes) or data to move. The block size of most storage devices is 512 bytes, but in certain storage devices, the block size can be different. This field is expressed as either Transfer Length (in blocks), Allocation Length (bytes moving to the host), or Parameter List Length (bytes moving to the device). Which name is used depends on the command itself. CDBs come in various sizes, typically 6, 10, 12, or 16 bytes total. Below is a figure of a 10-byte READ command to be sent to a hard drive. This command, if successful, will move one block (512 bytes) of data to the host computer system, from logical block address 100h (hex). All other bits or fields that are not labeled are set to zero. This measure is a good indicator for analyzing the traffic/load in this storage system.

	Successfully processed CDBs: Indicates the number of Command Descriptor Blocks that were successfully executed by the initiator during the last measurement period.	Number	A high value is desired for this measure. A low value indicates that there were too many unsuccessful CDB executions, which may have caused a processing bottleneck.
	CDBs with errors: Indicates the number of Command Descriptor Blocks that were processed by the initiators with errors during the last measurement period.	Number	Ideally, the value of this measure should be 0. A high value indicates that there were too many errors that occurred while processing the CDBs which may affect the performance of the storage system. Some of the common errors that occur while the CDBs are processed include the medium/hardware errors, providing illegal parameters for the CDB, accessing unauthorized data, volume overflow etc.
	Total errors: Indicates the total number of iSCSI errors that occurred during the last measurement period.	Number	Ideally, the value of this measure should be 0. Some of the common iSCSI errors that occur are digest errors, login/logout errors, PDU errors etc.
	Failed logins: Indicates the number of failed login attempts made by the initiator while creating new iSCSI sessions during the last measurement period.	Number	Ideally, the value of this measure should be 0.
	Failed logouts: Indicates the number of failed logouts while attempting to gracefully end the iSCSI sessions during the last measurement period.	Number	Ideally, the value of this measure should be 0.

	Failed tasks: Indicates the number of iSCSI tasks that failed during the last measurement period.	Number	
	Protocol errors: Indicates the number of protocol errors that occurred during the last measurement period.	Number	Ideally, the value of this measure should be 0. Protocol errors mainly occur due to the violation of protocol rules. The protocol errors occur in scenarios like violation of iSCSI PDU exchange sequences, duplication of protocol steps, invalid format/entries in protocol messages etc.
	Login requests: Indicates the number of login requests made during the last measurement period.	Number	This measure is an actual indicator of the users who are attempting to login to the storage system. Compare this value with the *Failed logins* measure to find out how well the user requests are processed in this storage system.
	Logout requests: Indicates the number of logout requests made during the last measurement period.	Number	This measure is an actual indicator of the users who are attempting to logout of the storage system. Compare this value with the *Failed logouts* measure to find out how well the user requests are processed in this storage system.
	Protocol Data Units rejected: Indicates the number of Protocol Data Units that were rejected by the initiator during the last measurement period.	Number	In a layered system such as iSCSI, a unit of data which is specified in a protocol of a given layer and which consists of protocol-control information and possibly user data of that layer is termed as a Protocol Data Unit. Ideally, the value of this measure should be 0. The Protocol Data Units are rejected due to iSCSI error conditions such as protocol errors, unsupported option etc., which may lead to connection/data loss, performance/processing bottleneck on the storage system etc.

1.7 The File Access Protocols Layer

The tests mapped to this layer monitors the CIFS and NFS operations on the NetApp storage system and reports I/O processing bottlenecks (if any).

Figure 8: The tests mapped to the File Access Protocols layer

1.7.1 CIFS Test

The Unified Storage Device (USD) exports data as files through two primary protocols, NFS and CIFS, which correspond to the UNIX and Windows processes.

Key features that CIFS offers are:

- **File Access with integrity**: CIFS supports the usual set of file operations; open, close, read, write and seek. CIFS also supports file and record lock and unlocking. CIFS allows multiple clients to access and update the same file while preventing conflicts by providing file sharing and file locking.
- **Optimization for Slow Links**: The CIFS protocol has been tuned to run well over slow-speed dial-up lines. The effect is improved performance for users who access the Internet using a modem.
- **Security**: CIFS servers support both anonymous transfers and secure, authenticated access to named files. File and directory security policies are easy to administer.
- **Performance and Scalability**: CIFS servers are highly integrated with the operating system, and are tuned for maximum system performance. CIFS supports all Microsoft platforms after Windows 95. It also supports other popular operation systems such as Unix, VMS, Macintosh, IBM LAN server etc.
- **Unicode File Names**: File names can be in any character set, not just character sets designed for English or Western European languages. Global File Names: Users do not have to mount remote file systems, but can refer to them directly with globally significant names, instead of ones that have only local significance.

By continuously tracking the status of the CIFS service and monitoring the read/write operations performed through the CIFS protocol, the **CIFS** test promptly detects and reports the non-availability of the service and provides you with a heads-up on probable latencies in the processing of I/O requests.

Purpose	For every protocol, this test reports the availability of the protocol service, the rate of I/O

	operations performed through each protocol, and the time taken by each protocol to process read-write requests, so that problem-prone protocols can be accurately identified
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results the NetApp storage system being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Service status:
Indicates the current status of the CIFS service.</td><td></td><td>The values reported by this measure and their numeric equivalents are available in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Starting</td><td>1</td></tr><tr><td>Started</td><td>2</td></tr><tr><td>Stopping</td><td>3</td></tr><tr><td>Stopped</td><td>4</td></tr></table>
Note:
This measure reports the Measure Values listed in the table above while indicating the current status of the CIFS service. However, in the graph of this measure, the status is indicated using only the Numeric Values listed in the above table.</td></tr>
<tr><td></td><td>Operations:
Indicates the rate at which operations were performed by users through CIFS protocol to access this NetApp Unified Storage system.</td><td>Ops/Sec</td><td></td></tr>
</table>

	Latency: Indicates the average time taken for performing the operations through the CIFS protocol.	Millisecs	A low value is desired for this measure.
	Read operations: Indicates the rate at which the read operations are performed across all LUNs of this storage system through the CIFS protocol.	Ops/Sec	Very high values for these measures are indicative of the existence of road-blocks to rapid reading/writing by the storage device. By observing the variations in these measures over time, you can understand whether the latencies are sporadic or consistent. Consistent delays in reading/writing could indicate that there are persistent bottlenecks (if any) in the storage device to speedy I/O processing.
	Read latency: Indicates the average time taken to perform read operations across all LUNs through the CIFS protocol.	Millisecs	
	Write operations: Indicates the rate at which the write operations were performed across all LUNs of this storage system through the CIFS protocol.	Ops/Sec	
	Write latency: Indicates the average time taken to perform write operations across all LUNs through the protocol.	Millisecs	

1.7.2 NetApp IGroup Config Mismatches Test

Initiator groups (igroups) are tables of host identifiers (FCP, WWPNs, or iSCSI node names) that are used to control hosts' access to LUNs.

igroups specify which initiators have access to which LUNs. igroups can be created either before or after LUNs are created, but they must be created before a LUN is mapped to an igroup. Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. However, a LUN can not be mapped to multiple igroups that have the same initiator.

An initiator cannot be a member of igroups of differing ostypes.

This test reveals if any mismatch of the cluster failover setting has occurred between the local and partner systems of a cluster and also the following:

- How many initiator groups have an incompatible operating system?
- How many initiator groups are with an invalid use partner setting?

- How many initiator groups have an operating system that is incompatible with the use partner setting and the VSA setting?
- How many initiator groups have the ALUA setting mismatch between the local and partner systems?

Purpose	Reveals if any mismatch of the cluster failover setting has occurred between the local and partner systems of a cluster and also following: • How many initiator groups have an incompatible operating system? • How many initiator groups are with an invalid use partner setting? • How many initiator groups have an operating system that is incompatible with the use partner setting and the VSA setting? • How many initiator groups have the ALUA setting mismatch between the local and partner systems?
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

<table>
<tr>
<td>test</td>
<td>Is CF mode mismatched?:
Indicates whether/not the cluster failover setting of the local system is different from that of the partner system.</td>
<td></td>
<td>In a cluster setup of the storage systems, the Cluster failover modes of the systems in that cluster need to match so that the cluster failover would function appropriately. This measure reports the value Yes if the cluster failover setting is different in the local system and the partner system and the value No if the settings are same in both the local system and the partner system.

The values reported by this measure and their numeric equivalents are available in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>No</td><td>0</td></tr>
<tr><td>Yes</td><td>1</td></tr>
</table>
Note:

This measure reports the Measure Values listed in the table above to indicate whether/not the CF mode is mistmatched. However, in the graph of this measure, the same is indicated using only the Numeric Values listed in the above table.

The Detailed Diagnosis of this measure shows the mismatching local and partner Cluster Failover modes.</td>
</tr>
<tr>
<td></td>
<td>Igroups with invalid OS type:
Indicates the number of initiator groups whose Operating system is not compatible with that of the configured fcp cfmode.</td>
<td>Number</td>
<td>Some host operating systems are compatible with certain selective fcp cfmodes only. In such a scenario, the OS type of the initiator group must match the fcp cfmode for proper functioning of the intiator group.

The detailed diagnosis of this measure indicates the name of the initiator group and the OS type of the initiator group.</td>
</tr>
<tr>
<td></td>
<td>Igroups with invalid use-partner setting:
Indicates the number of initiator groups with an invalid use-partner setting - i.e., the use partner setting of the initiator group is not compatible with that of the configured fcp cfmode.</td>
<td>Number</td>
<td>The use-partner setting indicates whether the initiators in the initiator group are allowed to use the partner's port ina cluster setup. Initiator groups with an invalid use-partner setting can result in some hosts losing LUNs during takeover.

The detailed diagnosis of this measure indicates the name of the affected initiator group and the use partner setting of that corresponding initiator group.</td>
</tr>
</table>

	Igroups with mismatching use-partner OS type setting: Indicates the number of initiator groups whose use partner setting is not compatible with their configured operating system.	Number	In a cluster setup, in order to ensure proper behavior of the Storage systems during failover, the host operating sytems are designed to support only certain use parameter settings. A difference in the use parameter setting may result in performance bottleneck. The detailed diagnosis of this measure highlights this incompatibility issue by listing out the initiator group, OS type and the use parameter setting of that initiator group.
	Igroups with invalid ALUA setting: Indicates the number of initiator groups for which the ALUA setting do not match between the local and the partner unified storage systems.	Number	ALUA is a T10 standard that specifies the access characteristics (in terms of performance and supported SCSI commands) of a Logical Unit that can be accessed through more than one target port. ALUA is typically used by host multi-path software to recognize primary and secondary paths to a Logical Unit when more than one path are available to the Logical Unit. If the ALUA setting does not match between the local and partner filers, it would affect the host multi-path software's ability to distinguish between the primary and secondary paths, which could lead to the poor performance of the system. The detailed diagnosis of this measure provides the name of the affected initiator group and the status of the ALUA (whether enabled or disabled) in the local and partner storage systems.
	Igroups with invalid VSA setting: Indicates the number of initiator groups with the VSA setting that do not match with that of the configured operating system.	Number	The Volume Set Addressing (VSA) setting is enabled only for the initiator groups that are configured with the HPUX operating system. Incorrect settings of the VSA may deny hosts access to some/all LUNs. The detailed diagnosis of this measure indicates the name of the initiator group, the operating system of the initiator groups and the status (whether enabled/disabled) of the VSA setting.

1.7.3 NetApp NFS I/O Test

NFS (Network File System) is a protocol used by Unix system to access data on the storage system.

This test auto-discovers the versions of NFS used on the storage system, and reports the following for each NFS version:

- The status of the NFS server;
- Whether all NFS messages have been drained from the NFS queue and the server has been disabled;
- The rate of read-write requests processed by the NFS server and latencies in I/O processing (if any)

Purpose	Auto-discovers the versions of NFS used on the storage system, and reports the following for each NFS version: • The status of the NFS server; • Whether all NFS messages have been drained from the NFS queue and the server has been disabled; • The rate of read-write requests processed by the NFS server and latencies in I/O processing (if any)
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each version of NFA on the NetApp storage system being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Is the NFS server running?:
Indicates whether the NFS server is currently running in this storage system.</td><td></td><td>The values reported by this measure and their numeric equivalents are available in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>No</td><td>0</td></tr><tr><td>Yes</td><td>1</td></tr></table>
Note:
This measure reports the Measure Values listed in the table above to indicate current state of the NFS server. However, in the graph of this measure, the same is indicated using only the Numeric Values listed in the above table.</td></tr>
</table>

<table>
<tr><td></td><td>Have all messages been drained?:
Indicates whether all the NFS messages have been cleared off from the NFS queue and the NFS server has been disabled.</td><td></td><td>The values reported by this measure and their numeric equivalents are available in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>No</td><td>0</td></tr><tr><td>Yes</td><td>1</td></tr></table>
Note:

This measure reports the Measure Values listed in the table above to indicate whether the NFS messages have been cleared off and the NFS server is disabled. However, in the graph of this measure, the same is indicated using only the Numeric Values listed in the above table.</td></tr>
<tr><td></td><td>Operations:
Indicates the total number of NFS operations per second for this NFS version.</td><td>Ops/Sec</td><td></td></tr>
<tr><td></td><td>Average operations latency:
Indicates the average time taken for any NFS operation that has happened for this NFS version.</td><td>Milliseconds</td><td>Ideally, the value of this measure should be low. A higher value is an indication of too many NFS messages waiting in the NFS queue thus leading to a processing bottleneck.</td></tr>
<tr><td></td><td>Read operations:
Indicates the rate at which the NFS read operations were performed for this NFS version.</td><td>Ops/Sec</td><td rowspan="2">Very high values for these measures are indicative of the existence of road-blocks to rapid reading/writing by the storage device.

By observing the variations in these measures over time, you can understand whether the latencies are sporadic or consistent. Consistent delays in reading/writing could indicate that there are persistent bottlenecks (if any) in the storage device to speedy I/O processing.</td></tr>
<tr><td></td><td>Read latency:
Indicates the average time taken for the NFS read operation for this NFS version.</td><td>Milliseconds</td></tr>
</table>

	Write operations: Indicates the rate at which the NFS write operations were performed for this NFS version.	Ops/Sec	
	Write latency: Indicates the average time taken for the NFS write operation for this NFS version.	Milliseconds	

1.8 The Logical Storage Layer

Using the tests associated with this layer, the following can be monitored:

- Usage of volumes to isolate the over-used and overloaded volumes;
- Snapshot usage to identify the snapshots that can be deleted to conserve space;
- Status of clone operations
- Disk and file usage quotas
- Space usage in LUNs
- LUN config errors

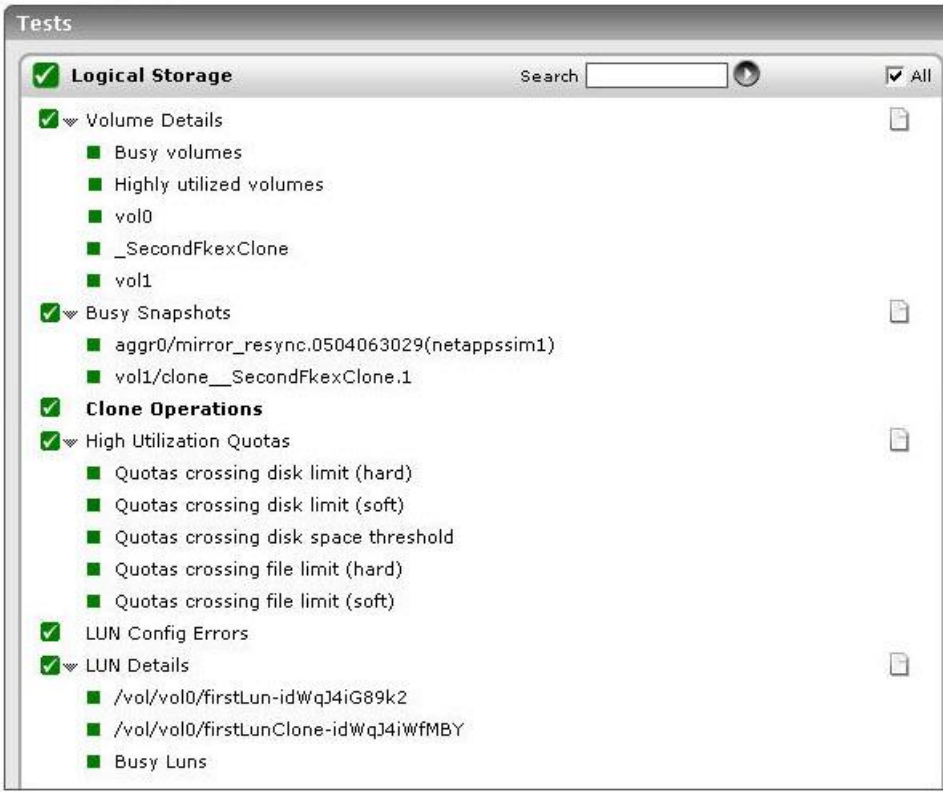

Figure 9: The tests mapped to the Logical Storage layer

1.8.1 NetApp Volume Details Test

Volumes contain file systems that hold user data that is accessible using one or more of the access protocols supported by Data ONTAP, including NFS, CIFS, HTTP, FTP, FC, and iSCSI.

For users to be able to read from/write data into volumes quickly, adequate space must be available in the volumes and the I/O requests should be processed rapidly by the volumes. Slowdowns in data storage/retrieval can be attributed to storage space contentions experienced by one/more volumes or I/O processing bottlenecks. In the event of such slowdowns, administrators need to swiftly isolate the following:

- Which volumes are over-utilized?
- Which volumes are overloaded?
- Which volumes are experiencing serious latencies?
- When were these latencies observed most frequently – while reading or writing?
- What type of operations registered the maximum latency – CIFS, NFS, or iSCSI?

The **NetApp Volume Details** test provides accurate answers to these questions. With the help of these answers, you can quickly diagnose the root-cause of slowdowns when reading from/writing into a volume.

Purpose	Helps quickly identify problematic volumes and accurately diagnose the root-cause of slowdowns when reading from/writing into a volume
Target of the test	A NetApp Unified Storage

Agent deploying the test	An external/remote agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

	10. **VFILERNAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box. 11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds. 12. **USED PERCENTAGE THRESHOLD** – This test not only reports a set of metrics for each volume on the storage device, but also reports metrics for the following descriptors: *Busy volumes*, *Slow volumes*, and *Highly utilized volumes*. By default, the *Highly utilized volumes* descriptor will report metrics for those volumes in which over 80% of space has already been utilized. This is why, the **USED PERCENTAGE THRESHOLD** is set to **80** by default. You can change this threshold by specifying a different percentage value against **USED PERCENTAGE THRESHOLD**. **This parameter is deprecated in v5.6.5 (and above).** 13. **OPERATIONS THRESHOLD** - This test not only reports a set of metrics for each volume on the storage device, but also reports metrics for the following descriptors: *Busy volumes*, *Slow volumes*, and *Highly utilized volumes*. The **OPERATIONS THRESHOLD** value (in operations/sec) you set determines which volumes will be counted as *Busy volumes* by this test. Typically, if the rate of operations to a volume exceeds the rate specified against **OPERATIONS THRESHOLD**, then the test will consider such a volume to be a *Busy volume*. **This parameter is deprecated in v5.6.5 (and above).** 14. **AVG LATENCY THRESHOLD** - This test not only reports a set of metrics for each volume on the storage device, but also reports metrics for the following descriptors: *Busy volumes*, *Slow volumes*, and *Highly utilized volumes*. The **AVG LATENCY THRESHOLD** value (in milliseconds) you set determines which volumes will be counted as *Slow volumes* by this test. Typically, if the latency registered by a volume falls exceeds the **AVG LATENCY THRESHOLD** you specify, then the test will consider such a volume to be a *Slow volume*. **This parameter is deprecated in v5.6.5 (and above).** 15. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.
	16. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Outputs of the test	One set of results for each volume on the NetApp storage system being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Number of volumes: Indicates the number of volumes that are currently highly utilized/slow/busy.	Number	a. This measure appears only for the *Highly utilized*, *Slow* and *Busy* volumes. In the case of *Highly utilized volumes*, the detailed diagnosis of this measure if enabled, lists the names of the highly utilized volumes and the percentage of space that is utilized in each volume. b. In the case of *Slow volumes*, the detailed diagnosis of this measure if enabled, lists the names of the slow volumes and the average latency i.e., the time taken to perform read/write operations on each volume. c. In the case of *Busy volumes*, the detailed diagnosis of this measure if enabled, lists the names of the busy volumes and the rate at which operations were performed on each volume. d. With the help of the detailed diagnosis information therefore, you can quickly identify the highly utilized, slow, and busy volumes. e. **This measure is deprecated in v5.6.5 (and above).**

<table>
<tr>
<td></td>
<td>State:
Indicates the current state of this volume.</td>
<td></td>
<td>The values that this measure can report and their corresponding numeric equivalents are shown in the table below:

<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Online</td></tr>
<tr><td>1</td><td>Creating</td></tr>
<tr><td>2</td><td>Restricted</td></tr>
<tr><td>3</td><td>Offline</td></tr>
<tr><td>4</td><td>Partial</td></tr>
<tr><td>5</td><td>Unknown</td></tr>
<tr><td>6</td><td>Failed</td></tr>
</table>

Note:

By default, this measure reports the above-mentioned Measure Values while indicating the current state of a volume. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents only.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Is volume in error?
Indicates whether/not this volume is error-prone.</td>
<td></td>
<td>Generally, errors may be caused when the volume is inconsistent, unrecoverable or invalid. A volume is considered to be inconsistent if there exists known inconsistencies in the associated file system. An increase in the inconsistencies will render the volume unrecoverable. Unrecoverable volumes cannot be accessed. If mirroring has been enabled, Data ONTAP will automatically access the mirrored data of the unrecoverable volume. A volume is said to be invalid if a vol-copy or Snapmirror initial transfer has been aborted. Such invalid volumes are generally partially created and cannot be recovered fully. Operation errors are taken into account if this volume is a Single Instance Storage (SIS) volume.

This measure reports the value Yes if a volume is error-prone and the value No if it is error-free.

The numeric values that correspond to the above-mentioned values are represented in the table below:

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

Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether/not this volume is error-prone. However, in the graph of this measure the same will be represented using the corresponding numeric equivalents only.

The detailed diagnosis capability of this measure, if enabled, lists the type of the error. In the case of an SIS operation error, the actual SIS error message will also be displayed as part of the detailed diagnosis.

This measure is applicable only to individual volumes.</td>
</tr>
</table>

	Used space percentage: Indicates the percentage of space that is utilized in this volume.	Percent	Ideally, the value of this measure should be low. A high value or a consistent increase in the value of this measure is indicative of excessive space usage in a volume. This measure will be *0* for restricted and offline volumes.
	Total size: Indicates the total size of this volume.	MB	The value of this measure will not include the WAFL reserve and the volume snapshot reserve. This measure will be *0* for restricted and offline volumes.
	Reserve space: Indicates the space that is reserved for overwriting snapshotted data in this volume.	MB	This space can be utilized only by space reserved LUNs and files and only when the volume is **full**. This measure will be *0* for restricted and offline volumes.
	Actual reserved space used: Indicates the percentage of reserved space that is actually used by this volume.	Percent	A low value is desired for this measure. This measure will be *0* for restricted and offline volumes.
	Files used percentage: Indicates the percentage of inodes i.e., files that are currently utilized in this volume.	Percent	A high value indicates that the inodes in the volume may get exhausted soon. This measure will be *0* for restricted and offline volumes.
	Total operations: Indicates the rate at which operations (including read and write) were performed on this volume.	Ops/Sec	This measure is a good indicator of how busy the volume is. Comparing the value of this measure across volumes will enable you to quickly detect load-balancing irregularities (if any).
	Write operations: Indicates the rate at which write operations were performed on this volume.	Ops/Sec	
	Read operations: Indicates the rate at which read operations were performed from this volume.	Ops/Sec	

	Avg latency: Indicates the average time taken by the WAFL filesystem to process all the operations performed on this volume	Microseconds	The value of this measure excludes the request processing time and the network communication time of the volume. A high value of this measure is a cause for concern, as it indicates a processing bottleneck.
	Read latency: Indicates the average time taken by the WAFL filesystem to process the read requests of this volume.	Microseconds	The value of these measures exclude the request processing time and the network communication time of the volume. If the *Avg latency* of a volume is high, then you can compare the value of these measures for that volume to know when the latency occurred – while reading or writing?
	Write latency: Indicates the average time taken by the WAFL filesystem to process the write requests made to this volume.	Microseconds	
	Read data: Indicates the rate at which data bytes were read from this volume.	Bytes/Sec	
	Write data: Indicates the rate at which data bytes were written to this volume.	Bytes/Sec	
	CIFS operations: Indicates the rate at which the CIFS operations were performed on this volume.	Ops/Sec	This measure is inclusive of all the CIFS operations i.e., read, write and other miscellaneous CIFS operations. By comparing the value of this measure with that of the *NFS operations* and *SAN operations* measures for a volume, you can figure out which type of operation imposed the maximum load on that volume.
	NFS operations: Indicates the rate at which the NFS operations were performed on this volume.	Ops/Sec	This measure is inclusive of all the NFS operations i.e., read, write and other miscellaneous NFS operations. By comparing the value of this measure with that of the *CIFS operations* and *SAN operations* measures for a volume, you can figure out which type of operation imposed the maximum load on that volume.

	SAN operations: Indicates the rate at which the SAN operations were performed on this volume.	Ops/Sec	This measure is inclusive of all the SAN operations i.e., read, write and other miscellaneous SAN operations. By comparing the value of this measure with that of the *CIFS operations* and *NFS operations* measures for a volume, you can figure out which type of operation imposed the maximum load on that volume.

	CIFS latency: Indicates the average time taken for performing the CIF operations (including read, write and other miscellaneous CIF operations) on this volume.	Microseconds	The value of these measures exclude the request processing time and the network communication time of the volume. Ideally, the value of these measure should be low. If the *Avg latency* of a volume is very high, then, you can compare the value of these measures for that volume to determine the reason for the latency – is it because of processing bottlenecks experienced by CIFS operations? NFS operations? Or SAN operations?
	NFS latency: Indicates the average time taken for performing the NFS operations (including read, write and other miscellaneous NFS operations) on this volume.	Microseconds	
	SAN latency: Indicates the average time taken for performing the block protocol operations (including read, write and other miscellaneous block protocols operations) on this volume.		

1.8.2 Busy Snapshots Test

A Snapshot copy is a point-in-time file system image. Low-overhead Snapshot copies are made possible by the unique features of the WAFL® (Write Anywhere File Layout) storage virtualization technology that is part of Data ONTAP®. Like a database, WAFL uses pointers to the actual data blocks on disk, but, unlike a database, WAFL does not rewrite existing blocks; it writes updated data to a new block and changes the pointer. A NetApp Snapshot copy simply manipulates block pointers, creating a "frozen" read-only view of a WAFL volume that lets applications access older versions of files, directory hierarchies, and/or LUNs (logical unit numbers) without special programming.

Whenever a volume/LUN/aggregate runs out of space, you may want to clear some space in that storage device so that there is no road-block to freely reading from and writing data into that device. To make room in a volume/LUN/aggregate for more data, you can start by deleting some snapshot copies from that storage device. Before attempting deletion however, you may want to determine the following:

- How much space is occupied by each snapshot on the storage system?
- Which snapshot copies, when deleted, will free more space? Will deleting the complete snapshot series make more space available?
- Which snapshot copies are way too old, and are hence ideal candidates for deletion?
- Which snapshots are easier to delete? – the snapshots containing LUN clones may take longer to delete as the LUN clones will first have to be deleted and then the snapshots.

Using the **Busy Snapshots** test, you can find quick and accurate answers for the questions above. This test auto-

discovers the snapshot copies on the storage system, and for each snapshot copy, reports the space used by the snapshot copy, the age of the copy, and whether the copy contains LUN clones or not. Deletion decisions can be taken based on the insights provided by this test.

Purpose	Auto-discovers the snapshot copies on the storage system, and for each snapshot copy, reports the space used by the snapshot copy, the age of the copy, and whether the copy contains LUN clones or not. Deletion decisions can be taken based on the insights provided by this test
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">
10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each snapshot copy on the NetApp storage system being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>sInterpretation</td></tr>
<tr><td></td><td>Total blocks:
Indicates the percentage of blocks that were owned by this snapshot with respect to the total number of blocks in this volume.</td><td>Percent</td><td>Using this measure, large snapshots can be identified easily and helps you to decide whether this snapshot can be deleted so that the snapshot reserve space can be reclaimed.</td></tr>
</table>

	Used blocks: Indicates the percentage of blocks that were owned by this snapshot with respect to the number of blocks that were currently used in this volume.	Percent	If a high percentage of used space in a volume/LUN/aggregate is in fact used up by a snapshot, then deleting such a snapshot can instantly reduce the space usage in that volume/LUN/aggregate, thereby enabling that storage medium to accommodate more data. Compare the value of this measure across snapshots to identify the snapshot that is occupying a lot of used space.
	Total blocks in volume: Indicates the percentage of total blocks that were owned by the snapshot series (including this snapshot) with respect to the total number of blocks in this volume.	Percent	Comparing the value of this measure across snapshots will instantly reveal the large-sized snapshots. These snapshots, when deleted, will release a large amount of snapshot reserve space.
	Used blocks in volume: Indicates the percentage of total blocks that were owned by the snapshot series (including this snapshot) with respect to the total number of blocks that were currently used in this volume.	Percent	If the containing volume is running out of space, then this measure is used to clearly indicate if too much of space is occupied by the snapshot series in this volume and helps you to identify the amount of space that can be reclaimed by deleting one or more snapshots from the snapshot series.
	Snapshot age: Indicates the number of days that have elapsed since this snapshot was created.	Days	Generally, the snapshots should not be older than *two* weeks. This measure helps you to identify the snapshots that are old enough to be deleted.

	Contains lun clones?: Indicates whether/not this snapshot contains LUN clones.		Lun clones are editable copies of LUNs which are backed by a snapshot. These snapshots cannot be deleted without deleting the associated LUN clones that are referenced by the snapshot. This measure reports the value *Yes* if LUN clones exist for a snapshot and *No* if the LUN clones do not exist for a snapshot. The corresponding numeric equivalents for the measures are detailed in the table below: **Numeric Value** / **Measure Value** 0 / No 1 / Yes **Note:** By default, this measure reports the above-mentioned **Measure Value**s while indicating whether/not this snapshot contains LUN clones. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.

1.8.3 NetApp High Utilization Quotas Test

Quotas are specified for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a user or group, or that can be contained by a qtree.
- To track the amount of disk space or the number of files used by a user, group, or qtree, without imposing a limit.
- To warn users when their disk usage or file usage is high

You specify quotas using the */etc/quotas* file. Quotas are applied to a specific volume or qtree.

When Data ONTAP receives a request to write to a volume, it checks to see whether quotas are activated for that volume. If so, Data ONTAP determines whether any quota for that volume (and, if the write is to a qtree, for that qtree) would be exceeded by performing the write operation. If any hard quota would be exceeded, the write operation fails, and a quota notification is sent. If any soft quota would be exceeded, the write operation succeeds, and a quota notification is sent.

This test reports the number of Windows/Unix users and Unix user groups that crossed the disk space (both hard and soft) and file usage quotas set. With the help of these metrics, you can promptly detect abnormal disk space and file usage at the volume/qtree-level.

Purpose	Reports the number of Windows/Unix users and Unix user groups that crossed the disk space

	(both hard and soft) and file usage quotas set
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">
10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of quota set at the volume/qtree-level</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>sInterpretation</td></tr>
<tr><td></td><td>Quotas:
Indicates the number of quotas of this type.</td><td>Number</td><td>The detailed diagnosis of this measure indicates whether the quota has been set for a user/group/qtree, the target of the quota, the volume on which the quota is applied, the qtree on which the quota is applied, the percentage of disk limit used, the percentage of file limit used, and the number of windows users, unix users and unix group users who violated each type of quota.</td></tr>
</table>

	Windows users: Indicates the number of windows users involved in quota violation of this quota type.	Number	Ideally, the value of these measures should be low. A high value indicates that there is space constraint in the disk/volume/LUN which in turn will affect the users who are accessing them.
	Unix users: Indicates the number of unix users involved in quota violation of this quota category.	Number	
	Unix groups: Indicates the number of unix groups involved in quota violation of this quota category.	Number	

1.8.4 NetApp Clone Operations Test

The cloning feature is based on WAFL block sharing and provides fast and almost 100% space efficient file and sub-file cloning, which can also be applied for LUN and sub-LUN cloning.

This test reports the number of clone operations that are currently running and the number of clone operations that have failed in the NetApp Unified Storage system.

Purpose	Reports the number of clone operations that are currently running and the number of clone operations that have failed in the NetApp Unified Storage system
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Running operations:
Indicates the number of clone operations that are currently running in this storage system.</td><td>Number</td><td>The detailed diagnosis of this measure lists the Operation ID, UUID of the cloned volume, type of the cloning operation (may be a file, sub file, LUN or the sub LUN of the NetApp Unified Storage), and percentage completion for each of the running cloning operation.</td></tr>
</table>

	Failed operations: Indicates the number of clone operations that have failed during the last measurement period.	Number	Cloning may fail due to insufficient disk space, permission issues etc., This failure information will be stored in a metadata file in the disk and have to be manually cleared. If the metadata file has been cleared during a measure period, this measure will be zero even if a cloning failure had occurred prior to the file getting cleared during the same measure period. The detailed diagnosis shows the Operation ID, UUID of the cloned volume, type of the cloning operation (may be a file, sub file, LUN or the sub LUN of the NetApp Unified Storage), reason for the failure, error code, and the percentage completion for each failed cloning operation.

f.

1.8.5 NetApp LUN Config Errors Test

LUN conflicts may result in various issues such as data inconsistencies in the LUN as hosts may overwrite each others data or may lead to LUN reservation issues.

Conflicts may arise when a LUN is mapped to both the FCP and iSCSI initiator groups with the ALUA setting being enabled on atleast one of the initiator groups. Further mapping to the conflicted LUN will be possible only when the conflicts are resolved.

This test reports the count of LUNs that are victims of such a conflict. Using the detailed diagnosis of this test, you can identify the affected LUNs as well.

Purpose	Reports the count of LUNs that are victims of such a conflict. Using the detailed diagnosis of this test, you can identify the affected LUNs as well
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>LUNs with ALUA conflicts:
Indicates the number of LUNs that were mapped to both FCP and iSCSI initiator groups with the ALUA setting being enabled in both/atleast one of these initiator groups.</td><td>Number</td><td>A non-zero value is indicative of the existence of LUN conflicts. These conflicts can either be resolved by unmapping one or more mappings from the conflicting LUN or by disabling the ALUA setting on FCP or iSCSI or both the initiator groups whichever is applicable.
The detailed diagnosis capability of this measure if enabled, lists out the path of the LUNs with this ALUA setting conflict.</td></tr>
</table>

1.8.6 NetApp LUNs Test

This test auto-discovers the LUNs configured on the NetApp Unified Storage system, monitors the availability, state,

and the processing ability of each LUN, and reports the following:

- Which LUNs are currently offline?
- Is any LUN experiencing a contention for storage space?
- Is I/O load uniformly balanced across all LUNs, or is any LUN overloaded? Is it causing the LUN to receive an increased number of *Queue Full* responses?
- Are the LUNs able to process the I/O requests quickly? Is any LUN experiencing processing bottlenecks?

Purpose	Auto-discovers the LUNs configured on the NetApp Unified Storage system, monitors the availability, state, and the processing ability of each LUN
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

	10. **VFILERNAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box. 11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds. 12. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**. 13. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for each LUN configured on the NetApp storage system being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr><td>test</td><td>Is LUN online?:
Indicates whether/not this LUN is online.</td><td></td><td>This measure is applicable only for the individual LUNs. This measure reports a value Yes if this LUN is currently available online and a value No if this LUN is not available online.

The numeric equivalents corresponding to the above-mentioned values are listed in the table below:

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

This measure reports the Measure Values listed in the table above to indicate the current state of a LUN. However, in the graph of this measure, the same is indicated using only the Numeric Values listed in the above table.</td></tr>
<tr><td></td><td>Size:
Indicates the size of this LUN in the active file system.</td><td>MB</td><td></td></tr>
<tr><td></td><td>Size used:
Indicates the currently used size of this LUN.</td><td>MB</td><td>A low value is desired for this measure. A high value indicates that the LUN is running out of space.</td></tr>
<tr><td></td><td>Read operations:
Indicates the rate at which the read operations were performed on this LUN.</td><td>Ops/Sec</td><td>A high value is desired for this measure. A consistent decrease in this value could indicate a processing bottleneck.</td></tr>
<tr><td></td><td>Write operations:
Indicates the rate at which the write operations were performed to this LUN.</td><td>Ops/Sec</td><td>A high value is desired for this measure. A consistent decrease in this value could indicate a processing bottleneck.</td></tr>
<tr><td></td><td>Total operations:
Indicates the rate at which the operations (incuding the read and write) were performed on this LUN.</td><td>Ops/Sec</td><td>A high value is desired for this measure. A consistent decrease in this value could indicate a processing bottleneck.</td></tr>
</table>

	Average latency: Indicates the average time taken for executing an operation in this LUN.	Milliseconds	A high value indicates that the LUN is taking too long to process the I/O requests to it. Compare the value of this measure across LUNs to isolate the *slow LUNs*.
	Queue full responses: Indicates the rate at which the queue full responses were received on this LUN.	Responses/Sec	This measure is a good indicator for detecting sudden/co=ordinated bursts of I/O from the initiators. A *Queue full* condition signals that the target/storage port is unable to process more I/O requests and thus the initiator will need to throttle I/O to the storage port. Some operating systems like AIX may not handle repeated *Queue full* responses gracefully i.e., will not throttle the I/O requests appropriately leading to I/O errors. These conditions can also be alleviated by reducing the LUN queue depth setting appropriately.
	Read data: Indicates the rate at which data is read from this LUN.	Bytes/Sec	A high value is desired for this measure.
	Write data: Indicates the rate at which data is written to this LUN.	Bytes/Sec	A high value is desired for this measure.
	Queue depth: Indicates the queue depth of this LUN.	Number	Queue Depth is the number of outstanding I/O requests a LUN will issue or hold before the LUN can trigger a Queue Full response i.e., the number of I/O operations that can run in parallel on the LUN. This is useful when compared to the number of Queue Full responses triggered by the LUN. Queue depth is usually set too high and hence could contribute significantly to latency if improperly set.

	Average read latency: Indicates the average time taken to execute a read request in this LUN.	Milliseconds	A low value is desired for this measure. A high value indicates that the requests take too long to execute which directly affects the performance of the LUNs.
	Average write latency: Indicates the average time taken to execute a write request in this LUN.	Milliseconds	

1.9 The NetApp System Layer

Besides reporting system status and overall performance, the tests mapped to this layer report the count of errors/warnings logged in the Syslog and events captured by the virus scanner.

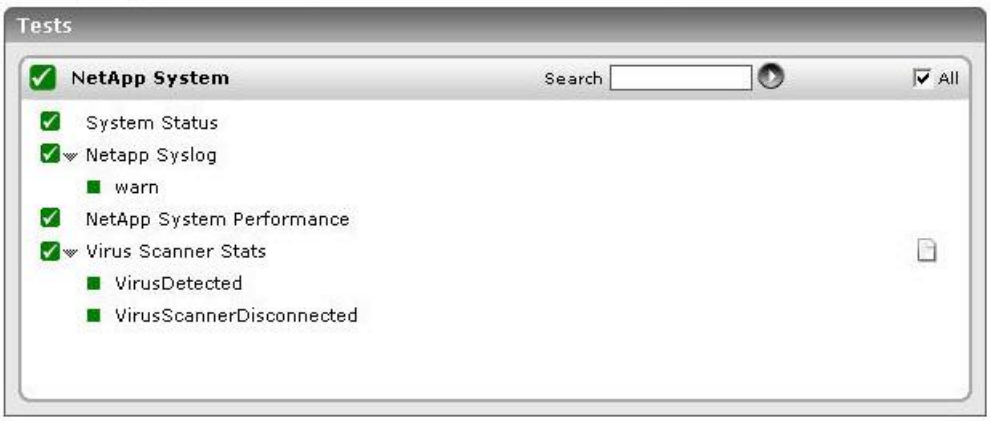

Figure 10: The tests mapped to the NetApp System layer

1.9.1 Virus Scanner Stats Test

A storage system can suffer performance setbacks or slowdowns if virus scanners detect malicious virus attacks on the system or if virus scanners are unable to access the system to run virus checks. This test promptly captures such failure events and intimates administrators of the same.

Purpose	Promptly captures failure events related to the virus scanner and intimates administrators of the same
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test

1. **TESTPERIOD** - How often should the test be executed
2. **HOST -** The host for which the test is to be configured.
3. **SOURCEADDRESS** - Specify a comma-separated list of IP addresses or address patterns of the hosts from which traps are considered in this test. For example, *10.0.0.1,192.168.10.**. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.
4. **OIDVALUE** - Provide a comma-separated list of OID and value pairs returned by the traps. The values are to be expressed in the form, *DisplayName:OID-OIDValue*. For example, assume that the following OIDs are to be considered by this test: *.1.3.6.1.4.1.9156.1.1.2* and *.1.3.6.1.4.1.9156.1.1.3*. The values of these OIDs are as given hereunder:

OID	Value
.1.3.6.1.4.1.9156.1.1.2	Host_system
.1.3.6.1.4.1.9156.1.1.3	NETWORK

In this case the **OIDVALUE** parameter can be configured as *Trap1:.1.3.6.1.4.1.9156.1.1.2-Host_system,Trap2:.1.3.6.1.4.1.9156.1.1.3-Network*, where *Trap1* and *Trap2* are the display names that appear as descriptors of this test in the monitor interface.

An * can be used in the OID/value patterns to denote any number of leading or trailing characters (as the case may be). For example, to monitor all the OIDs that return values which begin with the letter 'F', set this parameter to *Failed:*-F**.

Typically, if a valid value is specified for an OID in the *OID-value* pair configured, then the test considers the configured OID for monitoring only when the actual value of the OID matches with its configured value. For instance, in the example above, if the value of OID *.1.3.6.1.4.1.9156.1.1.2* is found to be *HOST* and not *Host_*system, then the test ignores OID *.1.3.6.1.4.1.9156.1.1.2* while monitoring. In some cases however, an OID might not be associated with a separate value – instead, the OID itself might represent a value. While configuring such OIDs for monitoring, your **OIDVALUE** specification should be: *DisplayName:OID-any*. For instance, to ensure that the test monitors the OID *.1.3.6.1.4.1.9156.1.1.5*, which in itself, say represents a failure condition, then your specification would be:

Trap5: .1.3.6.1.4.1.9156.1.1.5-any.

<table>
<tr><td></td><td colspan="3">5. SHOWOID – Specifying true against SHOWOID will ensure that the detailed diagnosis of this test shows the OID strings along with their corresponding values. If you enter false, then the values alone will appear in the detailed diagnosis page, and not the OIDs.

6. TRAPOIDS – By default, this parameter is set to all, indicating that the eG agent considers all the traps received from the specified SOURCEADDRESSes. To make sure that the agent considers only specific traps received from the SOURCEADDRESS, then provide a comma-separated list of OIDs in the TRAPOIDS text box. A series of OID patterns can also be specified here, so that the test considers only those OIDs that match the specified pattern(s). For instance, *94.2*,*.1.3.6.1.4.25*, where * indicates leading and/or trailing spaces.

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

8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each type of virus scanner-related failure event that occurred on the target storage system</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Number of messages:
Indicates the number of virus scanner related events that were captured during the last measurement period.</td><td>Number</td><td>The events may be generated due to the detection of virus in the storage system or the loss of connection between the virus scanner and the storage system.

Ideally, the value of this measure should be zero. A high value is an indication of performance degradation of the storage system.

The detailed diagnosis capability, if enabled provides you with a more detailed information about the virus scanner related events that were captured by this measure.</td></tr>
</table>

1.9.2 System Status Test

This test reports the overall health of the NetApp storage system and the current state of the AutoSupport feature.

Purpose	Reports the overall health of the NetApp storage system and the current state of the AutoSupport feature
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the Cisco Router. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **SNMPPORT** - The port number through which the Cisco router exposes its SNMP MIB. The default value is 161. 5. **SNMPVERSION** – By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the **SNMPVERSION** list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list. 6. **SNMPCOMMUNITY** – The SNMP community name that the test uses to communicate with the Cisco router. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the **SNMPVERSION** chosen is **v3**, then this parameter will not appear. 7. **USERNAME** – This parameter appears only when **v3** is selected as the **SNMPVERSION**. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against the **USERNAME** parameter. 8. **AUTHPASS** – Specify the password that corresponds to the above-mentioned **USERNAME**. This parameter once again appears only if the **SNMPVERSION** selected is **v3**. 9. **CONFIRM PASSWORD** – Confirm the **AUTHPASS** by retyping it here. 10. **AUTHTYPE** – This parameter too appears only if **v3** is selected as the **SNMPVERSION**. From the **AUTHTYPE** list box, choose the authentication algorithm using which SNMP v3 converts the specified **USERNAME** and **PASSWORD** into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: ➢ **MD5** – Message Digest Algorithm ➢ **SHA** – Secure Hash Algorithm 11. **ENCRYPTFLAG** – This flag appears only when **v3** is selected as the **SNMPVERSION**. By default, the eG agent does not encrypt SNMP requests. Accordingly, the **ENCRYPTFLAG** is set to **NO** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **YES** option. 12. **ENCRYPTTYPE** – If the **ENCRYPTFLAG** is set to **YES**, then you will have to mention the encryption type by selecting an option from the **ENCRYPTTYPE** list. SNMP v3 supports the following encryption types: ➢ **DES** – Data Encryption Standard ➢ **AES** – Advanced Encryption Standard 13. **ENCRYPTPASSWORD** – Specify the encryption password here. 14. **CONFIRM PASSWORD** – Confirm the encryption password by retyping it here.

<table>
<tr><td></td><td colspan="3">15. TIMEOUT - Specify the duration (in seconds) within which the SNMP query executed by this test should time out in the TIMEOUT text box. The default is 10 seconds.

16. DATA OVER TCP –By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the equalizer over TCP (and not UDP). For this, set the DATA OVER TCP flag to Yes. By default, this flag is set to No.

17. DD FREQUENCY - The DD FREQUENCY refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 2:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.

18. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG system 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 against DETAILED DIAGNOSIS. 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 bad and normal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored.</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

<table>
<tr>
<td>test</td>
<td>Global status:
Indicates the overall status of the NetApp Unified Storage system.</td>
<td></td>
<td>This measure reports the following values to indicate the overall status of the NetApp Unified Storage System:
• other
• unknown
• ok
• nonCritical
• critical
• nonRecoverable
The numeric values that correspond to the above-mentioned measure values are as follows:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>other</td><td>1</td></tr><tr><td>unknown</td><td>2</td></tr><tr><td>ok</td><td>3</td></tr><tr><td>nonCritical</td><td>4</td></tr><tr><td>critical</td><td>5</td></tr><tr><td>nonRecoverable</td><td>6</td></tr></table>Note:
By default, this measure reports the above-mentioned Measure Values while indicating the overall status of the NetApp Unified Storage system. However, in the graph of this measure, will be represented using the corresponding numeric equivalents i.e., 1 to 6.
The detailed diagnosis of this measure will provide a brief message stating the reason for the state mentioned in the table above.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Auto support status:
Indicates the status of the AutoSupport feature in this NetApp Unified Storage system.</td>
<td></td>
<td>Autosupport is a feature available in Data Ontap to monitor the Storage/Filer for any potential system problems and alerts. AutoSupport generates alert in one of the following situations

When events occur on the storage system that require corrective action from the system administrator or NetApp technical support
When the storage system reboots
When you initiate a test message using the autosupport.doit option
Once a week, early Sunday morning, at approximately midnight

This measure reports the following values to indicate the status of the AutoSupport feature of the NetApp Unified Storage system:

ok
smtpFailure
postFailure
smtpPostFailure
unknown

The numeric values that correspond to the above-mentioned measure values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>ok</td><td>1</td></tr>
<tr><td>smtpFailure</td><td>2</td></tr>
<tr><td>postFailure</td><td>3</td></tr>
<tr><td>smtpPostFailure</td><td>4</td></tr>
<tr><td>unknown</td><td>5</td></tr>
</table>
</td>
</tr>
</table>

			Note: By default, this measure reports the above-mentioned **Measure Value**s while indicating the overall status of the NetApp Unified Storage system. However, in the graph of this measure, states will be represented using the corresponding numeric equivalents i.e., *1 to 5*. The detailed diagnosis of this measure will provide a brief explanation regarding the various AutoSupport states.

The detailed diagnosis of the *Global status* measure displays a brief message describing the current state of the NetApp storage device and what caused it to switch to that state.

Figure 11: The detailed diagnosis of the Global status measure

1.9.3 NetApp Fiber Channel Adapters Test

This test instantly detects changes in the overall health, state/mode of the Host Bus Adapter (HBA), and immediately notifies administrators of the errors/problem conditions experienced by the HBA. Additionally the login and logout details through the HBA can also be monitored using this test.

Purpose	Instantly detects changes in the overall health, state/mode of the Host Bus Adapter (HBA), and immediately notifies administrators of the errors/problem conditions experienced by the HBA
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each Host Bus Adapter (HBA) of the NetApp storage system being monitored</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

<table>
<tr>
<td>test</td>
<td>State:
Indicates the current state of this Host Bus Adapter.</td>
<td></td>
<td>The values reported by this measure and their numeric equivalents are available in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Startup</td><td>1</td></tr>
<tr><td>Uninitialized</td><td>2</td></tr>
<tr><td>Initializing Firmware</td><td>3</td></tr>
<tr><td>Link Not Connected</td><td>4</td></tr>
<tr><td>Waiting For Link Up</td><td>5</td></tr>
<tr><td>Online</td><td>6</td></tr>
<tr><td>Link Disconnected</td><td>7</td></tr>
<tr><td>Resetting</td><td>8</td></tr>
<tr><td>Offline</td><td>9</td></tr>
<tr><td>Offlined by user/system</td><td>10</td></tr>
</table>
Note:

This measure reports the Measure Values listed in the table above to indicate the current state of this HBA. However, in the graph of this measure, the state is indicated using only the Numeric Values listed in the above table.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Is adapter on standby?:
Indicates whether/not this HBA is in standby mode.</td>
<td>Number</td>
<td>The values reported by this measure and their numeric equivalents are available in the table below:

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

Note:

This measure reports the Measure Values listed in the table above to indicate whether this HBA is in standby mode. However, the graph of this measure will be represented using only the Numeric Values listed in the above table.</td>
</tr>
<tr>
<td></td>
<td>Queue depth:
Indicates the number of I/O operations that can be run simultaneously i.e., in parallel in the ports of this HBA.</td>
<td>Number</td>
<td>A low value is desired for this measure. A high value is characterized by poor response time for the I/O operations and a queue full message. Too many I/O operations may fill the port queue to the maximum leading to a queue full message to the HBA. When a high value occurs, the host operating system may throttle the I/Os to a minimum or otherwise the I/Os may fail leading to performance bottleneck of the storage system.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Is SFP optical transceiver valid?:
Indicates whether/not the configuration of this small form-factor pluggable (SFP) optical transceiver valid.</td>
<td></td>
<td>The SFP optical transceiver serves as the interface to a fiber optic or copper networking cable. Installed SFPs that are not supported for the configuration become invalid and result in connection issues.

The values reported by this measure and their numeric equivalents are available in the table below:

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

Note:

This measure reports the Measure Values listed in the table above to indicate whether the configuration of this SFP optical transceiver is valid. However, in the graph of this measure, the validity of the SFP optical transceiver will be represented using only the Numeric Values listed in the above table.</td>
</tr>
</table>

	Selective LIP resets: Indicates the number of times the selective Reset LIP (Loop Initialization Primitive) occurred during the last measurement period.	Number	Loop Initialization is an essential process for allowing new devices onto the loop, assigning Aribrated Loop Physical Addresses (AL_PAs), providing notification of topology changes, and recovering from loop failure. Following loop initilaization, the loop enters a stable monitoring mode and resumes normal activity. Depending on the number of normal ports (NL_Ports) attached to the loop, an entire loop initialization may take a few milliseconds. A loop initialization can be triggered by a number of causes, the most common being the introduction of a new device. The new device could actually be a former device that has been powered on, or an active device that has been moved from one hub port to another. A number of ordered sets have been defined to cover the various conditions that an NL_port may sense as it launches the initialization process. These ordered sets, called **loop initialization primitive** sequences, are referred to collectively as LIPs. An NL_Port issues atleast 12 LIPs to start loop initialization. During loop initialization, each downstream device that are part of the loop receives the LIP stream and enters a state known as **Open-init**, which suspends any current operations and prepares the device for the loop initialization procedure. The LIPs are forwarded along the loop until all NL_ports, including the originator of the loop, are in **Open-init** state. At this point, a **temporary loop master** is selected for conducting the rest of the initialization procedure. The first task of the temporary loop master is to issue a series of four frames that will allow each device on the loop to select a *unique* AL_PA. A LIP reset is used to perform a vendor specific reset at the loop port specified by this AL-PA value. These LIP resets are used to temporarily cure connectivity issues. Prolonged resets should be noted and the underlying actual connectivity issues should be resolved.

	Total CRC errors: Indicates the number of Cyclic Redundancy Check (CRC) errors that occurred during data trafficking in the FC ports of this HBA, during the last measurement period.	Number	CRC or Cyclic Redundancy Check is a process that helps in identifying any errors that might occur during the data transmission process. Data is usually transmitted in small blocks, and a CRC value is assigned to each block and transmitted along with it. This CRC value is verified at the destination to ensure that it matches the CRC value transmitted from the source. A CRC error occurs when the two values (source and destination) do not match and the test fails. The main benefit of CRC is that it helps you ensure that data you have received or downloaded is not damaged or corrupt. By comparing the value of this measure across all FC ports, you can accurately identify the most error-prone FC ports.
	Discarded frames: Indicates the number of frames that were discarded during the last measurement period.	Number	Ideally, the value of this measure should be 0.
	Initiators connected: Indicates the number of initiators that were connected to this HBA during the last measurement period.	Number	
	Link breaks: Indicates the number of times the link failed (broke) during the last measurement period.	Number	Ideally, the value of this measure should be 0.
	Spurious interrupts: Indicates the number of spurious signals in the cable during the last measurement period.	Number	
	Protocol errors: Indicates the number of Fiber Channel Protocol (FCP) errors that occurred during the last measurement period.	Number	Ideally, the value of this measure should be 0.

	Dropped SCSI requests: Indicates the number of SCSI requests that were dropped since the last measurement period.	Number	
	Total logins: Indicates the total number of logins during the last measurement period.	Number	
	Logouts: Indicates the total number of logouts during the last measurement period.	Number	

1.9.4 NetApp Initiator Config Mismatches

Initiator groups (igroups) are tables of host identifiers (FCP, WWPNs, or iSCSI node names) that are used to control hosts' access to LUNs.

igroups specify which initiators have access to which LUNs. igroups can be created either before or after LUNs are created, but they must be created before a LUN is mapped to an igroup. Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. However, a LUN can not be mapped to multiple igroups that have the same initiator.

An initiator cannot be a member of igroups of differing ostypes.

Using this test, you can monitor the initiators of each igroup and determine the following:

- The number of initiators with ALUA setting mismatch
- The number of initiators with the OS type mismatch
- How many initiators are actually differing from the actual VSA setting? and
- Which are the initiators that are mapped to the LUNs with non unique ids?

Purpose	Monitors the initiators of each igroup and helps determine the following: • The number of initiators with ALUA setting mismatch • The number of initiators with the OS type mismatch • How many initiators are actually differing from the actual VSA setting? and • Which are the initiators that are mapped to the LUNs with non unique ids?
Target of the	A NetApp Unified Storage

test	
Agent deploying the test	An external/remote agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each iGroup on the NetApp storage system being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Initiators with mismatching ALUA setting:
Indicates the number of initiators in this igroup with an ALUA setting mismatch.</td><td>Number</td><td>If the ALUA setting does not match between the local and partner systems, it would affect the host multi-path software's ability to distinguish between primary and secondary paths. This could lead to incorrect system behavior. The containing igroups could be within the local system or between local and partner systems.
The detailed diagnosis of this measure reveals the name of the initiator, the initiator group to which the initiator belongs and the ALUA settings of the affected initiators.</td></tr>
</table>

	Initiators with mismatching OS type: Indicates the number of initiators in an initiator group with an operating system mismatch.	Number	An initiator cannot be a member of initiator groups of differing OS types i.e., the initiator can be a member of igroups that are of the same OS type. The detailed diagnosis of this measure reveals the name of the initiator, the igroup of the initiator and the OS type of the affected initiator.
	Initiators with mismatching VSA setting: Indicates the number of initiators in an initiator group with a differing VSA setting.	Number	In order to avoid unexpected performance related issues in the storage system, an initiator can be a member of initiator groups with the same VSA setting only. The detailed diagnosis of this measure reveals the name of the initiator, the initiator group to which the initiator belongs and the VSA setting of the affected initiator.
	Initiators with conflicting LUN mapping: Indicates the number of initiators that are mapped to the LUNs with non-unique LUN ids.	Number	Only one LUN in the cluster can be mapped to an initiator at a given LUN-id. Certain conflicts may arise if a LUN on each filer is mapped to the same initiator at the same LUN-id. These conflicts need to be resolved before a filer can be upgraded to run in the *'single_image'* fcp cfmode. The conflicts can be resolved by unmapping one LUN and remapping it to an unused LUN-id. The detailed diagnosis of this measure reveals the name of the initiator and the LUN id of the mapped LUN for the affected initiators.

1.9.5 NetApp Syslog Test

This test queries the target syslog file log for specific errors and warning messages in the and reports the number of such messages found.

Purpose	Queries the target syslog file log for specific errors and warning messages in the and reports the number of such messages found
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

10. **VFILERNAME** – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the **VFILERNAME** text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of *none* is displayed in this text box.

11. **TIMEOUT** - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.

12. **SYSLOG FULL PATH** - Specify the full path to the most recent syslog file in the **SYSLOG FULL PATH** text box. The default value displayed in this text box is */vol/vol0/etc/messages*

13. **SEARCHPATTERN** - Enter the specific patterns of alerts to be monitored. The pattern should be in the following format: *<PatternName>:<Pattern>*, where *<PatternName>* is the pattern name that will be displayed in the monitor interface and *<Pattern>* is the pattern that you need to search for in the log file. The *<Pattern>* can either be a text string or an expression of the form **expr**.

    For example, say you specify **Info_Msgs:info** in the **SEARCHPATTERN** text box. This indicates that **"Info_Msgs"** is the pattern name to be displayed in the monitor interface. The value **"info"** indicates that the test will monitor only those lines in the syslog which contain the string **"info"**. Similarly, if your pattern specification reads: **Error_Msgs:vol*error**, then it means that the pattern name is **Error_Msgs** and the test will monitor only those lines in the syslog which begin with the string **vol** and end with the string **error**.

    Multiple search patterns can be specified as a comma-separated list. For example: **Info_Msgs:info,Error_Msgs:vol*error**

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

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

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

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

Outputs of the test	One set of results for every **SEARCHPATTERN** configured		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages: Indicates the number of errors or warning messages of the configured **SEARCHPATTERN** that were found in the specified syslog file during the last measurement period.	Number	The detailed diagnosis of this measure lists all the individual messages.

1.9.6 NetApp System Performance Test

Using this test, the overall performance of the NetAapp Unified Storage system can be measured with key measures such as the following:

- The average latency for all the operations performed on the system;
- The disk throughput of the system;
- The network throughput of the system and
- The rate at which read and write operations were performed on the system

Purpose	Using this test, the overall performance of the NetAapp Unified Storage system can be measured with key measures such as the following: • The average latency for all the operations performed on the system; • The disk throughput of the system; • The network throughput of the system and • The rate at which read and write operations were performed on the system
Target of the test	A NetApp Unified Storage
Agent deploying the test	An external/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The IP address of the storage controller. 3. **PORT** - Specify the port at which the specified **HOST** listens in the **PORT** text box. By default, this is *NULL*. 4. **USER** – Here, specify the name of the user who possesses the following privileges: *login-http-admin,api-aggr-check-spare-low,api-aggr-list-info,api-aggr-mediascrub-list-info,api-aggr-scrub-list-info,api-cifs-status,api-clone-list-status,api-disk-list-info,api-fcp-adapter-list-info,api-fcp-adapter-stats-list-info,api-fcp-service-status,api-file-get-file-info,api-file-read-file,api-iscsi-connection-list-info,api-iscsi-initiator-list-info,api-iscsi-service-status,api-iscsi-session-list-info,api-iscsi-stats-list-info,api-lun-config-check-alua-conflicts-info,api-lun-config-check-cfmode-info,api-lun-config-check-info,api-lun-config-check-single-image-info,api-lun-list-info,api-nfs-status,api-perf-object-get-instances-iter*,api-perf-object-instance-list-info,api-quota-report-iter*,api-snapshot-list-info,api-vfiler-list-info,api-volume-list-info-iter**. If such a user does not pre-exist, then, you can create a special user for this purpose using the steps detailed in Section 1.1.4 of this document. 5. **PASSWORD** - Specify the password that corresponds to the above-mentioned **USER**. 6. **CONFIRM PASSWORD** - Confirm the **PASSWORD** by retyping it here. 7. **AUTHENTICATION MECHANISM** – In order to collect metrics from the NetApp Unified Storage system, the eG agent connects to the ONTAP management APIs over HTTP or HTTPS. By default, this connection is authenticated using the *LOGIN_PASSWORD* authentication mechanism. This is why, *LOGIN_PASSWORD* is displayed as the default **AUTHENTICATION MECHANISM**. 8. **USE SSL** - Set the **USE SSL** flag to **Yes**, if SSL (Secured Socket Layer) is to be used to connect to the NetApp Unified Storage System, and **No** if it is not. 9. **API PORT** - By default, in most environments, NetApp Unified Storage system listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled) only. This implies that while monitoring the NetApp Unified Storage system, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of the NetApp Unified Storage system - i.e., if the NetApp Unified Storage system is not SSL-enabled (i.e., if the **USE SSL** flag above is set to **No**), then the eG agent connects to the NetApp Unified Storage system using port 80 by default, and if the NetApp Unified Storage system is SSL-enabled (i.e., if the **USE SSL** flag is set to **Yes**), then the agent-NetApp Unified Storage system communication occurs via port 443 by default. Accordingly, the **API PORT** parameter is set to *default* by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **API PORT** parameter, you can specify the exact port at which the NetApp Unified Storage system in your environment listens, so that the eG agent communicates with that port for collecting metrics from the NetApp Unified Storage system.

<table>
<tr><td></td><td colspan="3">10. VFILERNAME – A vFiler is a virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. If the NetApp Unified Storage system is partitioned to accommodate a set of vfilers, specify the name of the vfiler that you wish to monitor in the VFILERNAME text box. In some environments, the NetApp Unified Storage system may not be partitioned at all. In such a case, the NetApp Unified Storage system is monitored as a single vFiler and hence the default value of none is displayed in this text box.
11. TIMEOUT - Specify the duration (in seconds) beyond which the test will timeout if no response is received from the device. The default is 120 seconds.
12. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capabili7ty. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the NetApp storage system being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Sys avg latency:
Indicates the time taken for performing all the operations in the NetApp Unified Storage system.</td><td>Milliseconds</td><td>g. A low value is desired for this measure. The responsiveness of the NetApp Unified storage system can be measured with this measure. The detailed diagnosis of this measure indicates the exact cause for the sudden slowdown in responsiveness of the system.</td></tr>
<tr><td>Disk data read:
Indicates the rate at which data is read from all the disks of the NetApp Unified Storage system.</td><td>KB/Sec</td><td>A high value is desired for this measure.</td></tr>
</table>

	Disk data written: Indicates the rate at which data is written to all the disks of the NetApp Unified Storage system.	KB/Sec	An abnormally high value indicates that the disk is taking too long to store the data which may be due to the disk being full or a processing bottleneck or a network slowdown.
	Data received: Indicates the rate at which data is received through the network to the NetApp Unified Storage system.	KB/Sec	A high value is desired for this measure. A low value indicates a processing bottleneck or a network slowdown.
	Data sent: Indicates the rate at which data is sent through the network to the NetApp Unified Storage system.	KB/Sec	
	HTTP operations: Indicates the rate at which HTTP operations were performed on the NetApp Unified Storage system.	Ops/Sec	HTTP operations include management operations. Usually log processing is done using the management web interface of the NetApp Unified Storage system. The value of this measure will increase if there is excessive log processing activity in the management web interface.
	Read operations: Indicates the rate at which read operations were performed on the NetApp Unified Storage system.	Ops/Sec	
	Write operations: Indicates the rate at which write operations were performed to the NetApp Unified Storage system.	Ops/Sec	
	Raid read latency: Indicates the average time taken for all read operations from WAFL to the RAID of the NetApp Unified Storage system.	Milliseconds	Ideally, the value of this measure should be very low.

Chapter 2

Conclusion

This document has clearly explained how eG Enterprise monitors IBM pSeries servers and the AIX LPARs configured on them. We can thus conclude that eG Enterprise, with its ability to provide in-depth insight into the performance of AIX LPAR infrastructures, is the ideal solution for monitoring such environments. For more information on eG Enterprise, please visit our web site at www.eginnovations.com or write to us at sales@eginnovations.com.