



# ***Monitoring IBM Storwize v7000***

***eG Enterprise v6***

**Restricted Rights Legend**

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

**Trademarks**

Microsoft Windows, Windows NT, Windows 2000, Windows 2003 and Windows 2008 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

**Copyright**

©2014 eG Innovations Inc. All rights reserved.

# Table of Contents

<b>MONITORING IBM STORWIZE V7000 STORAGE SYSTEM .....</b>	<b>1</b>
1.1 The IBM v7000 VDisks Layer .....	3
1.1.1 v7000 VDisks Test .....	4
1.1.2 v7000 VDisk Status Test .....	6
1.2 The IBM v7000 Nodes Layer .....	9
1.2.1 v7000 Port Traffic Test .....	10
1.2.2 v7000 Node Traffic Test .....	14
1.2.3 v7000 CPU usage Test .....	16
1.3 The IBM v7000 Cache Layer .....	17
1.3.1 v7000 VDisk Cache Test .....	18
1.4 The IBM v7000 MDisks Layer .....	21
1.4.1 v7000 MDisks Test .....	22
1.4.2 v7000 MDisk Status Test .....	25
1.5 The IBM v7000 Drives Layer .....	31
1.5.1 v7000 Drives Test .....	31
1.5.2 v7000 Drive Status Test .....	33
1.6 The Hardware Layer .....	35
1.6.1 v7000 Traps Test .....	36
1.6.2 v7000 Enclosures Test .....	38
1.6.3 v7000 Enclosure Slots Test .....	41
1.6.4 v7000 Enclosure PSU Status Test .....	43
1.6.5 v7000 Enclosure Canisters Test .....	45
1.6.6 v7000 Enclosure Battery Status Test .....	46
<b>CONCLUSION .....</b>	<b>51</b>

# Table of Figures

Figure 1.1: The architecture of IBM Storwize v7000 .....	1
Figure 1.2: The layer model of IBM Storwize v7000 storage system .....	2
Figure 1.3: The tests mapped to the IBM v7000 VDisks layer .....	4
Figure 1.4: The tests mapped to the IBM v7000 Nodes layer .....	10
Figure 1.5: The tests mapped to the IBM v7000 Cache layer .....	18
Figure 1.6: The tests mapped to the IBM v7000 MDisk layer .....	22
Figure 1.7: The tests mapped to the IBM v7000 Drives layer .....	31
Figure 1.8: The tests mapped to the Hardware layer .....	36

# Monitoring IBM Storwize v7000 storage system

The IBM Storwize V7000 solution provides a modular storage system that includes the capability to virtualize external SAN-attached storage and its own internal storage. The IBM Storwize V7000 solution is built upon the IBM SAN Volume Controller technology base and uses technology from the IBM System Storage DS8000 family. The architecture (see Figure 1.1) of the SAP HANA consists of four different components.

An IBM Storwize V7000 system provides a number of configuration options that are aimed at simplifying the implementation process. It also provides automated wizards, called *Directed Maintenance Procedures (DMP)*, to assist in resolving any events that might occur. An IBM Storwize V7000 system is a clustered, scalable, and midrange storage system, and an external virtualization device.




Figure 1.1: The architecture of IBM Storwize v7000

## Monitoring IBM Storwize v7000 storage system

The IBM Storwize V7000 solution consists of a one or two control enclosures and, optionally, up to 18 expansion enclosures (and supports the intermixing of the different expansion enclosures). Within each enclosure are two canisters. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

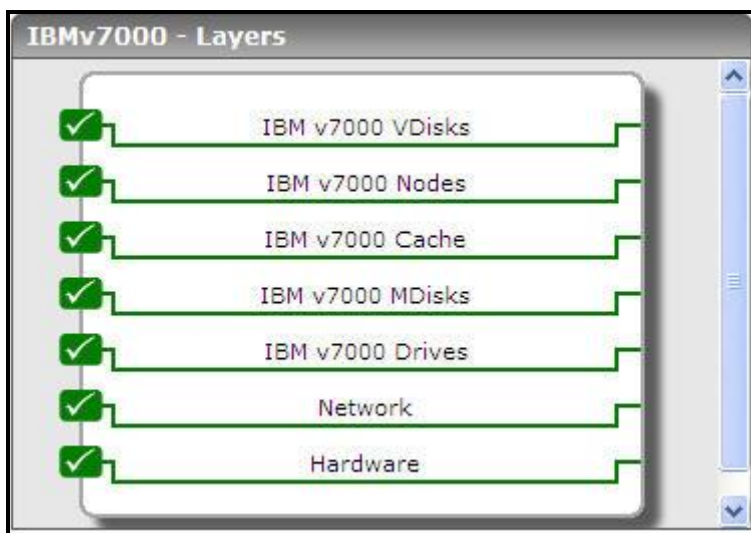


Figure 1.2: The layer model of IBM Storwize v7000 storage system

The eG Enterprise Suite provides you with extensive monitoring capabilities for the IBM Storwize v7000 storage system. A single eG agent is capable of monitoring the IBM Storwize v7000 storage system. Every layer of Figure 1.1 is mapped to a wide variety of tests that a single agent executes and extracts loads of performance metrics from the IBM Storwize v7000 storage system.

The metrics thus collected would be useful to figure out accurate answers to the following performance queries:

## Monitoring IBM Storwize v7000 storage system

- Are the VDisks processing I/O requests quickly, or has any processing bottleneck been noticed in any VDisk?
- Which VDisk is the slowest in I/O processing?
- Is any VDisk in a degraded state currently? If so, which one?
- How is the VDisk cache? Is the cache of any VDisk corrupt?
- Is any port on a node overloaded? Which one is it and how well is it processing read-write requests?
- Is any FC port experiencing errors in operation?
- Are all nodes in a cluster processing I/O requests at optimal speed, or is there any node that is experiencing significant latencies during processing?
- Which node canister is hogging the CPU?
- How are the VDisk caches utilized? Which VDisk is under-utilizing its cache and why? Is it because the cache is not adequately sized?
- Which VDisk cache is accumulating data and why?
- Which MDisk is the slowest when reading from and/or writing to the disk?
- Is any MDisk offline or degraded?
- The RAID status of which MDisk is degraded?
- Is I/O load uniformly balanced across all drives, or is any drive overloaded with requests?
- Is any drive processing I/O at a rate slower than the other drives? If so, which drive is it and why?
- Which drive is currently in a degraded state and what is its capacity?
- Which enclosure is degraded? What is its composition?
- Which enclosure slots have drives and which ones don't?
- Does any enclosure have power supplies in a degraded state?
- Are all canisters online and operating normally?
- How much charge is left with enclosure batteries? Which battery is fast-approaching its end-of-life?
- Does any battery require reconditioning?

The **Network** layer of an *IBM Storwize v7000* storage system model is similar to that of a *Windows Generic* server model. Since the tests associated with this layer have been dealt with in the *Monitoring Unix and Windows Servers* document, the below mentioned sections focus only on the layers pertaining to the IBM Storwize v7000 storage system alone.

### 1.1 The IBM v7000 VDisks Layer

The **IBM v7000 VDisks** layer auto discovers the VDisks of the storage system and reports how well the I/O operations were performed on each VDisk. This layer also tracks the status of each VDisk. The following sections provide more information on the tests and measures reported by them.



Figure 1.3: The tests mapped to the IBM v7000 VDisks layer

### 1.1.1 v7000 VDisks Test

A volume or a VDisk is a logical disk that the clustered system presents to a host connected over a Fibre Channel or Ethernet network. These VDisks enable administrators to more efficiently manage resources. However, if a vDisk experiences a processing bottleneck, it will not be able to process I/O requests quickly, thereby causing prolonged delays in data access for users. Similarly, a vdisk that is overloaded will not be able to perform at peak capacity, thus affecting the user experience with the storage system. Administrators hence have to continuously track the load on and the I/O processing speed of each of the vdisks, so that potential overload conditions and probable processing delays can be detected proactively and pre-emptively treated. The **v7000 VDisks** test helps administrators with this. This test monitors the level of traffic on each vdisk on an IBM v7000 storage system, and helps isolate irregularities in load balancing across the vdisks. Alongside, the test also helps identify which vdisk is experiencing processing bottlenecks (if any), so that the bottleneck is cleared before users complain of slowdowns.

<b>Purpose</b>	Monitors the level of traffic on each vdisk on an IBM v7000 storage system, and helps isolate irregularities in load balancing across the vdisks. Alongside, the test also helps identify which vdisk is experiencing processing bottlenecks (if any), so that the bottleneck is cleared before users complain of slowdowns.
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>
<b>Outputs of the test</b>	One set of results for every VDisk being monitored



Measurements made by the test	Measurement	Measurement Unit	Interpretation
	<b>Read operations:</b> Indicates the rate at which read operations were performed by this VDisk.	Reads/Sec	A consistent rise in the value of these measures indicates a high level of I/O activity on the VDisk. Comparing the value of each of these measures across VDIsks helps you in determining whether/not the I/O load is uniformly distributed across the VDIsks or one/more VDIsks are handling more load than the rest. This comparative analysis could shed light on irregularities in load balancing across the VDIsks.
	<b>Write operations:</b> Indicates the rate at which write operations were performed by this VDisk.	Writes/Sec	
	<b>Read data:</b> Indicates the rate at which data blocks were read from this VDisk.	MB/Sec	A consistent drop in this value could indicate a gradual dip in the rate at which read requests are processed by the VDisk. Compare the value of this measure across VDIsks to identify those VDIsks that are reading slowly.
	<b>Write data:</b> Indicates the rate at which the data blocks were written to this VDisk of this node.	MB/Sec	A consistent drop in this value could indicate a gradual dip in the rate at which write requests are processed by the VDisk. Compare the value of this measure across VDIsks to identify those VDIsks that are writing data slowly.
	<b>Average read latency:</b> Indicates the average time taken by this VDisk to respond to read requests.	Millisec	Comparing the value of these measures across VDIsks helps you in identifying the VDIsks that are the least responsive to I/O requests and when they are the slowest – when reading? or when writing?.
	<b>Average write latency:</b> Indicates the average time taken by this VDisk to respond to write requests.	Millisec	
	<b>Average transfer latency:</b> Indicates the average time taken by this VDisk to transfer requests during the last measurement period.	Microseconds	
	<b>Maximum read response time:</b> Indicates the maximum time taken by this VDisk to respond to a read request during the last measurement period.	Millisec	Compare the value of this measure across the VDIsks to identify the VDisk that is extremely slow in responding to read requests.

	<b>Maximum write response time:</b> Indicates the maximum time taken by this VDisk to respond to a write request during the last measurement period.	Millisec	Compare the value of this measure across the VDisks to identify the VDisk that is extremely slow in responding to write requests.
--	---	----------	---

## 1.1.2 v7000 VDisk Status Test

A volume or a VDisk is a logical disk that the clustered system presents to a host connected over a Fibre Channel or Ethernet network. These VDisks enable administrators to more efficiently manage resources. If any of these VDisks is in an offline or degraded state, it can cause write data that has been modified to be pinned in the SAN Volume Controller cache. This prevents volume failover and causes a loss of input/output (I/O) access. I/O loss can also occur if the cache of a VDisk is corrupt. To prevent or at least minimize such losses, administrators need to swiftly detect the abnormal state of the VDisk and/or its cache and instantly initiate measures to remove the abnormality, so that normalcy is restored soon. This is where the **v7000 VDisk Status** test helps. This test reports the current status of each VDisk of the IBM Storwize v7000 storage system and also reports the cache state of every VDisk, so that the abnormal state of the VDisk and/or the cache can be promptly detected and speedily resolved.

<b>Purpose</b>	Reports the current status of each VDisk of the IBM Storwize v7000 storage system and also reports the cache state of every VDisk, so that the abnormal state of the VDisk and/or the cache can be promptly detected and speedily resolved
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li><b>TEST PERIOD</b> - How often should the test be executed</li> <li><b>HOST</b> - The host for which the test is to be configured.</li> <li><b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li><b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> <li><b>DETAILED DIAGNOSIS</b> To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>The eG manager license should allow the detailed diagnosis capability</li> <li>Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul> </li> </ol>

<b>Outputs of the test</b>	One set of results for each VDisk being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Status:</b> Indicates the current status of this VDisk.		<p>This measure can take any of the follo:</p> <ul style="list-style-type: none"> <li>➤ offline</li> <li>➤ online</li> <li>➤ degraded</li> </ul> <p>A VDisk is offline and unavailable if one of the following takes place:</p> <ul style="list-style-type: none"> <li>➤ Both nodes in the I/O group are missing.</li> <li>➤ None of the nodes in the I/O group that are present can access the VDisk.</li> <li>➤ All synchronized copies for this VDisk are in storage pools that are offline.</li> <li>➤ The VDisk is formatting.</li> </ul> <p>A VDisk is reported as degraded if any of the following occurs:</p> <ul style="list-style-type: none"> <li>➤ One of the nodes in the I/O group is missing.</li> <li>➤ One of the nodes in the I/O group cannot access all the MDisk in the storage pool that the VDisk spans. In this case MDisk are shown as degraded and the fix procedures for MDisk should be followed to resolve the problem.</li> <li>➤ The fast write cache pins data for one or more VDisks in the I/O group and is unable to perform a failback until the situation is resolved. An error log indicating that the cache has pinned data is displayed. Follow the fix procedures for this error log to resolve the problem. The most common causes of pinned data are the following:</li> </ul>

- One or more VDisks in an I/O group is offline due to an asymmetric failure and has pinned data in the cache. Asymmetric failures can occur because of Storwize® V7000 fabric faults or misconfiguration, back-end controller faults or misconfiguration or because repeated errors has led to the system excluding access to a MDisk through one or more nodes.
- One or more VDisks in an I/O group is offline due to a problem with a FlashCopy® mapping.

The numeric values that correspond to the above-mentioned measure values are as follows:

Measure Value	Numeric Value
offline	0
online	1
degraded	2

**Note:**

By default, this measure reports the above-mentioned **Measure Values** while indicating the status of a VDisk. However, in the graph of this measure, VDisk status is represented using the corresponding numeric equivalents only

The detailed diagnosis of this measure reveals the VDisk ID, the VDisk IO GROUP ID, the VDisk IO GROUP NAME, MDISK ID, MDISK NAME, the VDisk TYPE and the FAST WRITE STATUS of the VDisk. From the detailed diagnostics, you can glean the name of the I/O group to which the VDisk belongs and the MDisks (i.e., the management disks) in the storage pool that the VDisk spans. In the event that the VDisk is offline or degraded, you can use the I/O group and MDisk ID to investigate the reason for the degradation or unavailability of the VDisk – is it because the I/O group has a missing node? or is it because the MDisk is degraded?

	<b>Capacity:</b> Indicates the total capacity of this VDisk.	TB											
	<b>Fast write status:</b> Indicates the cache status of this VDisk.		<p>This measure reports any of the values listed below:</p> <ul style="list-style-type: none"><li>➤ corrupt</li><li>➤ repairing</li><li>➤ empty</li><li>➤ not empty</li></ul> <p>A cache state of <b>corrupt</b> indicates that the VDisk requires recovery by using one of the <i>recovervdisk</i> commands. A cache state of <b>repairing</b> indicates that repairs initiated by a <i>recovervdisk</i> command are in progress.</p> <p>The numeric values that correspond to each of the measure values listed above are mentioned in the table below:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>corrupt</td><td>1</td></tr><tr><td>repairing</td><td>2</td></tr><tr><td>empty</td><td>3</td></tr><tr><td>not empty</td><td>4</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the cache status of a VDisk. However, in the graph of this measure, cache state is represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	corrupt	1	repairing	2	empty	3	not empty	4
Measure Value	Numeric Value												
corrupt	1												
repairing	2												
empty	3												
not empty	4												

## 1.2 The IBM v7000 Nodes Layer

The tests associated with this layer (see Figure 1.4) monitor the nodes of the IBM Storwize v7000 storage system and reports how long the node canister was busy, the percentage of time the CPU was busy, the traffic on each node, the traffic through each port of the nodes, etc.



Figure 1.4: The tests mapped to the IBM v7000 Nodes layer

### 1.2.1 v7000 Port Traffic Test

IBM Storwize V7000 storage system can have two or four hardware components called nodes or node canisters that provide the virtualization of internal and external volumes, and cache and copy services (Remote Copy) functions. Storwize V7000 type 100 node canisters contain four ports for Fibre Channel connection and two ports for 1 Gbps Ethernet connection. Type 300 node canisters contain four ports for Fibre Channel connection, two ports for 1 Gbps Ethernet connection, and an HBA that provides two additional ports for 10 Gbps Ethernet connection. Each node presents a volume to the SAN through four Fibre Channel ports or two FCoE ports. These ports therefore are the primary handlers of I/O requests from the SAN. I/O load on the ports directly translate into load on the volumes. This is why, administrators need to continuously monitor the data and commands processed by each port, so that overloaded ports can be quickly identified and the load-balancing algorithm fine-tuned accordingly. Moreover, since port-related errors can deny hosts access to the data stored in the SAN, port monitoring is imperative to enable administrators to quickly detect such errors and fix them to ensure the normal functioning of the SAN. All this and more can be achieved using the **v7000 Port Traffic** test. For each port on a node, this test reports the rate at which data and commands are handled by each node and the number and nature of errors/failures encountered by each FC port. This way, administrators can be proactively alerted to potential port overloads and error conditions (with FC ports), and thus enabled to rapidly initiate remedial measures to avoid an impending storage system slowdown.

<b>Purpose</b>	For each port on a node, this test reports the rate at which data and commands are handled by each node and the number and nature of errors/failures encountered by each port
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent

<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li><b>TEST PERIOD</b> - How often should the test be executed</li> <li><b>HOST</b> - The host for which the test is to be configured.</li> <li><b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li><b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>		
<b>Outputs of the test</b>	One set of results for each port available in the node of the IBM Storwize v7000 storage system being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Data transmitted to hosts:</b> Indicates the rate at which data is transmitted to the hosts through this port.	MB/Sec	Compare the value of these measures across the ports to identify the port that is busy transmitting /receiving data to the hosts - thus load-balancing irregularities on the ports can be detected easily.
	<b>Data received from hosts:</b> Indicates the rate at which data is received from the hosts through this port.	MB/Sec	
	<b>Data transmitted to controllers:</b> Indicates the rate at which data is transmitted to the controllers through this port.	MB/Sec	Compare the value of these measures across the ports to identify the port that is busy transmitting /receiving data to the controllers - thus load-balancing irregularities on the ports can be detected easily.
	<b>Data received from controllers:</b> Indicates the rate at which data is received from the controllers through this port.	MB/Sec	
	<b>Commands initiated to controllers:</b> Indicates the rate at which commands are initiated to the controllers through this port.	Commands/Sec	Compare the value of these measures across the ports to identify the port that is busy receiving/initiating commands – this way, load-balancing irregularities on the ports can be detected easily.
	<b>Commands received from hosts:</b> Indicates the rate at which commands are received from the hosts through this port.	Commands/Sec	

	<b>Link failures:</b> Indicates the number of link failures experienced by this FC port.	Number	<p>Ideally, the value of this measure should be zero. A non-zero value indicates that Fibre Channel connectivity with the port was "broken" that many times. This is likely an indicator for a faulty connector or cable. These are also caused when the device connected to the port is restarted, replaced or being serviced when the Fibre Channel cable connected to the port is temporarily disconnected.</p> <p><b>This measure is applicable to FC Ports.</b></p>
	<b>Loss-of-synchronizations:</b> Indicates the number of times this FC port failed to synchronize.	Number	<p>Ideally, the value of this measure should be zero. A non-zero value for this measure indicates that port went into the "loss of synchronization" state, where it encountered continuous Disparity errors.</p> <p>This is likely an indicator for a faulty connector or cable. These are also caused when the device connected to the port is restarted, replaced or being serviced when the Fibre Channel cable connected to the port is temporarily disconnected.</p> <p>If the port is in the "loss of synchronization" state for longer than a specific period, the port will get into the link failure state which could degrade the performance of the Fibre Channel link.</p> <p><b>This measure is applicable only to FC Ports.</b></p>
	<b>Loss-of-signal:</b> Indicates the number of times the signal was lost on this FC port.	Number	<p>Ideally, the value of this measure should be zero. A non-zero value for this measure indicates that the port detected a loss of the electrical or optical signal used to transfer data on the port.</p> <p>This is likely an indicator for a faulty connector or cable. These are also caused when the device connected to the port is restarted, replaced or being serviced when the Fibre Channel cable connected to the port is temporarily disconnected.</p> <p>If the port is in the "loss of signal" state for longer than a specific period, the port will get into the link failure state which could degrade the performance of the Fibre Channel link.</p> <p><b>This measure is applicable only to FC Ports.</b></p>



	<b>Primitive sequence protocol errors:</b> Indicates the number of Primitive Sequence protocol errors that occurred on this FC port.	Number	<p>Ideally, the value of this measure should be zero.</p> <p><b>This measure is applicable only to FC ports.</b></p>
	<b>Invalid transmission words:</b> Indicates the number of invalid words that were transmitted through this FC port.	Number	<p>Transmission Words are either data Transmission Words or control Transmission Words. The first two bits of a Transmission Word are the synchronization header, and are set to either 01h or 10h. The remaining 64 bits of the Transmission Word are the output of a scrambler applied to the Transmission Word body. The Transmission Word body is eight bytes that represent a pair of words and/or Special Functions.</p> <p>An invalid Transmission Word shall be recognized by the receiver when one of the following conditions is detected:</p> <ul style="list-style-type: none"> <li>➤ A code violation, as specified by the 8B/10B transmission code (see 5.2), is detected within a Transmission Word. This is referred to as a code violation condition;</li> <li>➤ A K30.7 special character is detected in any character position of a Transmission Word. This indicates an error condition has been detected at a lower implementation level within the receiver;</li> <li>➤ Any valid special character is detected in the second, third, or fourth character position of a Transmission Word. This is referred to as an invalid special code alignment condition; or</li> <li>➤ A defined Ordered Set is received with improper beginning running disparity.</li> </ul> <p>Ideally, the value of this measure should be zero.</p> <p><b>This measure is applicable only to FC ports.</b></p>

	<p><b>Invalid CRCs:</b></p> <p>Indicates the number of invalid CRCs that occurred on this FC port.</p>	Number	<p>This refers to the number of Fibre Channel frames handled by the port that contains checksum errors.</p> <p>Ideally, the value of this measure should be zero.</p> <p>These are usually recoverable errors and will not degrade system performance unless their occurrence is sustained when the data cannot be relayed after retransmissions.</p> <p><b>This measure is applicable only to FC ports.</b></p>
	<p><b>No Buffer credit timer:</b></p> <p>Indicates the time duration for which this FC port was unable to send frames due to the lack of buffer credit.</p>	Microseconds	<p><b>Buffer credits</b>, also called buffer-to-buffer credits (BBC) are used as a flow control method by Fibre Channel technology and represent the number of frames a port can store.</p> <p>Each time a port transmits a frame that port's BB Credit is decremented by one; for each R_RDY received, that port's BB Credit is incremented by one. Transmission of an R_RDY indicates that the port has processed a frame, freed a receive buffer, and is ready for one more. If the BB Credit is zero, the corresponding node cannot transmit until an R_RDY is received back. A high value for this measure therefore indicates that an R_RDY was not received by the FC port for a long time. This is a cause for concern, as until the R_RDY is received, the FC port will not resume communication.</p> <p>The solution for this problem is to allocate optimal buffer credits to the FC port. The optimal number of buffer credits is determined by the distance (frame delivery time), the processing time at the receiving port, the link signaling rate, and the size of the frames being transmitted. As the link speed increases, the frame delivery time is reduced and the number of buffer credits must be increased to obtain full link utilization, even in a short-distance environment. Smaller frame sizes need more buffer credits</p> <p><b>This measure is applicable only to FC ports.</b></p>

## 1.2.2 v7000 Node Traffic Test

A node is a single processing unit that provides virtualization, cache, and copy services for the cluster. SAN Volume

## Monitoring IBM Storwize v7000 storage system

Controller nodes are deployed in pairs called I/O groups. One node in the cluster is designated the configuration node. At any point in time, only the configuration node can operate as the focal point for configuration and monitoring requests. It is the only node that takes the active cluster IP addresses, and is the only node that receives cluster management requests. You can use one or more of these addresses to access the system through the management GUI or the command-line interface (CLI).

To understand how well the cluster manages I/O requests, you need to monitor the data sent and received by each node in the cluster in response to these requests and the time taken by every node to process the I/O requests. This is exactly what the **v7000 Node Traffic** test does! This test auto-discovers the nodes configured on the IBM Storwize v7000 storage system, and for each node, reports the data sent and received by that node and the latency of the node when receiving and sending data. Besides revealing the load on each node, this will also point you to those nodes that are most likely to experience an overload soon, and those nodes that are sending/receiving data much slower than the other nodes. Based on the results reported by this test, you can investigate the reason why certain nodes delay I/O processing, and initiate measures to eliminate the reasons; in addition, you can also fine-tune the load-balancing algorithm of the cluster to ensure that load is uniformly distributed across nodes.

<b>Purpose</b>	Auto-discovers the nodes configured on the IBM Storwize v7000 storage system, and for each node, reports the data sent and received by that node and the latency of the node when sending and receiving data		
<b>Target of the test</b>	An IBM Storwize v7000 storage system		
<b>Agent deploying the test</b>	A remote agent		
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"><li>1. <b>TEST PERIOD</b> - How often should the test be executed</li><li>2. <b>HOST</b> - The host for which the test is to be configured.</li><li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li><li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li></ol>		
<b>Outputs of the test</b>	One set of results for each node of the IBM Storwize v7000 storage system being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Messages or bulk data received:</b> Indicates the rate at which messages or bulk data is received on this node.	Msg/sec	Compare the value of these measures across the nodes to identify the node that is overloaded i.e., the node that is busy sending/receiving messages. This way, you could identify the irregularities in load balancing across the nodes.
	<b>Messages or bulk data sent:</b> Indicates the rate at which messages or bulk data is sent through this node.	Msg/sec	

	<b>Data received:</b> Indicates the rate at which data is received on this node.	MB/Sec	Compare the value of these measures across the nodes to identify the node that is overloaded - i.e., the node that is busy sending/receiving messages. This way, you could identify the irregularities in load balancing across the nodes.
	<b>Data sent:</b> Indicates the rate at which data is sent through this node.	MB/Sec	
	<b>Average receive latency including inbound queue:</b> Indicates the average time taken by this node to receive messages, including the time spent by the messages in the inbound queue.	Microsec/msg	Compare the value of each of these measures across nodes to identify that node which is the most latent – i.e., slow – when receiving messages.  You can then compare the <i>Average receive latency including inbound queue</i> and <i>Average receive latency excluding inbound queue</i> measures for that node, to understand where the messages spent maximum time – in the inbound queue? or in the node? This will point you to where the bottleneck is.
	<b>Average receive latency excluding inbound queue:</b> Indicates the average time taken by this node to receive messages, excluding the time spent by the messages in the inbound queue.	Microsec/msg	
	<b>Average send latency including outbound queue:</b> Indicates the average time taken by this node to send messages, including the time spent by the messages in the outbound queue.	Microsec/msg	Compare the value of each of these measures across nodes to identify that node which is the most latent – i.e., slow – when sending messages.  You can then compare the <i>Average send latency including outbound queue</i> and <i>Average send latency excluding outbound queue</i> measures for that node, to understand where the messages spent maximum time – in the inbound queue? or in the node? This will point you to where the bottleneck is.
	<b>Average send latency excluding outbound queue:</b> Indicates the average time taken by this node to send messages, excluding the time spent by the messages in the outbound queue.	Microsec/msg	

### 1.2.3 v7000 CPU Usage Test

A Node canister is a hardware unit that includes the node hardware, fabric and service interfaces, and serial-attached SCSI (SAS) expansion ports. Each control enclosure of the IBM Storwize v7000 system consists of two node canisters. In a clustered system therefore, where a collection of control enclosures are managed as a single system, multiple node canisters are bound to exist. At the heart of each node canister is an Intel Jasper Forrest (Sandy Bridge) based quad core CPU. Optimal usage of this CPU is key to reducing latencies in I/O processing. In other words, if a node canister utilizes its processing power excessively, it could cause I/O processing delays, which in turn

## Monitoring IBM Storwize v7000 storage system

can scar the user experience with the storage system. This is why, administrators should keep an eye on the CPU usage of each canister, so that abnormal usage patterns can be detected proactively and treated, before end-users complain. This is where the **v7000 CPU Usage** test helps. This test monitors the CPU usage of each node canister of the IBM Storwize v7000 storage system and reports how long every node canister used its CPU. In the process, the test points to those canisters that are over-utilizing their processing power.

<b>Purpose</b>	Monitors the CPU usage of each node canister of the IBM Storwize v7000 storage system and reports how long every node canister used its CPU		
<b>Target of the test</b>	An IBM Storwize v7000 storage system		
<b>Agent deploying the test</b>	A remote agent		
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"><li>1. <b>TEST PERIOD</b> - How often should the test be executed</li><li>2. <b>HOST</b> - The host for which the test is to be configured.</li><li>3. <b>PORT</b> - The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li><li>4. <b>TIMEOUT</b> -Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li></ol>		
<b>Outputs of the test</b>	One set of results for every node canister being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>CPU busy:</b> Indicates the time for which the CPU of this node canister was busy.	Millisec	A high value for this measure indicates that the CPU is busy for a longer time and the processor is working overtime to complete the job.
	<b>CPU busy Percentage:</b> Indicates the percentage of time the CPU of this node canister was busy.	Percent	A value close to 100 indicates a processing bottleneck.

## 1.3 The IBM v7000 Cache Layer

The **IBM v7000 Cache** layer monitors each type of transaction and reports how many transactions are currently open, blocked and the maximum time taken for execution of a transaction.



Figure 1.5: The tests mapped to the IBM v7000 Cache layer

### 1.3.1 v7000 VDisk Cache Test

Each node canister in the control enclosure has the cache memory, the internal drives to store software and logs, and the processing power to run the Storwize V7000 virtualizing and management software. The two nodes within a canister make an I/O group that is attached to the SAN fabric. VDisks are also associated with the I/O group. I/O traffic for a particular VDisk is, at any one time, managed exclusively by the nodes in a single I/O group. You can even set a particular node in an I/O group as the preferred node for the VDisk. In this case, read I/O to that VDisk is processed by first referencing the cache in the preferred node of the VDisk. If the preferred node is inaccessible, then the cache of the partner node is referenced. Likewise, write operations to a VDisk are first targeted at the preferred node of that VDisk. The write operation is then cached in the preferred node and a copy of the data is made in the cache of the partner node in the I/O group. Only then is the write deemed as complete. The data so written is destaged to the disk at a later time.

The key purpose of the VDisk cache is to buffer I/O, allow the quick completion of I/O operations, and to avoid (or at least, minimize) direct disk accesses. If the cache does not contain the data that the read I/O references and does not have the space to accommodate the data blocks written to it by the write I/O, direct disk accesses are bound to increase. Reading from and writing directly to disk are expensive operations that consume time and resources. If this is to be avoided, the VDisk cache has to be optimally sized and utilized. But, how does one know how well the cache is being used? For this, you can use the **v7000 Disk Cache** test. For each VDisk in the IBM Storwize v7000 storage system, this test monitors the I/O load on the VDisk, reports how well the cache services the read and write I/O load, and in the process, indicates whether/not the cache has been effectively utilized. I/O processing bottlenecks (if any) in the VDisk and inoptimal cache usage can thus be detected and the reasons for the same investigated.

<b>Purpose</b>	For each VDisk in the IBM Storwize v7000 storage system, this test monitors the I/O load on the VDisk, reports how well the cache services the read and write I/O load, and in the process, indicates whether/not the cache has been effectively utilized
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent

<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>		
<b>Outputs of the test</b>	One set of results for the cache of each VDisk of the IBM Storwize v7000 storage system being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Sectors read:</b> Indicates the rate at which the sectors are read from this VDisk.	Sectors/Sec	Comparing these measures across the VDIsks helps you to identify the VDIsks that are currently busy processing the I/O operations.
	<b>Sectors written:</b> Indicates the rate at which the sectors are written on this VDisk.	Sectors/Sec	
	<b>Prestage sector count:</b> Indicates the rate at which prestage sectors were initiated by the cache of this VDisk.	Sectors/Sec	To reduce latency, the storage systems typically uses prestage operations to retrieve data from a disk drive into a cache before the data is retrieved by the next host I/O request. The value of this measure therefore indicates how much data per second was read-ahead into the cache by the storage system. Ideally, this value should be high, so that enough data is prestaged.
	<b>Read cache sector hits:</b> Indicates the number of sectors (of prestaged or non-prestaged data) that were read from the cache every second to service read I/O requests to this VDisk.	Sectors/Sec	Ideally, the value of this measure should be high. A consistent drop in this value is a cause for concern, as it could indicate that read I/O requests are not being frequently serviced by the cache. This in turn implies low cache usage and high disk usage.
	<b>Read cache prestaged sector hits:</b> Indicates the number of sectors (of prestaged data) that were read from the cache every second to service read I/O requests to this VDisk.	Sectors/Sec	Ideally, the value of this measure should be high. A high value could indicate that many host I/O requests were serviced by data that was read into the cache, well ahead of the requests. Since prestaging reduces latency to a great extent, a high value is a sign of good cache health.

	<b>Read cache sector misses:</b> Indicates the rate at which the sectors were read for a read I/O request to this VDisk that has cache misses.	Sectors/Sec	A low value is desired for this measure. A high value indicates too many cache misses, which in turn, is indicative of poor cache configuration.
	<b>Track destage sector count:</b> Indicates the rate at which the sectors were written for the track writes initiated by this VDisk's cache.	Sectors/Sec	
	<b>Flush through sector writes:</b> Indicates the rate at which the sectors were written using flush through write mode for the write requests to this VDisk.	Sectors/Sec	
	<b>Write through sector writes:</b> Indicates the rate at which the sectors were written using write through write mode for the write requests to this VDisk.	Sectors/Sec	In the write-through mode, write is done synchronously to both the cache and to the disk. Write-through is the mode to set where workload is high.
	<b>Fast write sector writes:</b> Indicates the rate at which the sectors were written using fast write mode for the write requests to this VDisk.	Sectors/Sec	<p>Fast write cache (FWC) is an optional nonvolatile cache that provides redundancy with the standard adapter cache. The FWC tracks writes that have not been committed to disk.</p> <p>Fast write cache can significantly improve the response time for write operations. However, care must be taken not to flood the cache with write requests faster than the rate at which the cache can destage its data. FWC can also adversely affect the maximum I/O rate because additional processing is required in the adapter card to determine if the data that is being transferred is in the cache.</p>



	<b>Lack of memory fast write track sec:</b> Indicates the rate at which the sectors that were processed in fast-write mode were written using write-through mode owing to lack of memory in this VDisk's cache.	Sectors/Sec	
	<b>Dirty data write hit sectors:</b> Indicates the rate at which the write requests to this VDisk were serviced by sectors from dirty data in the cache.	Sectors/Sec	
	<b>Write cached data:</b> Indicates the amount of modified/dirty data that is held in this VDisk cache.	MB	A consistent increase in the value of this measure indicates that data destaging is not occurring frequently – i.e., data is not being flushed to the disk at regular intervals. If the situation persists, the cache may fill up to capacity and may reject subsequent entries. This can increase cache misses and degrade I/O performance of the storage system.
	<b>Read and write cached data:</b> Indicates the amount of data that is held in this VDisk cache.	MB	A consistent increase in the value of this measure could indicate one or both of the following: <ul style="list-style-type: none"> <li>• Old data is not being discarded;</li> <li>• Data written to the cache is not being flushed to the disk at regular intervals</li> </ul> If the situation persists, the cache may fill up to capacity and may reject subsequent entries. This can increase cache misses and degrade I/O performance of the storage system. To avoid such an eventuality, you may want to fine-tune your cache algorithm to make it more efficient.

## 1.4 The IBM v7000 MDisks Layer

The **IBM v7000 MDisks** layer tracks the status of each MDisk , the I/O operations performed on each MDisk, the capacity of each MDisk etc.



Figure 1.6: The tests mapped to the IBM v7000 MDisks layer

### 1.4.1 v7000 MDisks Test

A managed disk (MDisk) is a logical unit of physical storage. MDisks are either arrays (RAID) from internal storage or volumes from external storage systems. MDisks are not visible to host systems.

An MDisk might consist of multiple physical disks that are presented as a single logical disk to the storage area network (SAN). Each MDisk is divided into a number of extents, which are numbered, from 0, sequentially from the start to the end of the MDisk.

Real capacity to the VDisks/volumes is provided by groups of MDisks, known as storage pools. This is because, volumes are created from the extents that are available in these pools. It is hence evident that I/O requests to a volume are in effect serviced by the MDisks in the pools from which that volume was created. This means that a processing bottleneck in an MDisk can significantly delay reading from and writing to the VDisk that overlays it, and can thus adversely impact the overall I/O performance of the storage system. This is why, it is important that administrators track the I/O load on and the I/O processing speed of each MDisk in each node of the IBM Storwize v7000 system. The **v7000 MDisks** test helps administrators achieve this. This test monitors the I/O operations on each MDisk in each node of the storage system, reports how well every MDisk processes the I/O, and reveals processing latencies (if any) in that MDisk. This way, MDisks that are overloaded, those that are slow in processing I/O, and the nodes that such MDisks are mapped to, can be accurately identified.

<b>Purpose</b>	Monitors the I/O operations on each MDisk in each node of the storage system, reports how well every MDisk processes the I/O, and reveals processing latencies (if any) in that MDisk.
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent

<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li><b>TEST PERIOD</b> - How often should the test be executed</li> <li><b>HOST</b> - The host for which the test is to be configured.</li> <li><b>PORT</b> - The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li><b>TIMEOUT</b> -Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>		
<b>Outputs of the test</b>	One set of results for every MDisk in each node of the storage system being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Read operations:</b> Indicates the rate at which read operations were performed by this MDisk of this node.	Reads/Sec	Comparing the value of each of these measures across MDisks helps you in identifying the overloaded MDisks - it could shed light on irregularities on load balancing across the MDisks.
	<b>Write operations:</b> Indicates the rate at which write operations were performed by this MDisk of this node.	Writes/Sec	
	<b>Read data:</b> Indicates the rate at which data blocks were read from this MDisk of this node.	MB/Sec	By comparing the value of each of these measures across MDisks, you can identify the MDisk that is the slowest in reading and writing. The reason for the slowness has to be determined and eliminated to ensure the high availability and performance of the storage system.
	<b>Write data:</b> Indicates the rate at which the data blocks were written to this MDisk of this node.	MB/Sec	
	<b>Average read external response:</b> Indicates the average time taken by this MDisk to respond to read requests.	Millisec	For each SCSI Read command sent to the Managed disk, a timer is started when the command is issued across the fibre channel from the host system to the IBM Storwize v7000 storage system. When the command completes, the elapsed time since the timer was started is added to the cumulative counter.  A low value is desired for this counter. A high value is indicative of slowness when responding to read requests.  The least responsive MDisk can be identified by comparing the value of this measure across MDisks.

	<b>Average write external response:</b> Indicates the average time taken by this MDisk to respond to write requests.	Millisec	<p>A low value is desired for this counter. A high value is indicative of slowness when responding to write requests.</p> <p>To know which MDisk is taking an unreasonably long time to service write requests, compare the value of this measure across MDisk.</p>
	<b>Peak read external response:</b> Indicates the maximum time taken by this MDisk to respond to a read request during the last measurement period.	Millisec	<p>Compare the value of this measure across the MDisk to identify the MDisk that was the slowest in responding to read requests during the last measurement period. If the same MDisk tops this comparison consistently, it could indicate a read I/O processing bottleneck in that MDisk.</p>
	<b>Peak write external response:</b> Indicates the maximum time taken by this MDisk to respond to a write request during the last measurement period.	Millisec	<p>Compare the value of this measure across the MDisk to identify the MDisk that was the slowest in responding to write requests during the last measurement period. If the same MDisk tops this comparison consistently, it could indicate a write I/O processing bottleneck in that MDisk.</p>
	<b>Peak read queued response:</b> Indicates the maximum time a read request was waiting in the queue before being sent to this MDisk during the last measurement period.	Millisec	<p>The value of this measure includes the time spent by the read request in the queue and the time taken for execution of this request by the MDisk.</p> <p>By comparing the value of this measure with that of the <i>Peak read external response</i> measure, you can understand where the read requests could have spent maximum time – in the queue? or in the MDisk, being processed?</p>
	<b>Peak write queued response:</b> Indicates the maximum time a write request was waiting in the queue before being sent to this MDisk during the last measurement period.	Millisec	<p>The value of this measure includes the time spent by the write request in the queue and the time taken for execution of this request by the storage controller.</p> <p>By comparing the value of this measure with that of the <i>Peak write external response</i> measure, you can understand where the read requests could have spent maximum time – in the queue? or in the MDisk, being processed?</p>

## 1.4.2 v7000 MDisk Status Test

A managed disk (MDisk) refers to the unit of storage that IBM Storwize V7000 virtualizes. This unit could be a logical volume on an external storage array presented to the IBM Storwize V7000 or a RAID array consisting of internal drives. The IBM Storwize V7000 can then allocate these MDisk into various storage pools.

An MDisk is not visible to a host system on the storage area network, as it is internal or only zoned to the IBM Storwize V7000 system.

An MDisk has four modes:

- **Array:** Array mode MDisk are constructed from drives using the RAID function. Array MDisk are always associated with storage pools.
- **Unmanaged:** Unmanaged MDisk are not being used by the system. This situation might occur when an MDisk is first imported into the system, for example.
- **Managed:** Managed MDisk are assigned to a storage pool and provide extents so that volumes can use it.
- **Image:** Image MDisk are assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. This situation is normally used when importing logical volumes into the clustered system that already have data on them, which ensures that the data is preserved as it is imported into the clustered system.

If an MDisk assigned to a cluster cannot be accessed by any of the cluster nodes, then the nodes will not be able to service the host I/O requests they receive, resulting in poor I/O performance of the storage system. If this is to be avoided, then, the current state of every MDisk in the storage system should be tracked continuously, and abnormalities promptly brought to the attention of administrators. This will enable administrators to initiate corrective action and bring the state of the MDisk back to normal. To achieve this, administrators can use the **v7000 MDisk Status** test. This test reports the current state of each MDisk and the status of the RAID array that hosts the MDisk. In the process, administrators get to know which MDisk and RAID arrays are inaccessible to cluster nodes, investigate the reasons for the anomaly, and resolve it, before it affects the overall I/O performance of the storage system.

<b>Purpose</b>	Reports the current state of each MDisk and the status of the RAID array that hosts the MDisk
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> <li>5. <b>DETAILED DIAGNOSIS</b> To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>○ The eG manager license should allow the detailed diagnosis capability</li> <li>○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul> </li> </ol>		
Outputs of the test	One set of results for every MDisk being monitored		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	<p><b>Status:</b></p> <p>Indicates the current status of this MDisk.</p>	<p>This measure reports any of the following values:</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ excluded</li><li>➤ degraded</li></ul> <p>The table below describes each of the aforesaid measure values:</p> <table><tr><th>State</th><th>Description</th></tr><tr><td>Offline</td><td>The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN,storage system, or one or more physical disks connected to the storage system. The MDisk is reported as offline if all paths to the disk fail.</td></tr><tr><td>online</td><td>The MDisk can be accessed by all online nodes. TheMDisk is online when the following conditions are met:<ul style="list-style-type: none"><li>• All timeout error recovery procedures complete and report the disk as online.</li><li>• Logical unit number (LUN) inventory of the target ports correctly reported the MDisk.</li><li>• Discovery of this LUN completed successfully.</li><li>• All of the MDisk target ports report this LUN as available with no fault conditions.</li></ul></td></tr><tr><td>excluded</td><td>The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem.</td></tr></table>	State	Description	Offline	The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN,storage system, or one or more physical disks connected to the storage system. The MDisk is reported as offline if all paths to the disk fail.	online	The MDisk can be accessed by all online nodes. TheMDisk is online when the following conditions are met: <ul style="list-style-type: none"><li>• All timeout error recovery procedures complete and report the disk as online.</li><li>• Logical unit number (LUN) inventory of the target ports correctly reported the MDisk.</li><li>• Discovery of this LUN completed successfully.</li><li>• All of the MDisk target ports report this LUN as available with no fault conditions.</li></ul>	excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem.
State	Description									
Offline	The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN,storage system, or one or more physical disks connected to the storage system. The MDisk is reported as offline if all paths to the disk fail.									
online	The MDisk can be accessed by all online nodes. TheMDisk is online when the following conditions are met: <ul style="list-style-type: none"><li>• All timeout error recovery procedures complete and report the disk as online.</li><li>• Logical unit number (LUN) inventory of the target ports correctly reported the MDisk.</li><li>• Discovery of this LUN completed successfully.</li><li>• All of the MDisk target ports report this LUN as available with no fault conditions.</li></ul>									
excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem.									

			<div> <div>degraded</div> <div> <p>This can be owing to degraded paths or degraded ports. A degraded path can render an MDisk inaccessible to one/more nodes in the cluster. Degraded path status is most likely the result of incorrect configuration of either the disk controller or the fibre-channel fabric. However, hardware failures in the disk controller, fibre-channel fabric, or node could also be a contributing factor to this state. Complete the following actions to recover from this state:</p> <ul style="list-style-type: none"> <li>• Verify that the fabric configuration rules for storage systems are correct.</li> <li>• Ensure that you have configured the storage system properly.</li> <li>• Correct any errors in the event log.</li> </ul> <p>If the MDisk has one/more 1220 errors in the event log, then the degraded state is owing to degraded ports. The 1220 error indicates that the remote fibre-channel port has been excluded from the MDisk. This error might cause reduced performance on the storage controller and usually indicates a hardware problem with the storage controller. To fix this problem you must resolve any hardware problems on the storage controller and fix the 1220 errors in the event log.</p> </div> </div>
--	--	--	---



			<p>The numeric values that correspond to the above-mentioned states are as follows:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>offline</td><td>0</td></tr><tr><td>online</td><td>1</td></tr><tr><td>excluded</td><td>2</td></tr><tr><td>degraded</td><td>3</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of the MDisk. However, in the graph of this measure, this will be represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this measure reveals the MDisk ID, the MDisk GROUP ID, the MDisk GROUP NAME, the RAID STATUS, the RAID LEVEL, REDUNDANCY of this MDisk, the STRIP SIZE of the MDisk and the TIER to which the MDisk is associated with.</p>	Measure Value	Numeric Value	offline	0	online	1	excluded	2	degraded	3
Measure Value	Numeric Value												
offline	0												
online	1												
excluded	2												
degraded	3												
	<p><b>Capacity:</b></p> <p>Indicates the total capacity of this MDisk.</p>	TB											

	<p><b>Raid status:</b></p> <p>Indicates the current status of the RAID array hosting this MDisk.</p>	MB/Sec	<p>This measure reports any one of the following values:</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ excluded</li><li>➤ degraded</li></ul> <p>The numeric values that correspond to the above-mentioned measure values are as follows:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr><tr><td>offline</td><td>0</td><td>the array is offline on all nodes</td></tr><tr><td>degraded</td><td>1</td><td>the array has deconfigured or offline members; the array is not fully redundant</td></tr><tr><td>syncing</td><td>2</td><td>array members are all online, the array is syncing parity or mirrors to achieve redundancy</td></tr><tr><td>initting</td><td>3</td><td>array members are all online, the array is initializing; the array is fully redundant</td></tr><tr><td>online</td><td>4</td><td>array members are all online, and the array is fully redundant</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of the raid array hosting this MDisk. However, in the graph of this measure, the state will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	Description	offline	0	the array is offline on all nodes	degraded	1	the array has deconfigured or offline members; the array is not fully redundant	syncing	2	array members are all online, the array is syncing parity or mirrors to achieve redundancy	initting	3	array members are all online, the array is initializing; the array is fully redundant	online	4	array members are all online, and the array is fully redundant
Measure Value	Numeric Value	Description																			
offline	0	the array is offline on all nodes																			
degraded	1	the array has deconfigured or offline members; the array is not fully redundant																			
syncing	2	array members are all online, the array is syncing parity or mirrors to achieve redundancy																			
initting	3	array members are all online, the array is initializing; the array is fully redundant																			
online	4	array members are all online, and the array is fully redundant																			

## 1.5 The IBM v7000 Drives Layer

The **IBM v7000 Drives** layer tracks the various I/O activities of each drive supported by the IBM Storwize v7000 storage system, the capacity of each drive, the status of each drive etc.




Figure 1.7: The tests mapped to the IBM v7000 Drives layer

### 1.5.1 v7000 Drives Test

IBM Storwize V7000 enclosures currently support SSD, SAS, and Nearline SAS drive types. Each SAS drive has two ports (two PHYs) and I/O can be issued down both paths simultaneously. If even a single drive lags behind in I/O processing, the overall I/O performance of the storage system will suffer. It is therefore, imperative that administrators watch out for slowness in drives and proactively detect potential I/O processing bottlenecks in drives, so that end-users need not have to deal with slowness when reading from or writing into the storage system. The **v7000 Drives** test helps administrators with this. For each drive in the IBM Storwize v7000 storage system, this test reports the load on the drive and how well the drive handles the load. This way, overloaded drives and those experiencing processing slowdowns can be identified quickly and accurately.

<b>Purpose</b>	For each drive in the IBM Storwize v7000 storage system, this test reports the load on the drive and how well the drive handles the load. This way, overloaded drives and those experiencing processing slowdowns can be identified quickly and accurately
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>

<b>Outputs of the test</b>	One set of results for each drive of the IBM Storwize v7000 being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Read operations:</b> Indicates the rate at which read operations were performed on this drive.	Reads/Sec	Comparing the value of each of these measures across drives helps you in identifying the overloaded drives - it could shed light on irregularities on load balancing across the drives.
	<b>Write operations:</b> Indicates the rate at which write operations were performed to this drive.	Write/Sec	
	<b>Read data:</b> Indicates the rate at which data blocks were read from this drive.	MB/Sec	By comparing the value of each of these measures across drives, you can identify the drive that is the slowest in reading and writing. The reason for the slowness has to be determined and eliminated to ensure the high availability and performance of the storage system.
	<b>Write data:</b> Indicates the rate at which the data blocks were written to this drive.	MB/Sec	
	<b>Average read external response:</b> Indicates the average time taken by this drive to respond to read requests.	Millisec	A low value is desired for this counter. A high value is indicative of slowness when responding to read requests.  The least responsive drive can be identified by comparing the value of this measure across drives.
	<b>Average write external response:</b> Indicates the average time taken by this drive to respond to write requests.	Millisec	A low value is desired for this counter. A high value is indicative of slowness when responding to write requests.  To know which drive is taking an unreasonably long time to service write requests, compare the value of this measure across drives.
	<b>Peak read external response:</b> Indicates the maximum time taken by this drive to respond to a read request during the last measurement period.	Millisec	Compare the value of this measure across the drives to identify the drive that was the slowest in responding to read requests during the last measurement period. If the same drive tops this comparison consistently, it could indicate a read I/O processing bottleneck in that drive.

	<b>Peak write external response:</b> Indicates the maximum time taken by this drive to respond to a write request during the last measurement period.	Millisec	Compare the value of this measure across the drives to identify the drive that was the slowest in responding to write requests during the last measurement period. If the same drive tops this comparison consistently, it could indicate a write I/O processing bottleneck in that drive.
	<b>Peak read queued response:</b> Indicates the maximum time a read request was waiting in the queue before being sent to this drive during the last measurement period.	Millisec	<p>The value of this measure includes the time spent by the read request in the queue and the time taken for execution of this request by the drive.</p> <p>By comparing the value of this measure with that of the <i>Peak read external response</i> measure, you can understand where the read requests could have spent maximum time – in the queue? or in the drive, being processed?</p>
	<b>Peak write queued response:</b> Indicates the maximum time a write request was waiting in the queue before being sent to this drive during the last measurement period.	Millisec	<p>The value of this measure includes the time spent by the write request in the queue and the time taken for execution of this request by the drive.</p> <p>By comparing the value of this measure with that of the <i>Peak write external response</i> measure, you can understand where the read requests could have spent maximum time – in the queue? or in the drive, being processed?</p>

## 1.5.2 v7000 Drive Status Test

A drive that is offline or degraded will be inaccessible to I/O requests, until brought online or the source of the degradation is determined and removed. Using the **v7000 Drive Status** test, administrators can locate such drives instantly and have the problem fixed before end-users notice.

<b>Purpose</b>	Reports the status and capacity of each drive supported by the IBM Storwize v7000 storage system
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> <li>5. <b>DETAILED DIAGNOSIS</b> To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>○ The eG manager license should allow the detailed diagnosis capability</li> <li>○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul> </li> </ol>		
Outputs of the test	One set of results for each drive of the IBM Storwize v7000 storage system being monitored		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	<p><b>Status:</b></p> <p>Indicates the current status of this drive.</p>		<p>This measure reports the status of this drive as follows:</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ degraded</li></ul> <p>The numeric values that correspond to the above-mentioned states are as follows:</p> <table><tr><th>State</th><th>Numeric Value</th></tr><tr><td>offline</td><td>0</td></tr><tr><td>online</td><td>1</td></tr><tr><td>degraded</td><td>2</td></tr></table> <p>Note:</p> <p>By default, this measure reports the above-mentioned <b>States</b> while indicating the status of this drive. However, the graph of this measure will be represented using the corresponding numeric equivalents of the <b>States</b> as mentioned in the table above.</p> <p>The detailed diagnosis of this measure if enabled, lists the CURRENT ROLE of this drive, the ENCLOSURE ID and SLOT ID to which this drive is associated, the MEMBER ID of this drive, the MDISK NAME, ERROR SEQ NO and the DRIVE TECHNOLOGY USED.</p>	State	Numeric Value	offline	0	online	1	degraded	2
State	Numeric Value										
offline	0										
online	1										
degraded	2										
	<p><b>Capacity:</b></p> <p>Indicates the current capacity of this drive excluding quorum area.</p>	TB	<p>A quorum disk is a managed disk (MDisk) that contains a reserved area for use exclusively by the system. In the IBM Storwize V7000, internal drives can be considered as quorum candidates. The clustered system uses quorum disks to break a tie when exactly half the nodes in the system remain after a SAN failure.</p>								

## 1.6 The Hardware Layer

The **Hardware** layer provides you with a wealth of knowledge related to the status of the enclosure, the status of the enclosure battery, the status of the ports, the status of the enclosure canisters etc.



Figure 1.8: The tests mapped to the Hardware layer

### 1.6.1 v7000 Traps Test

This test intercepts the SNMP traps sent out by the IBM Storwize v7000 storage system, searches the traps for user-configured events, and reports the number of trap messages related to the configured events. Using the detailed diagnostics provided by this test, you can view the complete description of the trap messages that match your configuration, and thus capture critical errors/warnings thrown by the storage system.

<b>Purpose</b>	Intercepts the SNMP traps sent out by the IBM Storwize v7000 storage system, searches the traps for user-configured events, and reports the number of trap messages related to the configured events
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A 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
<i>.1.3.6.1.4.1.9156.1.1.2</i>	Host_system
<i>.1.3.6.1.4.1.9156.1.1.3</i>	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.*

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.

	<p>7. <b>DD FREQUENCY</b> - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. 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 <i>none</i> against <b>DD FREQUENCY</b>.</p> <p>8. <b>DETAILED DIAGNOSIS</b> - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the <b>On</b> option. To disable the capability, click on the <b>Off</b> option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>○ The eG manager license should allow the detailed diagnosis capability</li> <li>○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>		
<b>Outputs of the test</b>	One set of results for each of the traps configured for the storage system being monitored		
<b>Measurements made by the test</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>
	<b>Number of events:</b> Indicates the number of trap messages of this event that were sent from this storage system.	Number	The detailed diagnosis of this measure if enabled, provides the details of the events.

## 1.6.2 v7000 Enclosures Test

Each IBM Storwize V7000 system has one or two (optional) control enclosures that contain node canisters, disk drives, and power supply units. If a control enclosure is not visible to the SAN, then users will not be able to access the storage system. This is why, administrators need to closely monitor the state of each enclosure, and ensure its continuous availability to the SAN. This can be achieved using the **v7000 Enclosures** test. This test monitors each enclosure and reports its current state and composition, so that administrators can be instantly alerted to the abnormal state of the enclosure.

<b>Purpose</b>	Monitors each enclosure and reports its current state and composition, so that administrators can be instantly alerted to the abnormal state of the enclosure
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> <li>4. <b>DETAILED DIAGNOSIS</b> To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>○ The eG manager license should allow the detailed diagnosis capability</li> <li>○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul> </li> </ol>		
Outputs of the test	One set of results for every enclosure being monitored		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	<p><b>Status:</b></p> <p>Indicates the current status of this enclosure.</p>		<p>This measure reports any one of the following values:</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ degraded</li></ul> <p>The numeric values that correspond to the above-mentioned measure values are as follows:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr><tr><td>offline</td><td>0</td><td>A managed or unmanaged enclosure is visible</td></tr><tr><td>online</td><td>1</td><td>A managed enclosure is not visible, and other fields hold their last known values</td></tr><tr><td>degraded</td><td>2</td><td>An enclosure is visible, but not down both strands</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of the enclosure. However, in the graph of this measure, enclosure state will be represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this measure if enabled, lists the type of the enclosure, the managed status of the enclosure, the IO group ID, the IO group name, the machine type and the serial number of the enclosure. This way, you can accurately determine the I/O groups that will be rendered inaccessible if an enclosure goes offline or is degraded.</p>	Measure Value	Numeric Value	Description	offline	0	A managed or unmanaged enclosure is visible	online	1	A managed enclosure is not visible, and other fields hold their last known values	degraded	2	An enclosure is visible, but not down both strands
Measure Value	Numeric Value	Description													
offline	0	A managed or unmanaged enclosure is visible													
online	1	A managed enclosure is not visible, and other fields hold their last known values													
degraded	2	An enclosure is visible, but not down both strands													
	<p><b>Total Canisters:</b></p> <p>Indicates the total number of canisters available in this enclosure.</p>	Number													

	<b>Online Canisters:</b> Indicates the number of canisters contained in this enclosure that are currently online.	Number	
	<b>Total PSUs:</b> Indicates the total number of power supply units available in this enclosure.	Number	
	<b>Online PSUs:</b> Indicates the number of power supply units contained in this enclosure that are currently online.	Number	
	<b>Drive slots:</b> Indicates the total number of slots that can accommodate drives in this enclosure.	Number	

### 1.6.3 v7000 Enclosure Slots Test

This test reports the status of the ports that are present on each of the enclosure slot provided on the IBM Storwize v7000 storage system. Further, this test reveals the information on whether a drive is inserted into this slot and if so, reports the ID of the inserted drive.

<b>Purpose</b>	Reports the status of the ports that are present on each of the Enclosure slot provided on the IBM Storwize v7000 storage system. Further, this test reveals the information on whether a drive is inserted into this slot and if so, reports the ID of the inserted drive.		
<b>Target of the test</b>	An IBM Storwize v7000 storage system		
<b>Agent deploying the test</b>	A remote agent		
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>		
<b>Outputs of the test</b>	One set of results for every enclosure slot being monitored		
<b>Measurements made by the</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>

test	<p><b>Port1 Status:</b></p> <p>Indicates the current status of port 1 on this Enclosure slot.</p>	<p>This measure reports the status of the port 1 as follows:</p> <ul style="list-style-type: none"> <li>➤ online</li> <li>➤ excluded_by_drive</li> </ul> <p>The numeric values that correspond to the above-mentioned states are as follows:</p> <table border="1" data-bbox="946 449 1409 758"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr> </thead> <tbody> <tr> <td>online</td><td>1</td><td>The port is online</td></tr> <tr> <td>excluded_by_drive</td><td>2</td><td>The drive excluded the port</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of the port 1. However, in the graph of this measure, the drive state will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	Description	online	1	The port is online	excluded_by_drive	2	The drive excluded the port
Measure Value	Numeric Value	Description									
online	1	The port is online									
excluded_by_drive	2	The drive excluded the port									
	<p><b>Port2 Status:</b></p> <p>Indicates the current status of port 2 of this Enclosure slot.</p>	<p>This measure reports the status of the port 2 as follows:</p> <ul style="list-style-type: none"> <li>➤ online</li> <li>➤ excluded_by_drive</li> </ul> <p>The numeric values that correspond to the above-mentioned states are as follows:</p> <table border="1" data-bbox="938 1274 1433 1583"> <thead> <tr> <th>Measue Value</th><th>Numeric Value</th><th>Description</th></tr> </thead> <tbody> <tr> <td>online</td><td>1</td><td>The port is online</td></tr> <tr> <td>excluded_by_drive</td><td>2</td><td>The drive excluded the port</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned states while indicating the status of port 2. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p>	Measue Value	Numeric Value	Description	online	1	The port is online	excluded_by_drive	2	The drive excluded the port
Measue Value	Numeric Value	Description									
online	1	The port is online									
excluded_by_drive	2	The drive excluded the port									

	<p><b>Is drive present?:</b></p> <p>Indicates whether/not a drive is inserted into this Enclosure slot.</p>		<p>This measure reports the value <i>Yes</i> if the drive is present and the value <i>No</i> if otherwise.</p> <p>The numeric equivalents of these values are mentioned in the table below:</p> <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> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the presence of the drive in this Enclosure slot. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								
	<p><b>Drive Id:</b></p> <p>Indicates the ID of the drive that is currently inserted into this Enclosure slot.</p>	Number							

## 1.6.4 v7000 Enclosure PSU Status Test

Each IBM Storwize V7000 system has one or two (optional) control enclosures that contain two node canisters each, disk drives, and two power supplies. A failure of any of these power supplies can prove to be fatal to the storage system as a whole. Administrators hence need to be promptly alerted to such failures, so that corrective action can be taken immediately. The **v7000 Enclosure PSU Status** test captures sudden changes in the state of the each of the power supplies and brings it to the attention of the administrators immediately, so that the problem can be redressed before any permanent damage is done.

<b>Purpose</b>	Captures sudden changes in the state of the each of the power supplies and brings it to the attention of the administrators immediately, so that the problem can be redressed before any permanent damage is done.
<b>Target of the test</b>	An IBM Storwize v7000 storage system
<b>Agent deploying the test</b>	A remote agent
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>TIMEOUT</b> -Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> </ol>

Outputs of the test	One set of results for each Power Supply Unit of the control enclosure being monitored													
Measurements made by the test	Measurement	Measurement Unit	Interpretation											
	<b>Status:</b> Indicates the current status of this Power Supply Unit in this Control enclosure.	MB	<p>This measure reports any of the following values</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ degraded</li></ul> <p>The numeric values that correspond to the above-mentioned measure values are as follows:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr><tr><td>offline</td><td>0</td><td>The Power Supply Unit cannot be detected.</td></tr><tr><td>online</td><td>1</td><td>The Power Supply Unit is present and working normally.</td></tr><tr><td>degraded</td><td>2</td><td>The Power Supply Unit is present but not working normally.</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of a Power Supply Unit. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	Description	offline	0	The Power Supply Unit cannot be detected.	online	1	The Power Supply Unit is present and working normally.	degraded	2
Measure Value	Numeric Value	Description												
offline	0	The Power Supply Unit cannot be detected.												
online	1	The Power Supply Unit is present and working normally.												
degraded	2	The Power Supply Unit is present but not working normally.												



## 1.6.5 v7000 Enclosure Canisters Test

The IBM Storwize V7000 solution consists of one or two control enclosures and, optionally, up to 18 expansion enclosures (and supports the intermixing of the different expansion enclosures). Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

The expansion canister is a hardware unit that includes the serial-attached SCSI (SAS) interface hardware that enables the node hardware to use the drives of the expansion enclosure.

The node canister is a hardware unit that includes the node hardware, fabric and service interfaces, and serial-attached SCSI (SAS) expansion ports.

This test reveals the current status of each canister available in the enclosure of the IBM Storwize v7000 storage system, so that defective canisters can be promptly isolated and issues fixed.

<b>Purpose</b>	Reveals the current status of each canister available in the enclosure of the IBM Storwize v7000 storage system		
<b>Target of the test</b>	An IBM Storwize v7000 storage system		
<b>Agent deploying the test</b>	A remote agent		
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"> <li>1. <b>TEST PERIOD</b> - How often should the test be executed</li> <li>2. <b>HOST</b> - The host for which the test is to be configured.</li> <li>3. <b>PORT</b> – The port number at which the specified <b>HOST</b> listens to. By default, this is <i>NULL</i>.</li> <li>4. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li> <li>5. <b>DETAILED DIAGNOSIS</b> To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>○ The eG manager license should allow the detailed diagnosis capability</li> <li>○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul> </li> </ol>		
<b>Outputs of the test</b>	One set of results for each canister being monitored		
<b>Measurements made by the</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>

test	<b>Status:</b> Indicates the current status of this canister.	Percent	<p>This measure can report any of the following values:</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ degraded</li></ul> <p>The numeric values that correspond to the above-mentioned measure values are as follows:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr><tr><td>offline</td><td>0</td><td>The canister could not be detected.</td></tr><tr><td>online</td><td>1</td><td>The canister is present and working normally.</td></tr><tr><td>degraded</td><td>2</td><td>The canister is present but not working normally.</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of a canister. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this measure if enabled, lists the type of the canister, the node ID and the name of the node to which the canister is attached.</p>	Measure Value	Numeric Value	Description	offline	0	The canister could not be detected.	online	1	The canister is present and working normally.	degraded	2	The canister is present but not working normally.
	Measure Value		Numeric Value	Description											
	offline		0	The canister could not be detected.											
	online		1	The canister is present and working normally.											
	degraded		2	The canister is present but not working normally.											

### 1.6.6 v7000 Enclosure Battery Status Test

Within a control enclosure, each power supply unit (PSU) contains a battery. The battery is designed to enable the IBM Storwize V7000 system to perform a memory dump of the cache to internal disks in the event of both power inputs failing. If these batteries are defective or have reached end-of-life, the storage system may be unsuccessful in writing cached data to the disks in the event of a power failure; this may result in loss of critical data. To avoid this, potential battery failures should be proactively detected and intimated to administrators, so that the problem can be fixed before any data loss occurs. This is where the **v7000 Enclosure Battery Status** test helps. For each battery in the PSUs, this test reports the battery health, whether it has adequate charge or not, and whether/not it is reaching its

## Monitoring IBM Storwize v7000 storage system

end-of-life. This way, the test warns administrators of an impending battery failure, and enables them to take pre-emptive action to correct the issue.

<b>Purpose</b>	For each battery in the PSUs, this test reports the battery health, whether it has adequate charge or not, and whether/not it is reaching its end-of-life		
<b>Target of the test</b>	An IBM v7000 storage system		
<b>Agent deploying the test</b>	A remote agent		
<b>Configurable parameters for the test</b>	<ol style="list-style-type: none"><li>1. <b>TEST PERIOD</b> - How often should the test be executed</li><li>2. <b>HOST</b> - The host for which the test is to be configured.</li><li>3. <b>TIMEOUT</b> –Specify the duration (in seconds) beyond which the test will timeout in the <b>TIMEOUT</b> text box. The default value is 60 seconds.</li></ol>		
<b>Outputs of the test</b>	One set of results for each battery present in the enclosure of the IBM Storwize v7000 storage system being monitored		
<b>Measurements made by the</b>	<b>Measurement</b>	<b>Measurement Unit</b>	<b>Interpretation</b>

test	<b>Battery status:</b> Indicates the current health status of this battery.		<p>This measure can report any of the following values:</p> <ul style="list-style-type: none"><li>➤ offline</li><li>➤ online</li><li>➤ degraded</li></ul> <p>The numeric values that correspond to the above-mentioned states are as follows:</p> <table><tr><th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr><tr><td>offline</td><td>0</td><td>The battery could not be detected.</td></tr><tr><td>online</td><td>1</td><td>The battery is present and working as usual.</td></tr><tr><td>degraded</td><td>2</td><td>The battery is present but not working as usual.</td></tr></table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the status of this battery. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	Description	offline	0	The battery could not be detected.	online	1	The battery is present and working as usual.	degraded	2	The battery is present but not working as usual.
	Measure Value	Numeric Value	Description												
	offline	0	The battery could not be detected.												
	online	1	The battery is present and working as usual.												
	degraded	2	The battery is present but not working as usual.												

	<p><b>Charging Status:</b></p> <p>Indicates the charging status of this battery.</p>	<p>This measure can report any of the following values:</p> <ul style="list-style-type: none"> <li>➤ Idle</li> <li>➤ charging</li> <li>➤ discharging</li> </ul> <p>The numeric values that correspond to the above-mentioned measure values are as follows:</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th><th>Description</th></tr> </thead> <tbody> <tr> <td>Idle</td><td>0</td><td>The battery is neither charging nor discharging.</td></tr> <tr> <td>charging</td><td>1</td><td>The battery is charging.</td></tr> <tr> <td>discharging</td><td>2</td><td>The battery is discharging at present.</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the charging status of a battery. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	Description	Idle	0	The battery is neither charging nor discharging.	charging	1	The battery is charging.	discharging	2	The battery is discharging at present.
Measure Value	Numeric Value	Description												
Idle	0	The battery is neither charging nor discharging.												
charging	1	The battery is charging.												
discharging	2	The battery is discharging at present.												

	<p><b>Is recondition needed?:</b></p> <p>Indicates whether this battery needs to be reconditioned or not.</p>		<p>This measure reports the value <i>Yes</i> if the battery needs to be reconditioned and the value <i>No</i> if the battery requires reconditioning, but cannot be reconditioned due to one or more errors.</p> <p>The numeric equivalents of these values are mentioned in the table below:</p> <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> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating the reconditioning status of this battery. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								
	<p><b>Percent charged:</b></p> <p>Indicates the current charge percentage of this battery.</p>	Percent	<p>A high value is desired for this measure. A value close to 100% could indicate that the battery charge is about to die.</p>						
	<p><b>Is end of life?:</b></p> <p>Indicates whether this battery is reaching its end of life and needs to be replaced or not.</p>		<p>This measure reports the value <i>Yes</i> if the battery is reaching its end of life and a value <i>No</i> if otherwise.</p> <p>The numeric equivalents of these values are mentioned in the table below:</p> <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> <p><b>Note:</b></p> <p>By default, this measure reports the above-mentioned <b>Measure Values</b> while indicating whether this battery is reaching its end of life or not. However, in the graph of this measure, the same will be represented using the corresponding numeric only.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								

## Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to the **IBM Storwize v7000** storage system. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact [support@eginnovations.com](mailto:support@eginnovations.com). We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to [feedback@eginnovations.com](mailto:feedback@eginnovations.com).