



Monitoring Solaris Zones

eG Enterprise v6.0

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 2003, and Windows 2000 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

© 2012 eG Innovations, Inc. All rights reserved.

The copyright in this document belongs to eG Innovations, Inc. Complying with all applicable copyright laws is the responsibility of the user.

Table of Contents

MONITORING SOLARIS ZONES	1
1.1 What are Solaris Zones?.....	2
1.2 How does eG Enterprise Monitor Solaris Zones?	3
1.2.1 The Operating System Layer.....	5
1.2.2 The Network Layer	5
1.2.3 The Tcp Layer.....	6
1.2.4 The Zone Configuration Layer.....	10
1.2.5 The Zones Layer	12
1.3 Monitoring Applications Executing on Zones.....	23
CONCLUSION	28

Monitoring Solaris Zones

Solaris Containers are part of Sun Microsystems' comprehensive offering of virtualization technologies. These containers provide isolation between software applications or services using flexible, software-defined boundaries. Applications can be managed independently of each other, even while running in the same instance of the Solaris Operating System (Solaris OS). Solaris Containers create an execution environment within a single instance of the Solaris OS and provide:

- a. Full resource containment and control for more predictable service levels
- b. Software fault isolation to minimize fault propagation and unplanned downtime
- c. Security isolation to prevent unauthorized access as well as unintentional intrusions

In a nutshell, a Solaris container is an implementation of **operating system-level virtualization technology**, characterized by:

- Resource and namespace isolation
- A Scalable architecture
- The use of a single operating system, which eliminates redundant administration – the guest operating system is the same as the host operating system, but appears isolated!
- Hardware independence

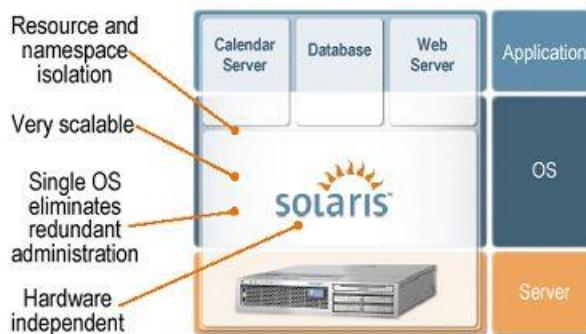


Figure 1: OS Virtualization technology

This is in stark contrast to **VMware's Para-virtualization technique**, where:

- a variety of operating systems can be mounted on the guests, and each guest functions as an independent kernel
- the modified guest operating systems use a special API
- hardware virtualization is performed in addition to OS virtualization

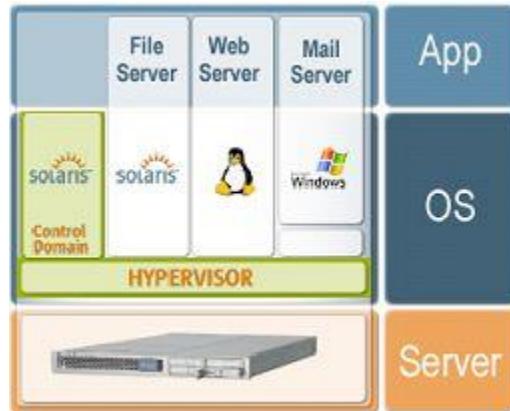


Figure 2: The Para-virtualization technology

1.1 What are Solaris Zones?

Solaris Containers are built using one or more of the following technologies:

- Solaris Resource Manager, for workload resource management;
- Resource Pools, for partitioning, and,
- Zones

A Solaris Zone is a complete execution environment for a set of software services — a separate, virtual Solaris environment within a Solaris instance. A Zone provides a virtual mapping from software services to platform resources, and allows application components to be isolated from each other even though they share a single Solaris OS instance. It establishes boundaries for resource consumption and provides isolation from other Zones on the same system. The boundaries can be changed dynamically to adapt to changing processing requirements of the applications running in the Zone.

The *global* zone encompasses the entire system and is comparable to a normal Solaris OS instance. It has access to the physical hardware and can see and control all processes. The administrator of the global zone can control the system as a whole. The global zone always exists, even when no other zones are configured. Inside the global zone are *local* zones. These zones are isolated from the physical hardware characteristics of the machine by the virtual platform layer. This layer provides the zones with a virtual network interface, one or more file systems and a virtual console. Even though the virtual network interfaces may map to the same physical network interface, applications in different zones are prevented from seeing traffic from applications in other zones. Every zone has its own process environment and runs its own set of core Solaris OS services, including `inetd(1M)`, `syslogd(1M)`, `rpcbind(1M)`, and more. Applications running in a zone are unable to see applications running in other zones because of this private process environment. All zones share the same operating system instance and therefore run the same Solaris OS version.

There are two general zone types to pick from during zone creation. They are:

- **Small zone** - (also known as a "Sparse Root zone")

This is a partial replica of the global zone using a loopback file system (LOFS) to access its shared libraries (see Figure 3). It inherits packages and patches distributed through the global zone.

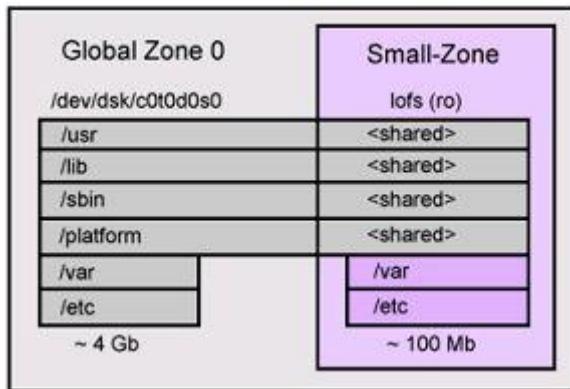


Figure 3: A small zone

- **Big zone** - (also known as a "Whole Root zone")

This is a whole replica of the global zone using its own physical copy of the system's shared libraries. It also inherits the entire package and patch database, and maintains its own physical copy of product contents.

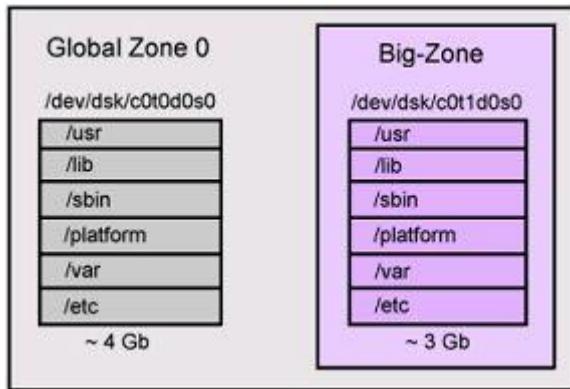


Figure 4: A big zone

Large virtualized Solaris environments are often characterized by a large number of zones (small or big) hosting critical server applications. Since the zones actively share resources with the base OS, any performance deficiency or resource contention at the base could adversely impact the performance of the server applications executing on the zone. It is therefore imperative to closely monitor the functioning of the virtual Solaris host and the zones executing on it from the perspective of resource usage, so that resource-intensive applications or processes can be instantly identified and corrective measures immediately initiated.

1.2 How does eG Enterprise Monitor Solaris Zones?

eG Enterprise currently offers out-of-the-box, an easy way to effectively monitor *Solaris zones*.

Towards this end, eG Enterprise prescribes a specialized *Solaris Virtual* server monitoring model (see Figure 5).

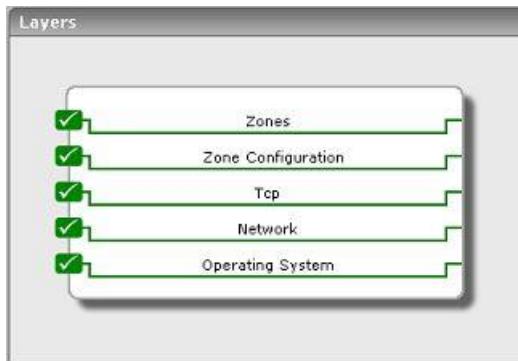


Figure 5: The layer model of a Solaris virtual server

eG Enterprise prescribes only an agent-based mechanism for monitoring a *Solaris virtual* server. In case of the agent-based approach, the **eG agent should be deployed on the global zone**. The procedure for installation is the same as installing an eG agent on a normal Solaris operating system. The steps for the same have been elaborately discussed in the *eG Installation Guide*.

The statistics collected by the eG agent from each of the zones on the *Solaris virtual server* are then presented in the eG monitoring console using the monitoring model of Figure 5 above. Every layer depicted by Figure 5 is associated with a wide variety of tests that measure the health of the Solaris host and each of the zones executing on it. The bottom three layers in Figure 5 represent the status of the base Solaris OS. The **Operating System** layer of the model above monitors the resource usage of the Solaris kernel – particularly with respect to CPU, memory, and disk resources. The **Network** layer monitors the network connectivity of the kernel, and the **Tcp** layer monitors TCP connections to and from the kernel. While the **Zone Configuration** layer in Figure 5 reveals the overall status of the zones that have been configured on the Solaris virtual server, the **Zones** layer reports the relative resource usage levels of every zone.

Since the Solaris virtual environment does not have a full blown OS for each zone, eG Enterprise does not represent two views for a zone (as is the case in a VMware virtual environment). The **Zones** layer represents the relative resource utilization levels across zones as seen by the base OS.

Using the metrics reported by eG Enterprise, administrators can find quick answers to the following questions:

- How many zones are running on each Solaris server, what is the IP address of each of the zones, and when was the zone powered on?
- Does the base operating system have sufficient memory available to support the zones that it is hosting?
- Are all the zones accessible over the network?
- What is the CPU utilization of the base operating system and which of the zones is taking up excessive CPU? Are any processes waiting for I/O utilizing excessive CPU at the host or zone-level?
- Is swap space usage optimal at the host and individual zone-level?
- Which application(s)/process running on each of the zones is taking up CPU, memory, and disk resources?
- Which are the directories to which a zone has full access? Is there sufficient disk space in each of these directories?
- Which of the zones is seeing the highest and lowest network traffic?
- Which is the busiest zone in terms of the number of TCP connections that it handles currently?
- Has any zone been running for a long time?

- What is the total uptime of the zones, and when was a zone booted up?

The sections to come will elaborate on the tests that the single eG agent executes on the Solaris virtual server, and the metrics it reports.

1.2.1 The Operating System Layer

The tests mapped to this layer report the percentage of physical resources used by the Solaris kernel.

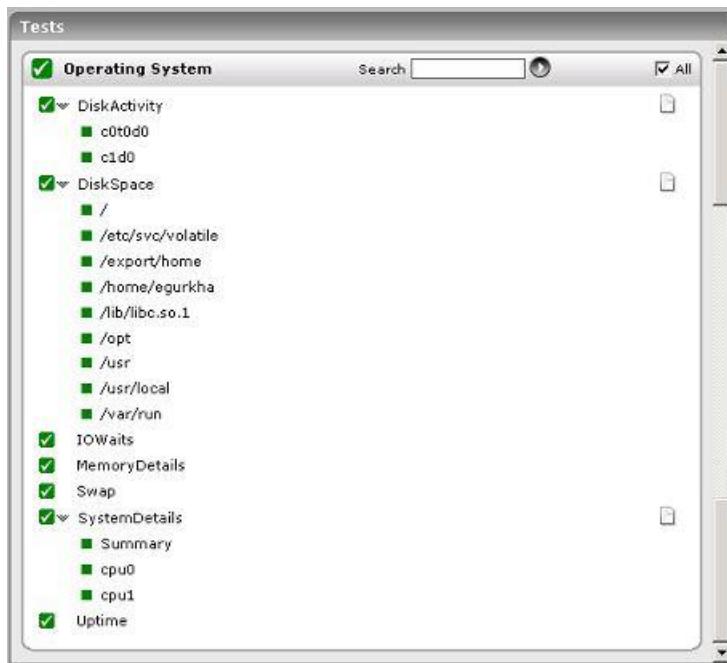


Figure 6: The tests mapped to the Operating System layer

These tests have been elaborately discussed in the *Monitoring Unix and Windows Servers* document. From this layer, administrators can view the resource in use at the base OS.

1.2.2 The Network Layer

The **Network** layer enables administrators to instantaneously detect network connection failures that could render the Solaris kernel inaccessible to users, and any unexpected / unusual increase in traffic to or from the kernel.



Figure 7: The tests mapped to the Network layer

These tests have been elaborately discussed in the *Monitoring Unix and Windows Servers* document.

In addition, the following tests also run on the layer, but are disabled by default.

1.2.2.1 Ethernet Link Aggregation Test

Link aggregation enables one or more network links to be aggregated together to form a link aggregation group. When network links are aggregated, more bandwidth resources become available to the network and load on the network can be uniformly distributed across the links. If one/more links in a group go down, it may limit the bandwidth resources available to the group and may even overload a few other links in the group. To avoid this, administrators must continuously track the state of every link within every aggregation group that is configured. This is where the **Ethernet Link Aggregation** test helps! This test periodically checks the status of each link aggregated within every link aggregation group and instantly alerts administrators when a link goes down or is in an unknown state. This way, the test enables administrators to ensure that all links in an aggregate group are available at all times for providing bandwidth resources and handling the network load.

Purpose	Periodically checks the status of each link aggregated within every link aggregation group and instantly alerts administrators when a link goes down or is in an unknown state
Target of the test	A Solaris virtual server
Agent deploying the test	An internal agent

Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured. 3. PORT – The port number at which the specified HOST listens. By default, the port is NULL. 4. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pexec command instead, select the pexec option from the CONNECT USING drop-down. 5. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>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"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 														
Outputs of the test	One set of results for each link in each link aggregation group														
Measurements made by the test	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center; width: 33.33%;">Measurement</th><th style="text-align: center; width: 33.33%;">Measurement Unit</th><th style="text-align: center; width: 33.33%;">Interpretation</th></tr> </thead> <tbody> <tr> <td style="padding: 5px;"> Link aggregation status: Indicates the current status of this link within this link aggregation group. </td><td style="padding: 5px;"></td><td style="padding: 5px;"> <p>The values that this measure can report and their corresponding numeric values are listed in the table below:</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 50%;">Measure Value</th><th style="width: 50%;">Numeric Value</th></tr> </thead> <tbody> <tr> <td>Up</td><td>0</td></tr> <tr> <td>Down</td><td>1</td></tr> <tr> <td>Unknown</td><td>2</td></tr> </tbody> </table> <p>Note:</p> <p>By default, this measure reports the Measure Values in the table above to indicate link status. In the graph of this measure however, the same is represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this test reveals the speed, operational state, and duplex state of a link.</p> </td></tr> </tbody> </table>	Measurement	Measurement Unit	Interpretation	Link aggregation status: Indicates the current status of this link within this link aggregation group.		<p>The values that this measure can report and their corresponding numeric values are listed in the table below:</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 50%;">Measure Value</th><th style="width: 50%;">Numeric Value</th></tr> </thead> <tbody> <tr> <td>Up</td><td>0</td></tr> <tr> <td>Down</td><td>1</td></tr> <tr> <td>Unknown</td><td>2</td></tr> </tbody> </table> <p>Note:</p> <p>By default, this measure reports the Measure Values in the table above to indicate link status. In the graph of this measure however, the same is represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this test reveals the speed, operational state, and duplex state of a link.</p>	Measure Value	Numeric Value	Up	0	Down	1	Unknown	2
Measurement	Measurement Unit	Interpretation													
Link aggregation status: Indicates the current status of this link within this link aggregation group.		<p>The values that this measure can report and their corresponding numeric values are listed in the table below:</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 50%;">Measure Value</th><th style="width: 50%;">Numeric Value</th></tr> </thead> <tbody> <tr> <td>Up</td><td>0</td></tr> <tr> <td>Down</td><td>1</td></tr> <tr> <td>Unknown</td><td>2</td></tr> </tbody> </table> <p>Note:</p> <p>By default, this measure reports the Measure Values in the table above to indicate link status. In the graph of this measure however, the same is represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this test reveals the speed, operational state, and duplex state of a link.</p>	Measure Value	Numeric Value	Up	0	Down	1	Unknown	2					
Measure Value	Numeric Value														
Up	0														
Down	1														
Unknown	2														

1.2.2.2 Ethernet Links Test

If a user complains that he/she is unable to connect to the server over the network, administrators will want to know whether any Ethernet link to the server is down presently. This can be determined using the **Ethernet Links** test. This test reports the state of every Ethernet link and notifies administrators if a link goes down.

Purpose	Reports the state of every Ethernet link and notifies administrators if a link goes down		
Target of the test	A Solaris virtual server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>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"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 		
Outputs of the test	One set of results for each Ethernet link		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	<p>Link status: Indicates the current status of this link.</p>		<p>The values that this measure can report and their corresponding numeric values are listed in the table below:</p> <table border="1" data-bbox="992 354 1465 551"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Up</td><td>0</td></tr> <tr> <td>Down</td><td>1</td></tr> <tr> <td>Unknown</td><td>2</td></tr> </tbody> </table> <p>Note: By default, this measure reports the Measure Values in the table above to indicate link status. In the graph of this measure however, the same is represented using the corresponding numeric equivalents only.</p> <p>The detailed diagnosis of this test reveals the speed and duplex state of a link.</p>	Measure Value	Numeric Value	Up	0	Down	1	Unknown	2
Measure Value	Numeric Value										
Up	0										
Down	1										
Unknown	2										

1.2.3 The Tcp Layer

To track the health of the TCP connects to and from Solaris kernel, the eG Enterprise suite uses the tests associated with the **Tcp** layer.



Figure 8: The tests associated with the Tcp layer

The metrics reported by this test have been discussed elaborately the *Monitoring Unix and Windows Servers* document.

1.2.4 The Zone Configuration Layer

Using the **ZoneStatus** test mapped to it, the **Zone Configuration** layer tracks the overall status of the zones.



Figure 9: The tests mapped to the Zone Configuration layer

1.2.4.1 ZoneStatus Test

This test tracks the overall status of the zones by reporting how many zones have been configured on the Solaris virtual server, how many have been registered, how many were newly added/removed, etc.

Purpose	Monitors the status of the zones configured on a Solaris virtual server
Target of the test	A Solaris virtual server
Agent deploying the test	An internal agent
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>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"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

MONITORING SOLARIS ZONES

Outputs of the test	One set of results for the Solaris virtual server being monitored		
Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Registered zones: Indicates the total number of zones that have been registered with the Solaris virtual server.	Number	To view the list of registered zone, click on the DIAGNOSIS icon against this measure.
	Added zones: Indicates the number of zones that were newly added with the Solaris virtual server during the last measurement period.	Number	
	Removed zones: Indicates the number of zones that were removed from the Solaris virtual server during the last measurement period.	Number	The detailed diagnosis of this measure, if enabled, provides the details of the removed zones.
	Running zones: Indicates the number of zones that are currently running on the Solaris virtual server.	Number	The detailed diagnosis of this measure, if enabled, provides the details of the zones that are running.
	Configured zones: Indicates the number of zones that are currently configured with the Solaris virtual server.	Number	The detailed diagnosis of this measure, if enabled, provides the details of the configured zones.
	Installed_zones: Indicates the number of zones that are currently installed on the Solaris virtual server.		The detailed diagnosis of this measure, if enabled, provides the details of the installed zones.

The detailed diagnosis of the **Running zones** measure, if enabled, provides the IP addresses of the zones, the zone names, and the version of Solaris they are executing on (see Figure 10).

MONITORING SOLARIS ZONES

Figure 10: The detailed diagnosis of the Running zones measure

The detailed diagnosis of the **Configured zones** measure, if enabled, provides the IP addresses of the zones, the zone names, and the version of Solaris they are executing on (see Figure 11).

Figure 11: The detailed diagnosis of the Configured zones measure

The detailed diagnosis of the **Installed zones** measure, if enabled, provides the IP addresses of the zones, the zone names, and the version of Solaris they are executing on (see Figure 12).

Figure 12: The detailed diagnosis of the Installed zones measure

1.2.5 The Zones Layer

The **Zones** layer provides the host operating system's view of the resource usage levels of each of the zones hosted on it. Using the information reported by this test, administrators can:

MONITORING SOLARIS ZONES

- Determine which of the zones is taking up more resources (CPU, memory, network, or disk) than the others. This information can help with load balancing or capacity planning. For example, if one of the zones is receiving a very high rate of requests compared to the others, this zone may be a candidate for migration to another host, so as to minimize the impact it has on the other zones on the host.
- Determine times when sudden or steady spikes in the physical resource utilization are caused by the zones

By default, clicking on the **Zones** layer invokes the **Zone View** depicted by Figure 13. Using this view, you can quickly determine the health of each of the zones (including *global* zones) configured on the monitored Solaris host.



Figure 13: The Zone view

To zoom into the performance of a particular zone, click on any zone in Figure 13. This will lead you to Figure 14, which displays all the metrics that the eG agent collected from that zone. You are thus enabled to cross-correlate across the various metrics, and quickly detect the root-cause of current/probable disturbances to the internal health of a guest. To view the time-of-day variations in a measure, simply click on a measure in Figure 14.

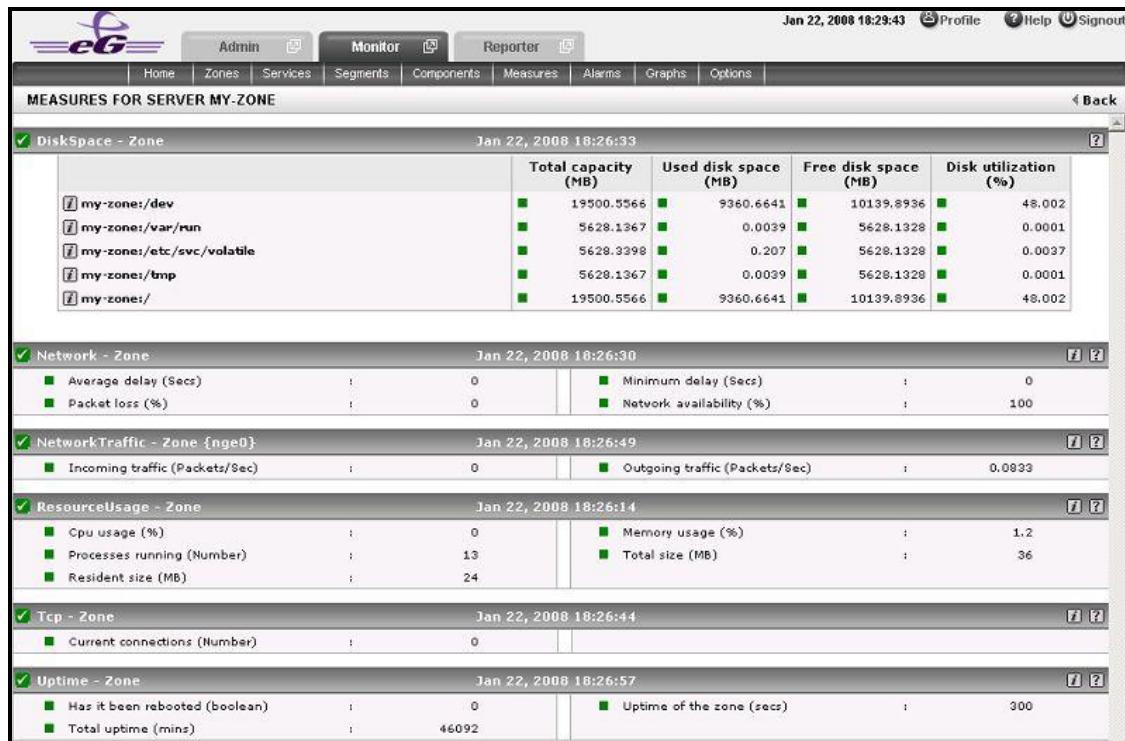


Figure 14: The measures pertaining to a zone

To view real-time graphs of pre-configured measures (pertaining to the Solaris host and the zones operating on it), click on the **LIVE GRAPH** link in Figure 13. The graph display that appears subsequently (see Figure 15) has been

organized in such a way that next to every host-pertinent measure graph, the closely related zone-specific measure graph appears. For instance, next to the graph of the 'Cpu usage' measure of the **Resource – Zone Test**, you will find a graph of the ' CPU utilization' measure of the **Processor** test. This way, you can easily compare and correlate how well the physical CPU resources are being utilized by both the Solaris kernel processes and those that are executing on the zones. On the basis of this analysis, you can proactively isolate potential performance issues, and also determine the root-cause of the issue - is it the host? or is it the zone? If you access this page from the **LIVE GRAPH** link in Figure 13, then, by default, you will view live graphs pertaining to the *Solaris Virtual* server. However, you can select a different virtualized component-type and a different virtualized component using the **Component Type** and **Component** lists (respectively) in Figure 15.

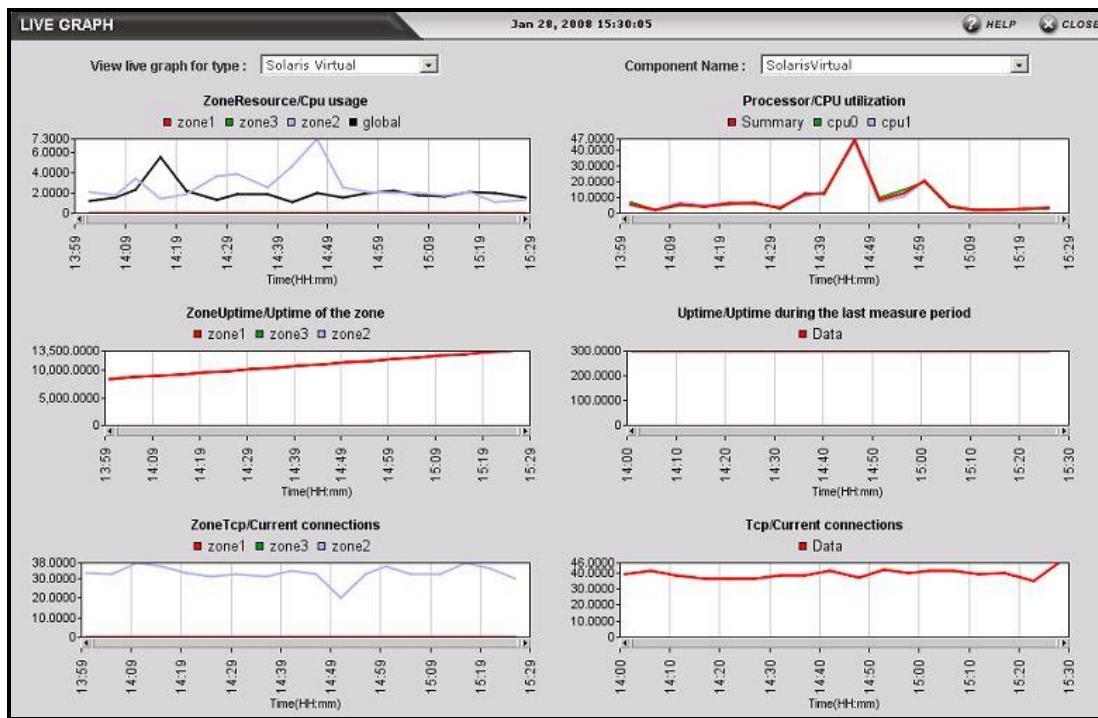


Figure 15: Live graph of Solaris Virtual server

If you prefer to view the tests associated with the **Zones** layer instead of zone-specific measures or live graphs, then, simply click on the **COMPONENT LAYERS** link in Figure 13. This will lead you back to the layer model page, wherein you can view the complete list of tests associated with the **Zones** layer (see Figure 16).

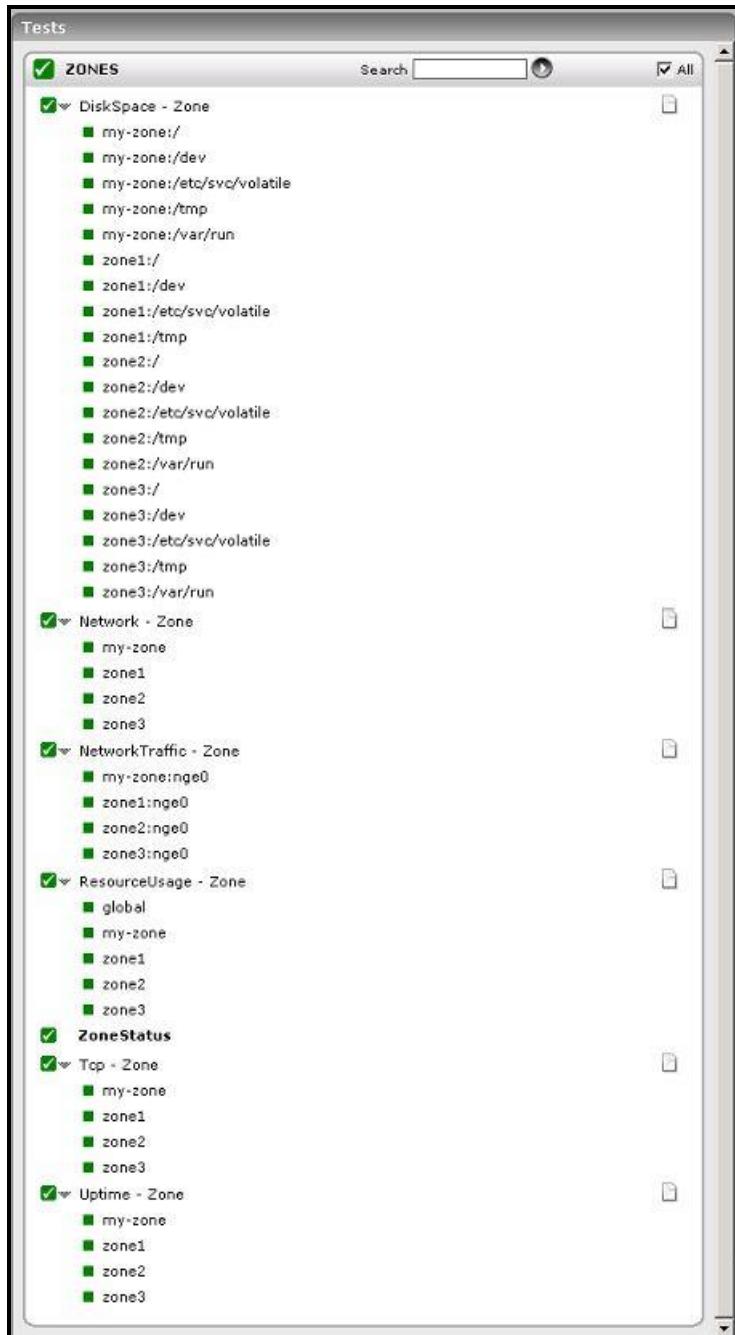


Figure 16: The tests mapped to the Zones layer

1.2.5.1 DiskSpace – Zone Test

This test measures the space usage of every disk partition on each zone of a Solaris virtual server.

Purpose	To measure the space usage of every disk partition on a host
Target of the test	Any host system

Agent deploying the test	An internal agent		
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. 		
Outputs of the test	One set of results for each disk partition on every zone monitored		
Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Total capacity: Indicates the total capacity of a disk partition.	MB	
	Used space: Indicates the amount of space used in a disk partition.	MB	
	Free space: Indicates the current free space available for each disk partition of a zone.	MB	
	Percent usage: Indicates the percentage of space usage on each disk partition of a system.	Percent	In case of drives that the small zones share with the base operating system, the value of this measure will be the same as that of the base. This is because, small zones have only 'read-only' access to the shared drives. In case of drives/directories to which the zone has full access, a value close to 100% indicates very high space usage. Due to this problem condition, applications executing on the zone may not be able to write data to the disk partition(s).

1.2.5.2 Network – Zone Test

This test monitors the network status of each zone on a Solaris virtual server.

Purpose	Monitors the network status of each zone on a Solaris virtual server			
Target of the test	A Solaris virtual server			
Agent deploying the test	An internal agent			
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. PACKETSIZE - The size of packets used for the test (in bytes) PACKETCOUNT – The number of packets to be transmitted during the test TIMEOUT - How long after transmission should a packet be deemed lost (in seconds) PACKETINTERVAL - Represents the interval (in milliseconds) between successive packet transmissions during the execution of the network test for a specific target. TARGETS - In the TARGETS text box, specify a comma-separated list of <i>name:IP address</i> pairs. While the <i>name</i> is just a display name, the <i>IP address</i> refers to the IP to be monitored. This specification will ensure that the test pings multiple IP addresses. For example - <i>mysql:192.168.0.102,egwebsite:209.15.165.127</i> CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. 			
Outputs of the test	One set of results for every zone monitored			
Measurements made by the test	<th>Measurement</th> <th>Measurement Unit</th> <th>Interpretation</th>	Measurement	Measurement Unit	Interpretation
	Average delay: Indicates the average delay between transmission of packet to a zone and receipt of the response to the packet at the source.	Secs	An increase in network latency could result from misconfiguration of the router(s) along the path, network congestion, retransmissions at the network, etc.	

	Minimum delay: Indicates the minimum time between transmission of a packet and receipt of the response back.	Secs	A significant increase in the minimum round-trip time is often a sure sign of network congestion.
	Packet loss: Indicates the percentage of packets lost during transmission from source to target and back.	Percent	Packet loss is often caused by network buffer overflows at a network router or by packet corruptions over the network.
	Availability: Indicates the availability of the network connection.	Percent	A value of 100 indicates that the zone is connected. The value 0 indicates that the zone is not connected. Typically, the value 100 corresponds to a packet loss of 0.

1.2.5.3 Network Traffic – Zone Test

This test measures the incoming and outgoing traffic through each zone of a Solaris virtual server.

Purpose	Measures the incoming and outgoing traffic through each zone of a Solaris virtual			
Target of the test	A Solaris virtual server			
Agent deploying the test	An internal agent			
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. 			
Outputs of the test	One set of results for every zone executing on the Solaris virtual server			
Measurements made by the	<table border="1"> <thead> <tr> <th>Measurement</th> <th>Measurement Unit</th> <th>Interpretation</th> </tr> </thead> </table>	Measurement	Measurement Unit	Interpretation
Measurement	Measurement Unit	Interpretation		

test	Incoming traffic: Indicates the rate at which data (including framing characters) is received on a network interface.	Packets/Sec	An abnormally high rate of incoming traffic may require additional analysis.
	Outgoing traffic: Represents the rate at which data (including framing characters) is sent on a network interface.	Packets/Sec	An abnormally high rate of incoming traffic may require additional analysis.

1.2.5.4 Resource Usage – Zone Test

This test monitors the CPU and memory usage of each zone (including Global zone) of a Solaris virtual server.

Purpose	Monitors the CPU and memory usage of each zone (including Global zone) of a Solaris virtual server		
Target of the test	A Solaris virtual server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. 		
Outputs of the test	One set of results for every zone executing on the Solaris virtual server		
Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Processes running: Indicates the number of processes currently executing on this zone.	Number	

	Cpu usage: Indicates the percentage of CPU used by all the processes executing on this zone.	Percent	A very high value could indicate the existence of one/more CPU-intensive processes. Using the detailed diagnosis of this measure, you can quickly and accurately identify the top-10 processes that are consuming CPU excessively.
	Memory usage: Indicates the percentage of memory utilized by the processes executing on this zone.	Percent	A sudden increase in memory utilization for processes may be indicative of memory leaks in the application.
	Total size: Indicates the total virtual memory size of mappings within the processes executing on the zone.	MB	
	Resident size: Indicates the resident set size of the processes executing on the zone - i.e., the amount of physical memory mapped into the processes.	MB	

The detailed diagnosis of the **Cpu usage** measure reveals the top-10 CPU-intensive processes executing on a zone.

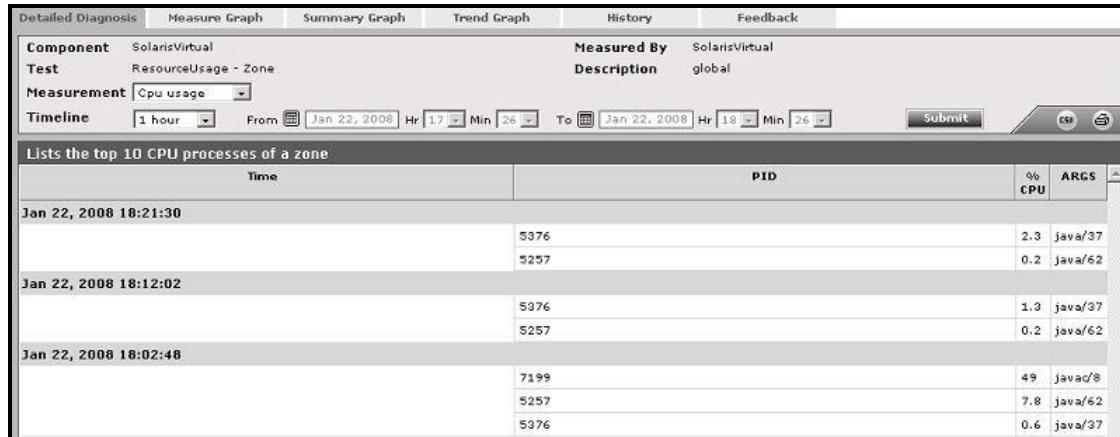


Figure 17: The detailed diagnosis of the Cpu usage measure

The detailed diagnosis of the **Memory usage** measure reveals the top-10 memory-intensive processes executing on a zone.

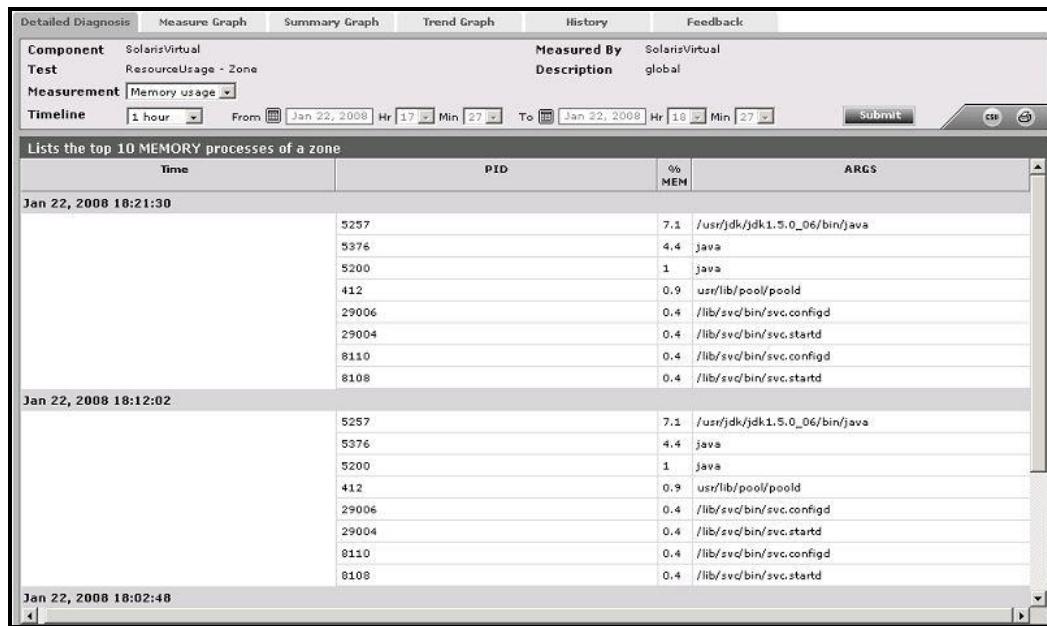


Figure 18: The detailed diagnosis of the Memory usage measure

1.2.5.5 Tcp – Zone Test

This test measures the health of the TCP layer of every zone of a Solaris virtual server.

Purpose	Measures the health of the TCP layer of every zone of a Solaris virtual server		
Target of the test	A Solaris virtual server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT – The port number at which the specified HOST listens. By default, the port is NULL. CONNECT USING – By default, this parameter is set to sudo. This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. 		
Outputs of the test	One set of results for every zone executing on the Solaris virtual server		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	Current connections: Indicates the currently established connections to the zone.	Number	A sudden increase in the number of connections established on a zone can indicate either an increase in load to one or more of the applications executing on the zone, or that one or more of the applications are experiencing a problem (e.g., a slow down).
------	---	--------	--

1.2.5.6 Uptime – Zone Test

This test monitors the uptime of every zone on a Solaris virtual server.

Purpose	Monitors the uptime of every zone on a Solaris virtual server		
Target of the test	A Solaris virtual server		
Agent deploying the test	An internal 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. PORT – The port number at which the specified HOST listens. By default, the port is NULL. 4. CONNECT USING – By default, this parameter is set to sudo . This indicates that, by default, this test uses the sudo command to connect to the Solaris virtual server and pull out metrics. If you want the test to use the pfexec command instead, select the pfexec option from the CONNECT USING drop-down. 5. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1 . This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY .		
Outputs of the test	One set of results for every zone executing on the Solaris virtual server		
Measurements made by the test	Measurement Has system been rebooted: Indicates whether the zone has been rebooted during the last measurement period or not.	Measurement Unit Boolean	Interpretation If this measure shows 1, it means that the zone was rebooted during the last measurement period. By checking the time periods when this metric changes from 0 to 1, an administrator can determine the times when this zone was rebooted.

	Uptime during the last measure period: Indicates the time period that the zone has been up since the last time this test ran.	Secs	If the zone has not been rebooted during the last measurement period and the agent has been running continuously, this value will be equal to the measurement period. If the zone was rebooted during the last measurement period, this value will be less than the measurement period of the test. For example, if the measurement period is 300 secs, and if the zone was rebooted 120 secs back, this metric will report a value of 120 seconds. The accuracy of this metric is dependent on the measurement period - the smaller the measurement period, greater the accuracy.
	Total uptime of the system: Indicates the total time that the zone has been up since its last reboot.	Mins	Administrators may wish to be alerted if a zone has been running without a reboot for a very long period. Setting a threshold for this metric allows administrators to determine such conditions.

Note:

If a value less than a minute is configured as the **TEST PERIOD** of the Uptime - Zone test, then, the **Uptime during the last measure period** measure will report the value 0 for the zones until the minute boundary is crossed. For instance, if you configure the Uptime - Zone test to run every 10 seconds, then, for the first 5 test execution cycles (i.e., $10 \times 5 = 50$ seconds), the **Uptime during the last measure period** measure will report the value 0 for all zones; however, the sixth time the test executes (i.e., when test execution touches the 1 minute boundary), this measure will report the value 60 seconds for the zones. Thereafter, every sixth measurement period will report 60 seconds as the uptime of the zones. This is because, Unix-based operating systems report uptime only in minutes and not in seconds.

1.3 Monitoring Applications Executing on Zones

An eG agent can be installed on a Solaris zone to monitor specific applications executing on it. For instance, say that an Oracle server is executing in a Solaris zone. You can deploy an eG agent on the zone and configure it to monitor the Oracle server. Figure 19 depicts the layer model, tests, and measurements of such an Oracle database server.

MONITORING SOLARIS ZONES

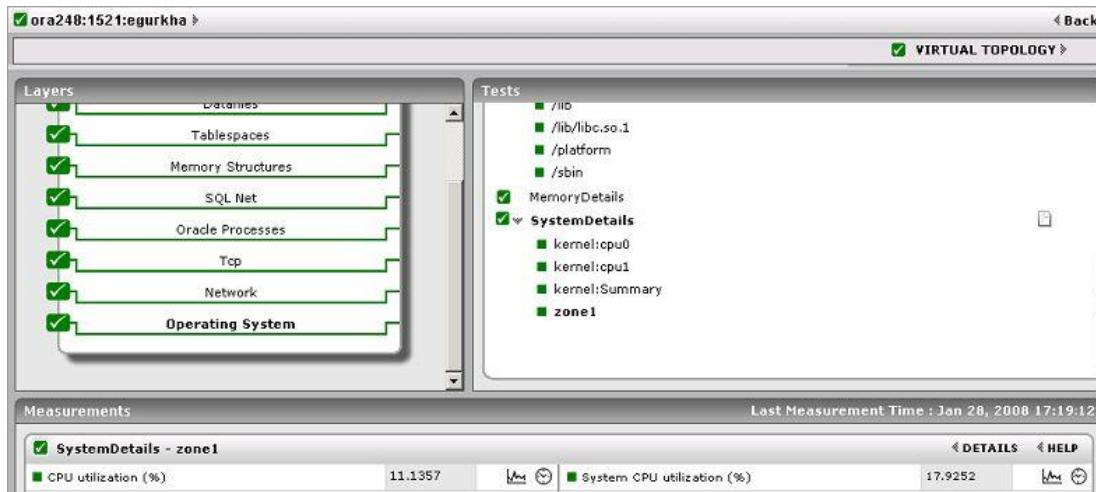


Figure 19: The layer model, tests, and measurements of the Oracle server executing on a Solaris zone

Note that eG Enterprise does not differentiate between applications executing in virtualized and non-virtualized Solaris environments, while reporting metrics for the application-specific layers. In the case of our example therefore, the top 6 layers of the layer model in Figure 19, will report the same metrics that the eG agent typically extracts from any Oracle server executing on a normal Solaris host.

The difference however lies in the resource-usage statistics reported by the operating system-specific layers – i.e., the bottom 3 layers of Figure 19 in our example. In a virtualized Solaris environment, both the zones and the Solaris kernel at the base share resources from a common resource pool. Therefore, while monitoring the CPU usage of an application that is executing on a Solaris zone, the eG agent reports the percentage of the total CPU resources that the zone has utilized, and also the percentage CPU consumption of each of the processors of the Solaris kernel at the base. Accordingly, you will find that the **SystemDetails** test mapped to the **Operating System** layer of Figure 95 reports one set of CPU usage metrics each for every *kernel* processor (in the format, *kernel:<Processor>*), and one set of metrics for the zone *zone1*. This way, whenever there is a spike in resource utilization, administrators can easily figure out what has contributed to the excessive usage – CPU-intensive processes executing on the zone? or one/more runaway processes executing on the Solaris kernel?

Similarly, the **Network** layer of the application will report network traffic statistics and network interface-related metrics (if available) for the kernel and the zone as well, thereby enabling you to accurately figure out where the traffic is more. Likewise, the **TCP** layer will report the TCP connections currently established with the kernel and the zone, so as to help you instantly determine the busier of the two entities.

Also, since the Oracle database server in our example is running on a virtual machine that is hosted on a Solaris virtual server, a link for **VIRTUAL TOPOLOGY** appears to the top of the right hand panel in Figure 19. This link is pre-fixed by a state indicator, which indicates whether/not there is a problem with the virtualized environment that this application is part of. In Figure 19, the **VIRTUAL TOPOLOGY** link indicates that the virtualized environment that the Oracle database server belongs to is experiencing no problems.

Note:

- If the current monitor user is not authorized to monitor the Solaris virtual server, but is associated with one/more of the server applications executing on the managed Solaris virtual server, then, while viewing the layer model of any of these server applications, the **VIRTUAL TOPOLOGY** link will **not appear** in the layer model page depicted by Figure 19.
- The **VIRTUAL TOPOLOGY** link also appears in the user interface that displays the layer model of a managed Solaris virtual server.

Clicking on the **VIRTUAL TOPOLOGY** link reveals Figure 20, which depicts the association between the Solaris virtual server and each of the server applications executing on it. eG Enterprise is able to automatically determine the mapping of applications to the Solaris virtual server.

Whether eG Enterprise automatically determines the mapping of applications to Solaris virtual servers or not is determined by the value of the **AutoVirtualMapping** variable in the **[MISC]** section of the **eg_external.ini** configuration file in the **<EG_INSTALL_DIR>\manager\config** directory of the eG manager. If the value of this variable is **true**, the eG manager auto-discovers the applications to Solaris virtual servers mapping.

Note:

- For **AutoVirtualMapping** to work, the detailed diagnosis frequencies set globally (i.e., using the Configure -> Diagnosis menu sequence) should not be set to *0:0*.
- As long as the **Identify agents only using nick names** flag in the **MANAGER SETTINGS** page of the eG administrative interface (Configure -> Manager Settings menu sequence) is **Yes** (which is the default), eG Enterprise can automatically identify the server applications executing on an ESX host, using the host/nick names that are mapped to the IP addresses discovered on the host. If the **Identify agents only using nick names** flag is set to **No** instead, then make sure that, while managing a server application executing in a virtualized environment, the hostname of the virtual machine is specified as the nick name of the corresponding server application. If more than one server application is executing on the same virtual machine, then any one of those server applications should have the virtual machine name as its nick name.



To disable auto-discovery, set this value to **false**. In such a case, once a *Solaris Virtual* server is added, when adding any new server application using the eG administrative interface, you will be prompted to manually set an association between the server application being added and the Solaris virtual server. This mapping (determined automatically or indicated manually) is used by eG Enterprise for correlation – e.g., since the application runs on a Solaris zone, which shares resources with the base operating system, it is most likely that a problem with the Solaris virtual server (at the base) will impact the performance of the application running on one of its zones.

Note:

The topology representation of Figure 20 will differ according to the monitoring rights of the current user. If the user has been assigned only some of the server applications executing on the managed Solaris virtual server, then Figure 20 that appears will depict the association between the Solaris virtual server and the assigned server applications only.

Apparently, besides the Oracle server in our example, a Web server, and a WebLogic server are also executing on virtual machines configured on the same Solaris virtual server.

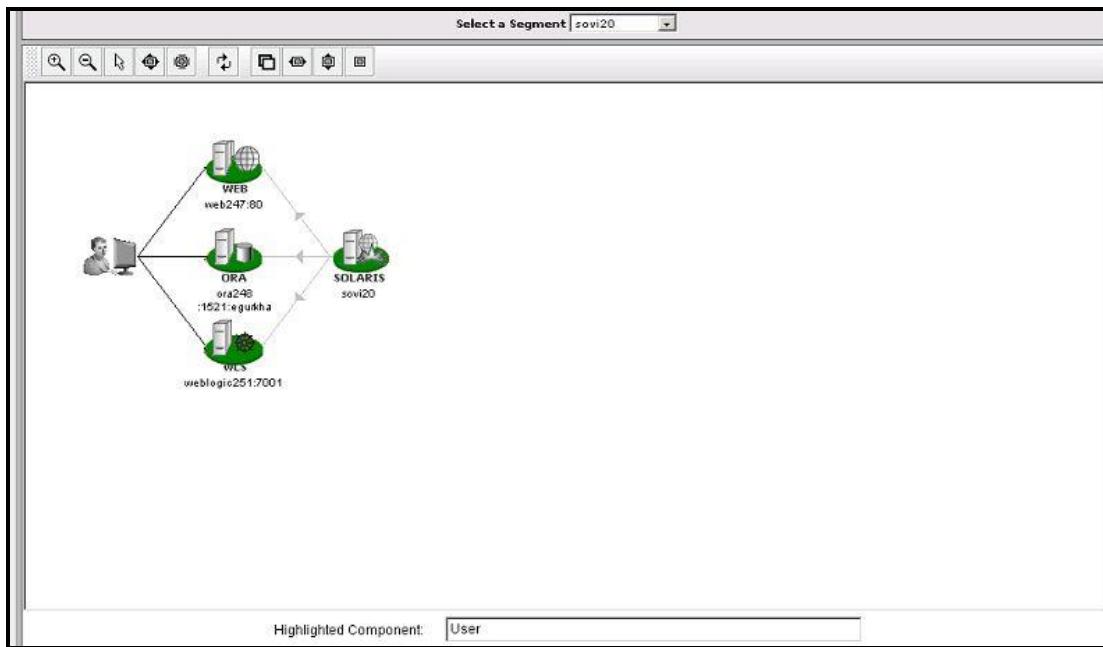


Figure 20: The association between the Solaris host and server applications executing on its zones

Whenever the virtualized environment is faced with performance issues, administrators can easily and accurately identify the root-cause of the issue – is it the base Solaris operating system? Or any of the configured Solaris zones?

- using the direction of dependency between the various components in the topology representation, and the color coding of the components. Clicking on the problem source so identified, allows the administrator to manually drill down to the layer model of the problem component, revealing exactly what the problem is.

Besides this manual exercise, the eG Enterprise system also automatically performs the correlation – between applications, zones, and the kernel - in real-time. As and when a new alert is detected, and a problem is narrowed down to one application, the eG Enterprise manager assesses the zone to Solaris virtual server mapping to determine which Solaris virtual server the impacted application is running on. The solution then reviews the real-time state of the other zones running on the same host, and the performance of the base operating system itself (using the dependency graph shown in Figure 20) to determine where the real problem lies.

a.

b.

Conclusion

eG Enterprise, with its ability to provide in-depth insight into the performance of virtualized Solaris infrastructures, is the ideal solution for monitoring such environments. For more information on eG Enterprise, please visit our web site at www.eginnovations.com or write to us at sales@eginnovations.com.