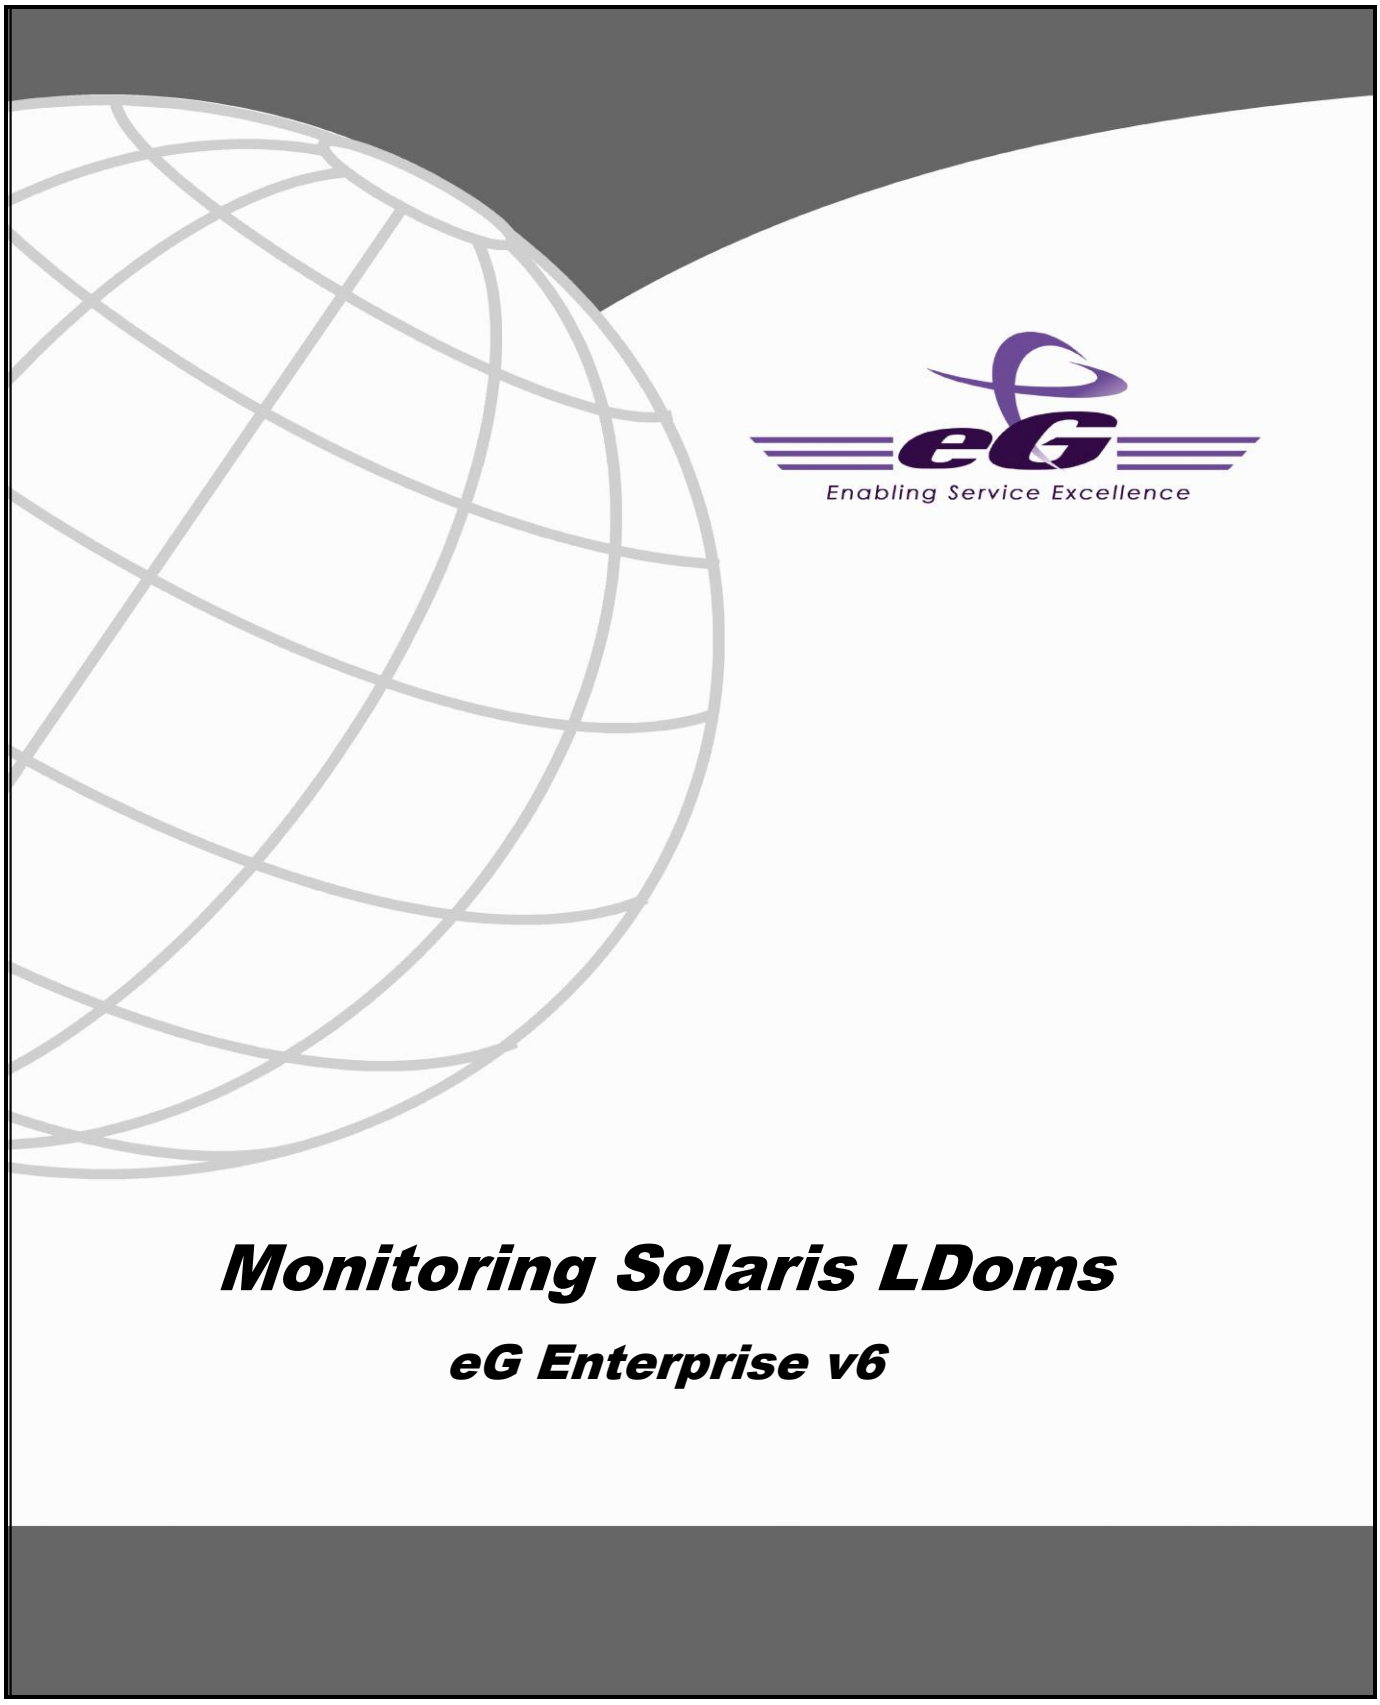

Monitoring Solaris LDoms

eG Enterprise v6

Table of Contents

MONITORING SOLARIS LDOMS 1

1.1 How eG Enterprise Monitors Solaris LDoms? 2

1.2 Installing the eG Agent 3

1.3 Pre-requisites for Monitoring a Solaris LDoms Server 3

1.4 The Solaris LDoms Server Monitoring Model 5

1.4.1 The Operating System Layer 5

1.4.2 The Network Layer 13

1.4.3 The LDoms Guests Layer 13

1.4.4 The Inside View of VMs Layer 18

1.5 Correlation between Applications in a Solaris LDoms Virtualized Environment 42

1.6 Troubleshooting Solaris LDoms Monitoring 43

CONCLUSION 45

Chapter 1

Monitoring Solaris LDoms

LDoms is a new server virtualization and partitioning technology from Sun for CoolThreads systems. By deeply integrating the industry-leading multithreading capability of the UltraSPARC® T1 processor and the Solaris 10 OS, the LDoms (or Logical Domains) technology allows you to allocate a system's various resources, such as memory, CPUs, and devices, into logical groupings and create multiple, discrete systems, each with their own operating system, resources, and identity within a single computer system.

Figure 1 depicts the architecture of a Solaris LDoms server.

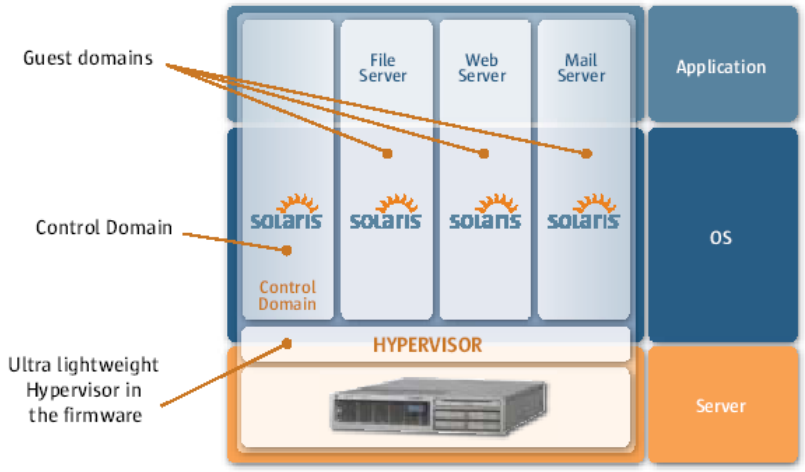

Figure 1: The architecture of a Solaris LDoms server

The hypervisor software/firmware of Figure 1 is responsible for maintaining separation (eg: visible hardware parts) between domains, and provides logical domain channels (LDCs) using which the domains communicate with each other. The control domain (see Figure 1) is the configuration platform for managing domains, and allows monitoring and reconfiguration of domains. The guest domains are where the applications are hosted, and can be independently powered on/off without affecting other domains. **Note that Solaris LDoms support only Solaris and Linux guests**.

As each of the guest domains can function like a full-blown operating system, this virtualization technology helps conserve space, reduce the physical size, and ease the management of large server-sprawls, typically characterized by hundreds of servers. How? Where ten physical hosts are required, a single Solaris LDoms server can now be used; ten guest domains can then be configured on the single Solaris LDoms server, with each guest playing host to a different operating system and a different application.

The applications on the logical domains use the virtual CPU/disk/memory resources that the hypervisor allocates to every guest. Typically, the hypervisor allocates subsets of the overall CPU, memory, and I/O resources of the server to a given logical domain. This implies that a resource contention at the host might impact the resource allocation to the guests, thereby affecting the performance of the applications executing on the guests. In a few other cases, the base Solaris host might be sized right, but the administrator might have improperly configured the size of a logical domain. A resource-intensive application executing on the domain would then leave a very high resource foot-print on the allocated resources, but would not affect the physical resource availability. Therefore, whenever an application executing on a domain experiences a slowdown, administrators might first need to determine where the root-cause of the problem lies – with the Solaris host or with the guest domain? To determine this, administrators need to continuously monitor the resource utilization of the Solaris operating system at the base, the control domain, and the logical domains on the Solaris server, so that they can proactively detect resource shortages or resource usage excesses, accurately indicate where the resource crunch has occurred, and help prevent adversities.

1.1 How eG Enterprise Monitors Solaris LDoms?

eG Enterprise presents a specialized *Solaris LDoms* (see Figure 2) monitoring model that proposes a patented 'In-N-Out' approach to monitoring Solaris LDoms. In this approach, the eG agent is deployed on the *control or primary domain* of the Solaris LDoms server, and is configured to extract critical resource-usage metrics pertaining to the Solaris host, primary domain, and guest domains. The metrics so collected reveal the following:

- The resource usage of the physical server – i.e., the Solaris host and the control domain
- The resource usage of each of the logical domains in relation to the physical resources available at the host-level (i.e., the "outside" view)
- The usage of allocated resources (i.e., virtual resources) by the guests (i.e., the "inside" view)

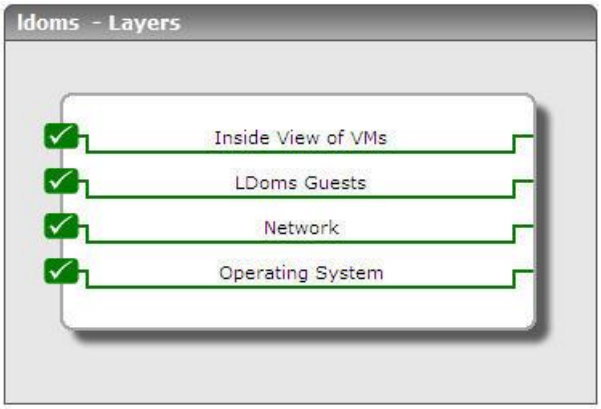

Figure 2: The layer model of a Solaris LDoms server

Each layer depicted by Figure 2 above reports a wide variety of statistics, which enable administrators to find quick and accurate answers for the following questions:

- Is the Solaris server available over the network? If so, how quickly is it responding to requests?
- Are any resource-intensive processes executing on the Solaris host? Does the host have adequate physical CPU and memory resources to support the guest domains configured on it?
- What percentage of the physical resources have been allocated to the guest domains?
- How effectively is the primary domain utilizing the allocated CPU and memory resources? Are any CPU/memory-intensive processes executing on the primary domain? If so, what are they?
- Are the primary domain's processes draining the physical CPU resources?
- How many logical domains are running, what is the IP address of each guest domain, which operating system is each guest domain running on, and when was the guest domain powered on?
- Which guest is consuming the physical CPU resources excessively?
- Have sufficient CPU and memory resources been allocated to the logical domains?
- Is any guest consuming too much CPU/memory/disk space?
- Which is the busiest guest in terms of incoming TCP connections?

1.2 Installing the eG Agent

To monitor a Solaris LDoms server inside-out, the eG agent should be installed on the *control domain*. The control domain is a Solaris-variant, and hence, you need to install a Solaris agent on it. The installation procedure is the same as that of a normal Solaris agent. For a detailed installation procedure, refer to the *eG Installation Guide*.

Note:

While configuring the eG agent using the **setup_agent** script, ensure that the script is executed only by a user who belongs to the **root** group.

1.3 Pre-requisites for Monitoring a Solaris LDoms Server

There are several pre-requisites for an eG agent to be able to monitor a Solaris LDoms server and the guest VMs hosted on it.

- Make sure that the firewall on the Solaris LDoms server is configured to allow outbound traffic from the eG agent to the eG management console. The port used for this communication is determined at the time the eG manager and agents are installed in your environment; port 7077 is the default. To configure the agent-manager communication, do the following:
  - Edit the **ipf.conf** file in the **/etc/ipf** directory on the Solaris LDoms server.
  - The sample contents of the file are as given below:

```
# Allow Outbound Traffic:
pass out quick proto tcp from any to any flags S keep state
pass out quick proto udp from any to any keep state
pass out quick proto icmp from any to any keep state
#
# SSHD - Allow Inbound from All Addresses to Port 22:
pass in quick proto tcp from any to any port = 22
```

```
#
# Block without Logging Inbound Ports 137, 138, 139, and 631:
block in quick proto udp from any to any port = 137
block in quick proto tcp from any to any port = 137
block in quick proto udp from any to any port = 138
block in quick proto tcp from any to any port = 138
block in quick proto udp from any to any port = 139
block in quick proto tcp from any to any port = 139
block in quick proto tcp/udp from any to any port = 631
block in quick proto http from any to any port = 7077
#
# Block and Log Everything Else Inbound (comment first line below to allow ping
requests):
block in log first proto icmp from any to any
block in log first proto tcp/udp from any to any
```

- If the firewall has been configured to block the port 7077, then an entry to that effect will be available in the **ipf.conf** file. For instance, in the sample contents above, the entry in **Bold** indicates that the port 7077 is currently blocked.
- To enable eG agent-manager communication, simply comment the entry in **Bold** by prefixing it with a #, as indicated below:

```
# Allow Outbound Traffic:
pass out quick proto tcp from any to any flags S keep state
pass out quick proto udp from any to any keep state
pass out quick proto icmp from any to any keep state
#
# SSHD - Allow Inbound from All Addresses to Port 22:
pass in quick proto tcp from any to any port = 22
#
# Block without Logging Inbound Ports 137, 138, 139, and 631:
block in quick proto udp from any to any port = 137
block in quick proto tcp from any to any port = 137
block in quick proto udp from any to any port = 138
block in quick proto tcp from any to any port = 138
block in quick proto udp from any to any port = 139
block in quick proto tcp from any to any port = 139
block in quick proto tcp/udp from any to any port = 631
# block in quick proto http from any to any port = 7077
#
# Block and Log Everything Else Inbound (comment first line below to allow ping
requests):
block in log first proto icmp from any to any
block in log first proto tcp/udp from any to any
```

- Then, save the **ipf.conf** file.
- To make sure that the changes take effect, execute the following command:

  svcadm restart network/ipfilter

- Also, ensure that the Solaris LDoms firewall allows the eG agent on the server to communicate with the Linux and Solaris guests of the server using SSH. If the firewall has blocked the SSH port (default is 22), then make sure that it is opened using the procedure discussed above.
- Ensure that SSH is enabled on all Linux and Solaris guests to be monitored.

- To enable the eG agent to communicate with the guest domains, an administrative account login and password must be provided when configuring the eG monitoring capabilities; different logins can be provided for different logical domains on the same Solaris server.

1.4 The Solaris LDoms Server Monitoring Model

This section will take a closer look at the LDoms monitoring model that eG Enterprise offers, the layers that constitute the model, the tests that map to each layer, and the metrics reported by the tests.

1.4.1 The Operating System Layer

The tests mapped to this layer report the following:

- the resource allocations to the primary domain
- how well the primary domain utilizes the allocated resources, and its impact on the physical resources
- how resource-efficient the Solaris host at the base is

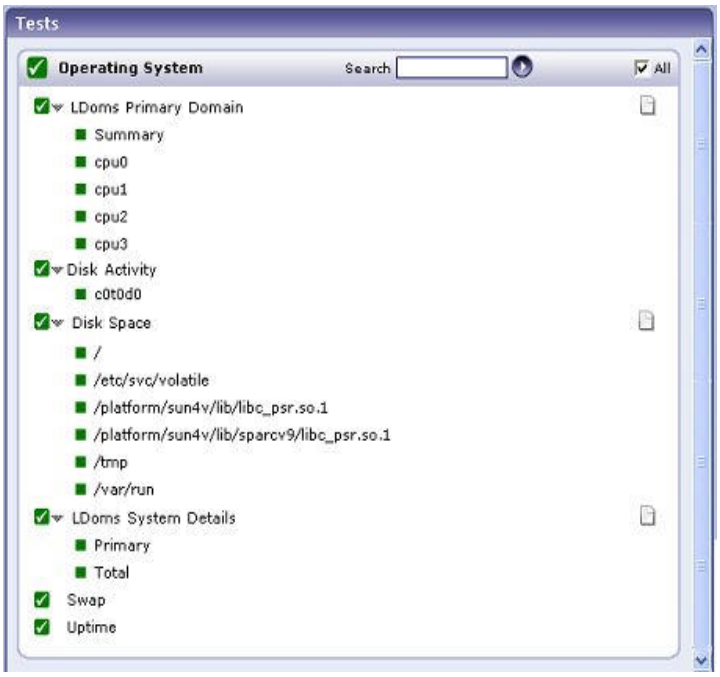

Figure 3: The tests mapped to the Operating System layer

The sections too come will discuss the LDoms Primary Domain test and the LDoms System Details tests alone, as all other tests have been dealt with elaborately in the *Monitoring Unix and Windows Servers* document.

1.4.1.1 LDoms Primary Domain Test

This LDoms Primary Domain test monitors the resource usage of the control domain (also referred to as the primary domain) of a Solaris LDoms server. The test reports the CPU and memory resources used by processes running in the primary domain. Using this test, administrators can determine whether any resource consuming processes are running in the primary domain.

Purpose	Monitors the resource usage of the primary domain of a Solaris LDoms server

Target of the test	A Solaris LDoms 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. **DURATION -** Determines the period (in seconds) for which the test runs each time. 4. **SUMMARY** – This attribute is applicable to multi-processor systems only. If the **Yes** option is selected, then the eG agent will report not only the CPU and memory utilization of each of the processors, but it will also report the summary (i.e., average) of the CPU and memory utilizations of the different processors. If the **No** option is selected, then the eG agent will report only the CPU usage of the individual processors. 5. **USEIOSTAT** – By default, the **USEIOSTAT** flag is set to **No**. This indicates that, by default, this test reports the CPU utilization of every processor on the system being monitored, and also provides the average CPU utilization across the processors. However, if you want the test to report only the average CPU utilization across processors and across user sessions, then set the **USEIOSTAT** flag to **Yes**. In such a case, the processor-wise breakup of CPU utilization will not be available. 6. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: ➢ The eG manager license should allow the detailed diagnosis capability. ➢ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for each processor supported by the primary domain on the Solaris LDoms host		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Virtual CPU utilization of console:** Indicates the percentage of the allocated CPU resources that were utilized by the control domain.	Percent	A very high value for this measure indicates excessive CPU utilization by that processor supported by the control domain. The CPU utilization may be high because a few processes are consuming a lot of CPU, or because there are too many processes contending for a limited resource. If adequate CPU resources are not available to the control domain, then the applications executing on the domain will experience slowdowns. To check whether resource-intensive processes are executing on the primary domain, use the detailed diagnosis capability of this measure; this will list the top 10 CPU consumers on the control domain.
	System usage of virtual CPU by console: Indicates the percentage of CPU time spent for system-level processing.	Percent	An unusually high value indicates a problem and may be due to too many system-level tasks executing simultaneously on the primary domain.
	Run queue length of console: Indicates the instantaneous length of the queue in which threads are waiting for the processor cycle. This length does not include the threads that are currently being executed.	Number	**In case of a control domain that supports multiple processors, this measure will be available only for the Summary descriptor of the test.** A value consistently greater than 2 indicates that many processes could be simultaneously contending for the processor.
	Blocked processes on console: Indicates the number of processes blocked for I/O, paging, etc.	Number	**In case of a control domain that supports multiple processors, this measure will be available only for the Summary descriptor of the test.** A high value could indicate an I/O problem on the control domain (e.g., a slow disk).

	Swap memory of console: This metric corresponds to the swap space currently available.	MB	**In case of a control domain that supports multiple processors, this measure will be available only for the Summary descriptor of the test.** An unusually high value for the swap usage can indicate a memory bottleneck. Check the memory utilization of individual processes to figure out the process(es) that has (have) maximum memory consumption and look to tune their memory usages and allocations accordingly.
	Free memory of console: Indicates the free memory available.	MB	**In case of a control domain that supports multiple processors, this measure will be available only for the Summary descriptor of the test.** A very low value of free memory is also an indication of high memory utilization on the control domain. The detailed diagnosis of this measure lists the top 10 processes responsible for maximum memory consumption on the control domain.
	Scan rate of console: Indicates the memory scan rate.	Pages/Sec	**In case of a control domain that supports multiple processors, this measure will be available only for the Summary descriptor of the test.** A high value is indicative of memory thrashing. Excessive thrashing can be detrimental to application performance.

The detailed diagnosis capability of the *Virtual CPU utilization of console* will list the CPU-intensive processes executing on the control domain. Resource intensive processes can be stopped to avoid excessive CPU usage. Alternatively, administrators can allocate additional CPU resources to the control domain to ensure that processes have adequate memory for execution.

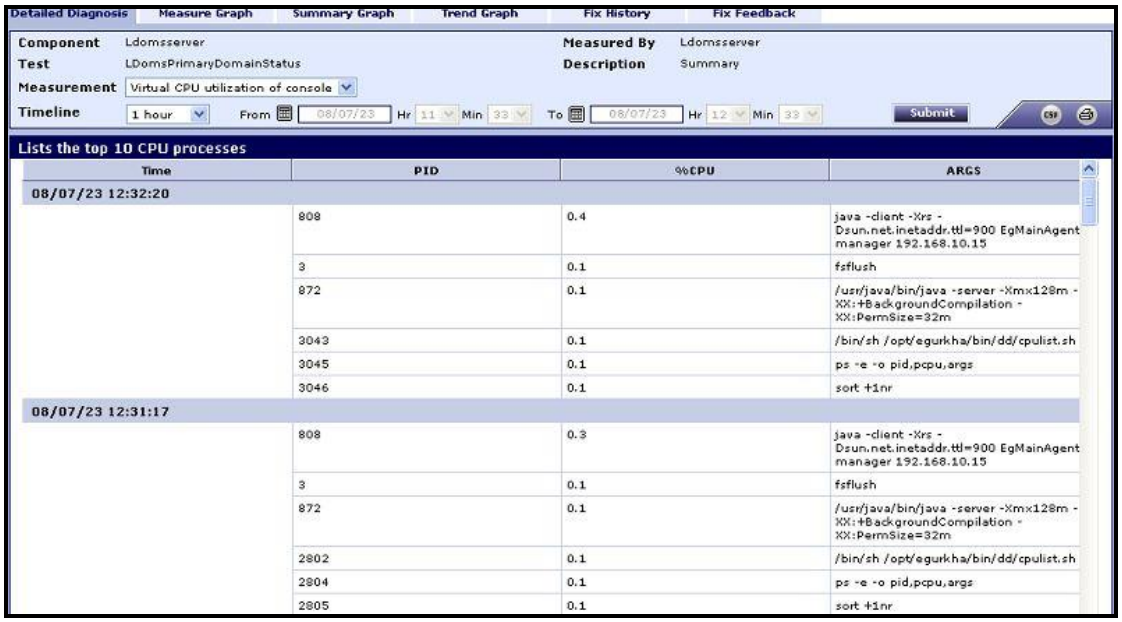

Figure 4: The detailed diagnosis of the Physical CPU usage measure

Note:

In case of multi-processor systems, this test typically reports the CPU usage for every processor. The *Summary* descriptor reports the average CPU usage across processors.

The detailed diagnosis of the *Free memory of console* measures indicates the top-1o memory-consuming processes executing on the control domain.

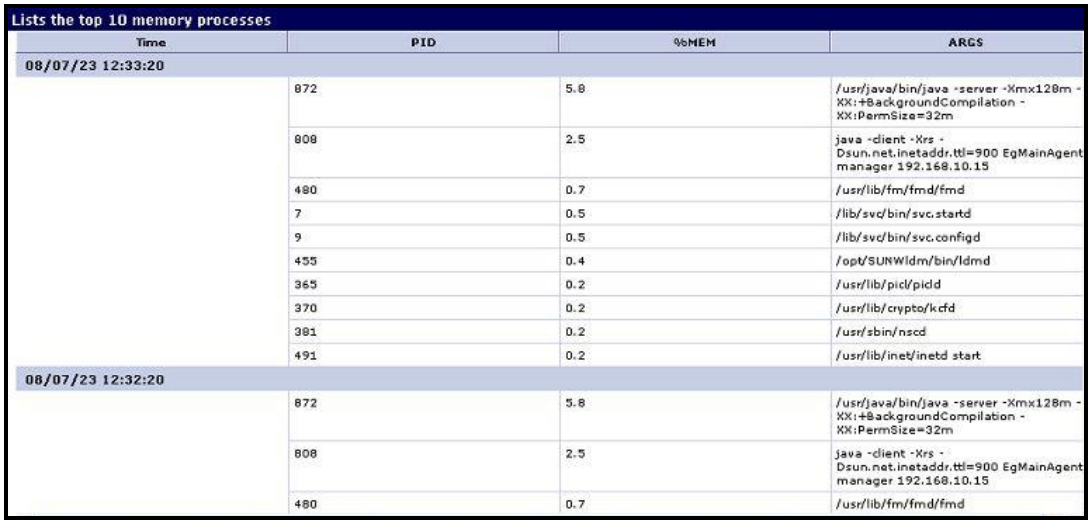

Figure 5: The detailed diagnosis of the *free memory of console* measure

1.4.1.2 LDoms System Details Test

This test gives an overview of CPU and memory resource usage on the Solaris physical server that is supporting one or more Logical Domains. This test reports the CPU and memory usage of the primary domain in relation to the physical resource availability at physical server (also referred to as the Solaris LDoms server). Using the results of this test, it is possible to determine the load on the Solaris physical server and also on the primary domain of the server.

Purpose	Reports the CPU and memory usage of the primary domain in relation to the physical resource availability at the host; also reveals how well the Solaris physical server as a whole utilizes the physical resources		
Target of the test	A Solaris LDoms server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TESTPERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL. 4. **KEONCOMMAND** - Keon is a security mechanism that can be used with a multitude of operating systems to provide a centralized base for user account and password management, user access and inactivity control, system integrity checking, and auditing. If the Keon security model is in use on the Solaris LDoms host being monitored, then this test may require special user privileges for executing the operating system commands. In such a case, specify the exact command that the test is permitted to execute, in the **KEONCOMMAND** text box. For example, if the keon command to be executed by the test is *sudo*, specify *sudo* in the **KEONCOMMAND** text box. Alternatively, you can even specify the full path to the *sudo* command in the **KEONCOMMAND** text box. On the other hand, if a Keon security model is not in place, then set the **KEONCOMMAND** parameter to *none*. 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. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: ➢ The eG manager license should allow the detailed diagnosis capability. ➢ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for the primary domain and one for the Solaris host as a whole		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Physical CPU utilization:** Indicates the percentage of physical CPU resources utilized. Two descriptors are reported for this test – for the "Primary" descriptor, this is the total CPU utilization of the primary (or control) domain, whereas the "Total" descriptor refers to the total CPU utilization on the Solaris server as a whole (including all the LDoms and the primary domain).	Percent	This measure clearly indicates where resource-intensive processes are executing – on the Solaris host or on the control domain? If the processes executing on the control domain or on the Solaris host erode the physical CPU resource-base, it is bound to affect the resource allocations to the other guest domains, and ultimately impact the performance of applications executing on those guest domains. Use the detailed diagnosis of this measure to determine the top CPU-consuming processes on the control domain or on the Solaris host, and kill rogue processes to ensure optimal CPU usage.
	Virtual CPU utilization: This metric is available only for the "Primary" descriptor and it indicates the percentage of allocated CPU resources that the primary domain has utilized.	Percent	A very high value for this measure indicates excessive CPU utilization by that processor supported by the control domain. The CPU utilization may be high because a few processes are consuming a lot of CPU, or because there are too many processes contending for a limited resource. If adequate CPU resources are not available to the control domain, then the applications executing on the domain will experience slowdowns. To check whether resource-intensive processes are executing on the primary domain, use the detailed diagnosis capability of this measure; this will list the top 10 CPU consumers on the control domain. **Note that this measure is available only for the Primary descriptor of this test.**

	Percent of virtual CPUs allocated to VMs: For the **Primary** descriptor, this indicates the ratio of the number of virtual CPUs allocated to the control domain to the total number of CPUs available on the Solaris physical server, expressed as a percentage. For the **Total** descriptor, this measure is the ratio of the number of virtual CPUs allocated to all the VMs (Logical domains) to the total number of CPUs available on the Solaris physical server, expressed as a percentage.	Percent	The values of these measures are a clear indicator to whether the hypervisor has over-allocated or under-allocated CPU resources to the domains.
	Virtual CPUs allocated to VMs: For the **Primary** descriptor, this indicates the number of virtual CPUs allocated to the primary domain. For the **Total** descriptor, this measure indicates the number of virtual CPUs allocated to all the guest domains (excluding the primary domain) on the Solaris physical server	Number	
	Physical memory allocated to VMs: For the **Primary** descriptor, this indicates the physical memory size allocated to the primary domain. For the **Total** descriptor, this measure indicates the total physical memory that is allocated to other guest domains on the Solaris host.	GB	The value of these measures indicates how efficiently the hypervisor has allocated memory resources to the control domain and other guest domains. These values also enable administrators to figure out if the Solaris LDoms server is adequately sized in terms of memory resources to handle the load imposed by the logical domains.

	Percent of physical memory allocated to VMs: For the **Primary** descriptor, this indicates the percentage of physical memory allocated to the primary domain. For the **Total** descriptor, this measure indicates the percentage of physical memory that is allocated to other guest domains on the Solaris host.	Percent	
	Total memory: Indicates the total memory on the Solaris host.	MB	

1.4.2 The Network Layer

The tests mapped to this layer indicate whether the Solaris LDoms server is available over the network, and also measures the network traffic to and from every network interface supported by the server.

Figure 6: The tests mapped to the Network layer

These tests have already been dealt with elaborately in the *Monitoring Unix and Windows Servers* document.

1.4.3 The LDoms Guests Layer

The tests mapped to this layer auto-discover the guest domains configured on the Solaris LDoms server and report their status. The layer also reports the resource usage of each discovered guest in relation to the physical resources available to the Solaris host. Using the information reported by this test, administrators can:

- Determine which of the guests 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 guests is receiving a very high rate of requests compared to the others, this guest may be a candidate for migration to another server, so as to minimize the impact it has on the other guests on the current server.

- Determine times when sudden or steady spikes in the physical resource utilization are caused by the guest machines

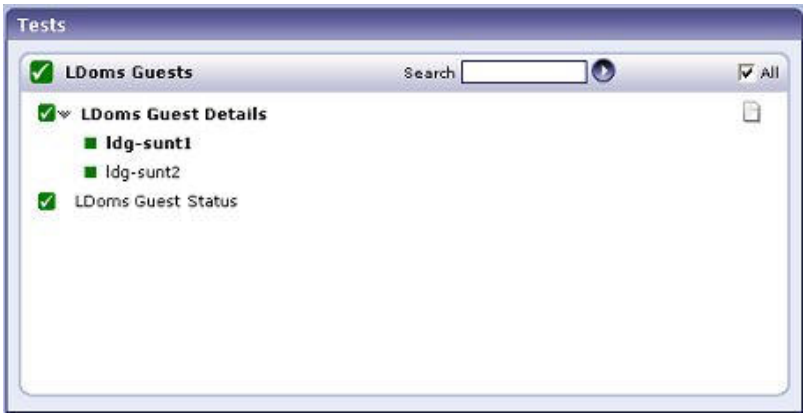

Figure 7: The tests mapped to the LDoms Guests layer

1.4.3.1 LDoms Guest Status Test

This test enables administrators to determine how many guests have registered with the Solaris LDoms server, and how many of these are currently running. In addition, the test also indicates whether any guests have been newly moved to the target Solaris LDoms server or removed from it. The numbers of guests in various states of activity are also revealed by this test.

Purpose	Reports the status of guests
Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL. 4. **KEONCOMMAND** - Keon is a security mechanism that can be used with a multitude of operating systems to provide a centralized base for user account and password management, user access and inactivity control, system integrity checking, and auditing. If the Keon security model is in use on the Solaris LDoms host being monitored, then this test may require special user privileges for executing the operating system commands. In such a case, specify the exact command that the test is permitted to execute, in the **KEONCOMMAND** text box. For example, if the keon command to be executed by the test is *sudo*, specify *sudo* in the **KEONCOMMAND** text box. Alternatively, you can even specify the full path to the *sudo* command in the **KEONCOMMAND** text box. On the other hand, if a Keon security model is not in place, then set the **KEONCOMMAND** parameter to *none*.

	5. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: ➢ The eG manager license should allow the detailed diagnosis capability. ➢ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for the Solaris LDoms server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Added guests: Indicates the number of guests that were newly added to the Solaris LDoms server during this measurement period.	Number	The detailed diagnosis of these measures, if enabled, lists the guest domains that were either newly created on the Solaris LDoms server or moved from another server.
	Removed guests: Indicates the number of guests that were newly removed from the Solaris LDoms server during this measurement period.	Number	
	Registered guests: Indicates the total number of guest domains that have been registered with the server.	Number	
	Active guests: Indicates the total number of guests that are currently running.	Number	
	Inactive_guests: Indicates the number of guests that are currently inactive on the server.	Number	When a guest domain is stopped and resources mapped to it are released (i.e., unbound), the guest domain automatically switches to the inactive state.

	Guests in bind state: Indicates the number of guests in the bind state.	Number	The LDom manager is responsible for mapping physical resources to logical domains. This resource mapping is referred to as 'binding'. When physical resources are assigned (i.e., bound) to a guest domain, the domain is said to be in a "bind" state.
	Guests in unbind state: Indicates the number of guests in the unbind state.	Number	If physical resources mapped (i.e., bound) to a logical domain are released, the domain switches to the *unbind* state.

The detailed diagnosis of the *Active guests* and *Inactive guests* measures, if enabled, lists the virtual machines that are in an active or inactive state (as the case may be) on the Solaris LDoms server (see Figure 8).

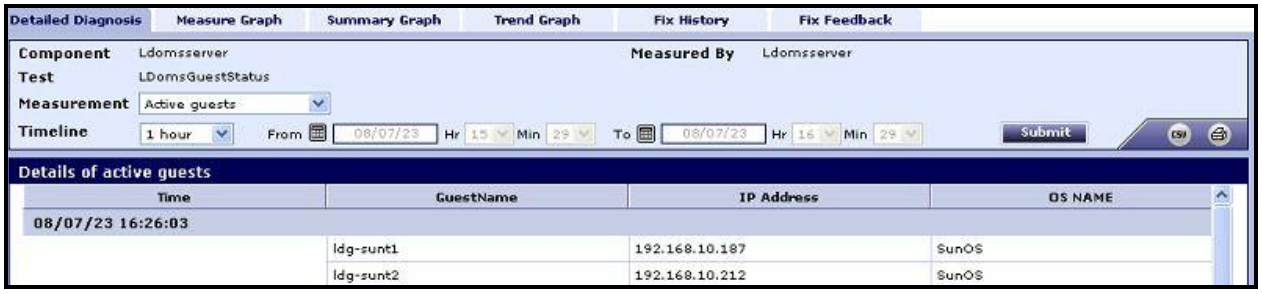

Figure 8: The detailed diagnosis of the Active guests measure

1.4.3.2 LDoms Guest Details Test

This test monitors the amount of the physical server's resources that each guest on a Solaris LDoms server is taking up. Using the metrics reported by this test, administrators can determine which virtual guest is taking up most CPU, which guest is taking up the maximum memory, etc. Note that the amount of resources taken up by a guest will be limited by the resource allocations that have been made by administrators. For example, an administrator could cap the amount of memory that a specific guest may take.

Purpose	Monitors the amount of the physical server's resources that each guest on an Solaris LDoms server is taking up
Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

<table>
<tr><td></td><td colspan="3">4. DOMAIN - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to none.

5. ADMIN USER - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against ADMIN USER, and specify his password against ADMIN PASSWORD. On the other hand, if the user credentials vary from one guest to another, then multiple ADMIN USERs and ADMIN PASSWORDs might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document.

6. ADMIN PASSWORD - The password of the ADMIN USER needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document.

7. CONFIRM PASSWORD - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.

8. KEONCOMMAND - Keon is a security mechanism that can be used with a multitude of operating systems to provide a centralized base for user account and password management, user access and inactivity control, system integrity checking, and auditing. If the Keon security model is in use on the Solaris LDoms host being monitored, then this test may require special user privileges for executing the operating system commands. In such a case, specify the exact command that the test is permitted to execute, in the KEONCOMMAND text box. For example, if the keon command to be executed by the test is sudo, specify sudo in the KEONCOMMAND text box. Alternatively, you can even specify the full path to the sudo command in the KEONCOMMAND text box. On the other hand, if a Keon security model is not in place, then set the KEONCOMMAND parameter to none.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each guest configured on the Solaris LDoms server</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Is VM powered on?:
Whether the logical domain is currently running on the Solaris LDoms server or no.</td><td>Boolean</td><td>The value 1 indicates that the guest is up and running. The value 0 could indicate that the guest has been powered-off; it could also indicate that the guest has been moved to a different server.

While the test reports a wide variety of other metrics too for virtual machines that are alive, only the powered on status is indicated for virtual machines that are currently not available.</td></tr>
</table>

	Physical CPU utilization of VM: Indicates the percentage of the physical CPU processing power of the physical server that is used by the guest.	Percent	A high value for this measure indicates a virtual machine that one/more processes executing on the guest domain are using up too much physical CPU. Beside draining the server of its physical CPU resources, this can also affect the CPU allocations to other guests, thereby adversely impacting the performance of the applications on those guest domains.
	Virtual CPU utilization of VM: Indicates the percentage of virtual CPU utilized by this guest domain.	Percent	A high value for this measure indicates that resource-intensive processes are executing on the guest domain.
	Percent of virtual CPUs allocated: Indicates the ratio of the virtual CPUs allocated to this guest domain to the total number of CPUs on the server, expressed as a percentage.	Percent	These measures help determine whether CPU allocations to the guest domain are commensurate to the current and expected processing load on the domain; over-allocations/under-allocations are thus brought to light.
	Virtual CPUs allocated to VM: Indicates the number of CPUs allocated to the guest domain.	Number	
	Physical memory allocated to VM: Indicates the amount of physical memory currently allocated to the guest.	MB	These measures help determine whether memory allocations to the guest domain are commensurate to the memory requirements of the domain; over-allocations/under-allocations are thus brought to light.
	Percent of physical memory allocated to VM: Indicates the percentage of total physical memory that is allocated to VM.	Percent	

1.4.4 The Inside View of VMs Layer

The **LDoms Guests** layer provides an "external" view of the different guest domains – the metrics reported at this layer are based on what the Solaris host is seeing about the performance of the individual guest domains. However, an external view of the guest operating system and its applications may not be sufficient. For instance, suppose one of the disk partitions of the guest operating system has reached capacity. This information cannot be gleaned from host operating system. Likewise, bottlenecks such as a longer process run queue or a higher disk queue length are

more visible using an internal monitor. Internal monitoring (from within the guest operating system) also provides details about the resource utilization of different application(s) or processes.

The tests mapped to the **Inside View of VMs** layer provide an "internal" view of the workings of each of the guests - these tests execute on an Solaris LDoms server host, but send probes into each of the guest operating systems to analyze how well each guest utilizes the resources that are allocated to it, and how well it handles load.

By default however, clicking on the **Inside View of VMs** layer, does not display the list of tests associated with that layer. Instead, Figure 9 appears which provides you with an overview of individual guest performance (see Figure 9).

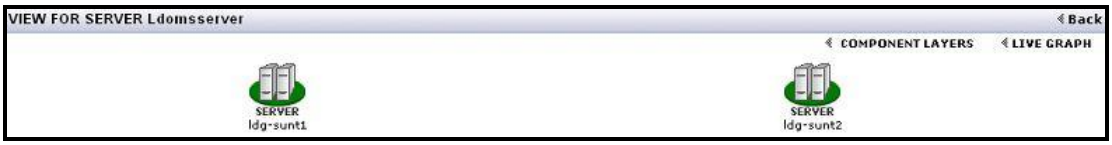

Figure 9: A list of guest operating systems on a Solaris LDoms server and their current state

To return to the layer model of the *Solaris LDoms* server and view the tests associated with the **Inside View of VMs** layer, as depicted by Figure 10 below.

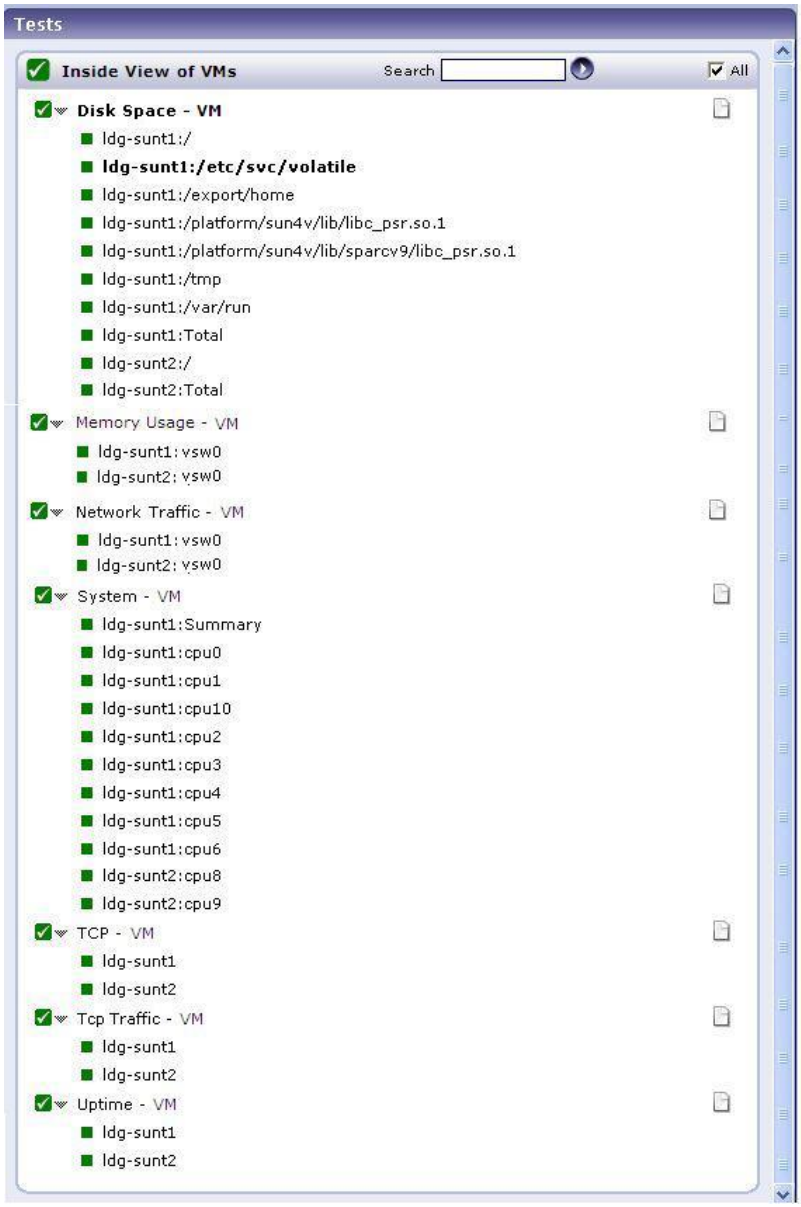

Figure 10: The tests associated with the Inside View of VMs layer

As indicated in Figure 10, the tests associated with this layer monitor different aspects of each virtual guest. Disk space utilization, CPU utilization, memory usage levels, TCP traffic, etc. are all monitored and reported for each virtual guest hosted on the Solaris LDoms server. Detailed diagnosis for these tests provide details of individual processes and their utilization levels.

The tests associated with this layer are described in detail below.

1.4.4.1 Disk Space - VM Test

This test reports statistics pertaining to the space usage of each physical disk on a guest, and across all disk partitions on a guest.

Purpose	Reports statistics pertaining to the space usage of each physical disk on a guest, and across all disk partitions on a guest

<table>
<tr><td>Target of the test</td><td colspan="3">A Solaris LDoms server</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT - The port at which the specified HOST listens. By default, this is NULL.
4. DOMAIN - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to none.
5. ADMIN USER - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against ADMIN USER, and specify his password against ADMIN PASSWORD. On the other hand, if the user credentials vary from one guest to another, then multiple ADMIN USERs and ADMIN PASSWORDs might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document.
6. ADMIN PASSWORD - The password of the ADMIN USER needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document.
7. CONFIRM PASSWORD - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every combination of virtual_guest:disk_partition; for every guest, a set of measures will be reported for the virtual_guest:Total descriptor as well – these measures will indicate the total disk space usage across all disk partitions on that guest.</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total capacity:
Indicates the total capacity of a disk partition.</td><td>MB</td><td></td></tr>
<tr><td>Used space:
Indicates the amount of space used in a disk partition.</td><td>MB</td><td></td></tr>
</table>

	Free space: Indicates the current free space available for each disk partition of a system.	MB	
	Percent usage: Indicates the percentage of space usage on each disk partition of a system.	Percent	A value close to 100% can indicate a potential problem situation where applications executing on the guest may not be able to write data to the disk partition(s) with very high usage.

1.4.4.1.1 Configuring Users for VM Monitoring

In order to enable the eG agent to connect to VMs and pull out metrics from them, the eG administrative interface provides a special page using which the different **ADMIN USER** names and **ADMIN PASSWORDS** can be specified. To access this page, just click on the **Click here** hyperlink in any of the VM test configuration pages.

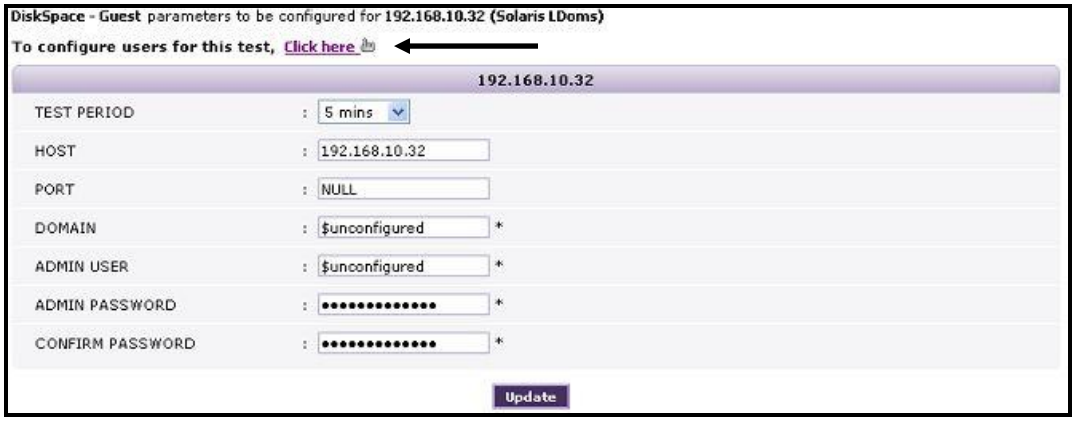

Figure 11: Configuring a VM test

Upon clicking, Figure 12 will appear, using which the VM user details can be configured.

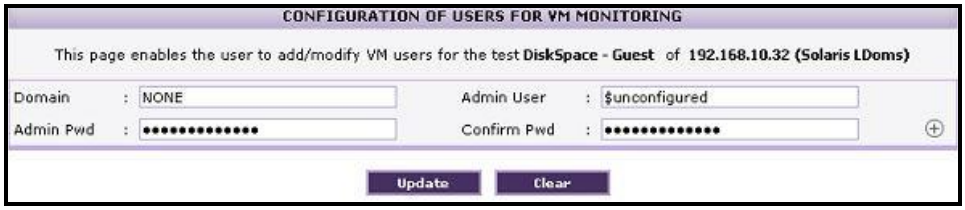

Figure 12: The VM user configuration page

To add a user specification, do the following:

1. By default, *NONE* will be displayed against **Domain**.
2. The eG agent must be configured with user privileges that will allow the agent to communicate with the VMs on a Solaris server and extract statistics. Specify of the name of such a user the **Admin User** text box.
3. The password of the specified **Admin User** should be mentioned in the **Admin Pwd** text box.
4. Confirm the password by retyping it in the **Confirm Pwd** text box.
5. To add more users, click on the ⊕ button in Figure 12. This will allow you to add one more user specification as

depicted by Figure 13.

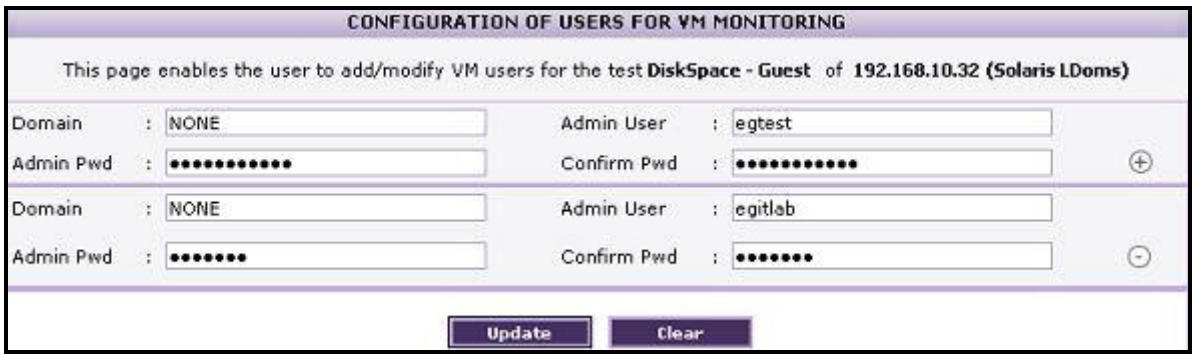

Figure 13: Adding another user

6. Similarly, multiple user specifications can be provided.

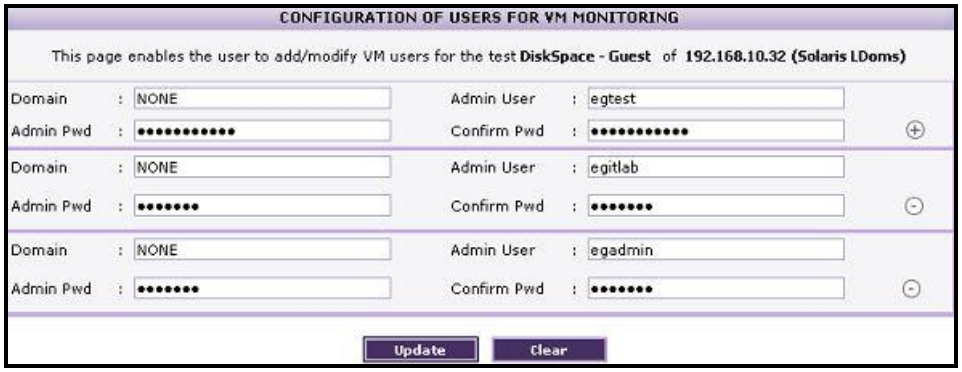

Figure 14: Adding multiple user specifications

7. To clear all the user specifications, simply click the **Clear** button in Figure 14.
8. To remove the details of a particular user alone, just click the ⊖ button in Figure 14.
9. To save the specification, just click on the **Update** button in Figure 14. This will lead you back to the test configuration page, where you will find the multiple domain names, user names, and passwords listed against the respective fields.

1.4.4.2 Memory Usage - VM Test

This test reports statistics related to the usage of physical memory of the LDOMs.

Purpose	Reports statistics related to the usage of physical memory of the LDOMs
Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT - The port at which the specified HOST listens. By default, this is NULL.
4. DOMAIN - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to none.
5. ADMIN USER - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against ADMIN USER, and specify his password against ADMIN PASSWORD. On the other hand, if the user credentials vary from one guest to another, then multiple ADMIN USERs and ADMIN PASSWORDs might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document.
6. ADMIN PASSWORD - The password of the ADMIN USER needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document.
7. CONFIRM PASSWORD - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.
8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
➢ The eG manager license should allow the detailed diagnosis capability.
➢ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every LDOM on the Solaris server being monitored.</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total physical memory:
Indicates the total physical memory of this VM.</td><td>MB</td><td></td></tr>
<tr><td>Used physical memory:
Indicates the used physical memory of this VM.</td><td>MB</td><td></td></tr>
</table>

	Free physical memory: Indicates the free physical memory of the VM.	MB	This measure typically indicates the amount of memory available for use by applications running on the target VM. On Unix operating systems (AIX and Linux), the operating system tends to use parts of the available memory for caching files, objects, etc. When applications require additional memory, this is released from the operating system cache. Hence, to understand the true free memory that is available to applications, the eG agent reports the sum of the free physical memory and the operating system cache memory size as the value of the *Free physical memory* measure while monitoring AIX and Linux guest operating systems.
	Physical memory utilized: Indicates the percent usage of physical memory by this VM.	Percent	Ideally, the value of this measure should be low. While sporadic spikes in memory usage could be caused by one/more rogue processes on the VM, a consistent increase in this value could be a cause for some serious concern, as it indicates a gradual, but steady erosion of valuable memory resources. If this unhealthy trend is not repaired soon, it could severely hamper VM performance, causing anything from a slowdown to a complete system meltdown. You can use the detailed diagnosis of this measure to figure out which processes on the VM are consuming memory excessively.

	Available physical memory: Indicates the amount of physical memory, immediately available for allocation to a process or for system use.	MB	Not all of the *Available physical memory* is *Free physical memory*. Typically, *Available physical memory* is made up of the Standby List, Free List, and Zeroed List. When Windows wants to trim a process' working set, the trimmed pages are moved (usually) to the Standby List. From here, they can be brought back to life in the working set with only a soft page fault (much faster than a hard fault, which would have to talk to the disk). If a page stays in the standby List for a long time, it gets freed and moved to the Free List. In the background, there is a low priority thread (actually, the only thread with priority 0) which takes pages from the Free List and zeros them out. Because of this, there is usually very little in the Free List. All new allocations always come from the Zeroed List, which is memory pages that have been overwritten with zeros. This is a standard part of the OS' cross-process security, to prevent any process ever seeing data from another. If the Zeroed List is empty, Free List memory is zeroed and used or, if that is empty too, Standby List memory is freed, zeroed, and used. It is because all three can be used with so little effort that they are all counted as "available". A high value is typically desired for this measure. **This measure will be available for Windows 2008 VMs only.**

	Modified memory: Indicates the amount of memory that is allocated to the modified page list.	MB	This memory contains cached data and code that is not actively in use by processes, the system and the system cache. This memory needs to be written out before it will be available for allocation to a process or for system use. Cache pages on the modified list have been altered in memory. No process has specifically asked for this data to be in memory, it is merely there as a consequence of caching. Therefore it can be written to disk at any time (not to the page file, but to its original file location) and reused. However, since this involves I/O, it is not considered to be Available physical memory. **This measure will be available for Windows 2008 VMs only.**
	Standby memory: Indicates the amount of memory assigned to the standby list.	MB	This memory contains cached data and code that is not actively in use by processes, the system and the system cache. It is immediately available for allocation to a process or for system use. If the system runs out of available free and zero memory, memory on lower priority standby cache page lists will be repurposed before memory on higher priority standby cache page lists. Typically, Standby memory is the aggregate of Standby Cache Core Bytes,Standby Cache Normal Priority Bytes, and Standby Cache Reserve Bytes. Standby Cache Core Bytes is the amount of physical memory, that is assigned to the core standby cache page lists. Standby Cache Normal Priority Bytes is the amount of physical memory, that is assigned to the normal priority standby cache page lists. Standby Cache Reserve Bytes is the amount of physical memory, that is assigned to the reserve standby cache page lists. **This measure will be available for Windows 2008 VMs only.**
	Cached memory: This measure is an aggregate of *Standby memory* and *Modified memory*.	MB	**This measure will be available for Windows 2008 VMs only.**

Note:

While monitoring Linux/AIX guest operating systems, you may observe discrepancies between the value of the *Physical memory utilized* measure and the memory usage percentages reported per process by the detailed diagnosis of the same measure. This is because, while the *Physical memory utilized* measure takes into account the memory in the OS cache of the Linux/AIX VM, the memory usage percent that the detailed diagnosis reports per process does not consider the OS cache memory.

1.4.4.3 System Details - VM Test

This test collects various metrics pertaining to the CPU and memory usage of every processor supported by a guest. The details of this test are as follows:

Purpose	To measure the CPU and memory usage of each guest of a Solaris LDoms server
Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT - The port at which the specified HOST listens. By default, this is NULL.
4. DOMAIN - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to none.
5. ADMIN USER - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against ADMIN USER, and specify his password against ADMIN PASSWORD. On the other hand, if the user credentials vary from one guest to another, then multiple ADMIN USERs and ADMIN PASSWORDs might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document.
6. ADMIN PASSWORD - The password of the ADMIN USER needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document.
7. CONFIRM PASSWORD - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.
8. ENABLE MEMORY DIAGNOSIS - By default, the ENABLE MEMORY DIAGNOSIS flag is set to NO, indicating that detailed diagnosis will not be available for the Free memory in VM measure reported by this test by default. If you want to view the detailed diagnosis of the Free memory in VM measure - i.e., to view the top 10 processes on the target VM that are utilizing memory excessively - you can change this flag to YES.
9. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
➢ The eG manager license should allow the detailed diagnosis capability.
➢ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every combination of virtual_guest:processor; for every guest, a set of measures will be reported for the virtual_guest:Summary descriptor as well – these measures will indicate the total CPU usage across all processors supported by that guest.</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

test	**Virtual CPU utilization:** This measurement indicates the percentage of CPU utilized by the processor.	Percent	A high value could signify a CPU bottleneck. The CPU utilization may be high because a few processes are consuming a lot of CPU, or because there are too many processes contending for a limited resource. The detailed diagnosis of this test reveals the top-10 CPU-intensive processes on the guest.
	System usage of virtual CPU: Indicates the percentage of CPU time spent for system-level processing.	Percent	An unusually high value indicates a problem and may be due to too many system-level tasks executing simultaneously.
	Run queue in VM: Indicates the instantaneous length of the queue in which threads are waiting for the processor cycle. This length does not include the threads that are currently being executed.	Number	A value consistently greater than 2 indicates that many processes could be simultaneously contending for the processor. **In case of guest domains that support multiple processors, this measure will be available only for the 'Summary' descriptor.**
	Blocked processes in VM: Indicates the number of processes blocked for I/O, paging, etc.	Number	A high value could indicate an I/O problem on the guest (e.g., a slow disk). **In case of guest domains that support multiple processors, this measure will be available only for the 'Summary' descriptor.**
	Swap memory in VM: Denotes the committed amount of virtual memory. This corresponds to the space reserved for virtual memory on disk paging file(s).	MB	An unusually high value for the swap usage can indicate a memory bottleneck. Check the memory utilization of individual processes to figure out the process(es) that has (have) maximum memory consumption and look to tune their memory usages and allocations accordingly. **In case of guest domains that support multiple processors, this measure will be available only for the 'Summary' descriptor.**

	Free memory in VM: Indicates the free memory available.	MB	This measure typically indicates the amount of memory available for use by applications running on the target VM. On Unix operating systems (AIX and Linux), the operating system tends to use parts of the available memory for caching files, objects, etc. When applications require additional memory, this is released from the operating system cache. Hence, to understand the true free memory that is available to applications, the eG agent reports the sum of the free physical memory and the operating system cache memory size as the value of the *Free memory in VM* measure while monitoring AIX and Linux guest operating systems. **In case of guest domains that support multiple processors, this measure will be available only for the 'Summary' descriptor.** The detailed diagnosis of this measure, if enabled, lists the top 10 processes responsible for maximum memory consumption on the target VM.
	Scan rate in VM: Indicates the memory scan rate.	Pages/Sec	A high value is indicative of memory thrashing. Excessive thrashing can be detrimental to guest performance. **In case of guest domains that support multiple processors, this measure will be available only for the 'Summary' descriptor.**

The detailed diagnosis capability of the *Virtual CPU utilization* measure, if enabled, provides a listing of the top 10 CPU-consuming processes (see Figure 15). In the event of a Cpu bottleneck, this information will enable users to identify the processes consuming a high percentage of CPU time. The users may then decide to stop such processes, so as to release the CPU resource for more important processing purposes.

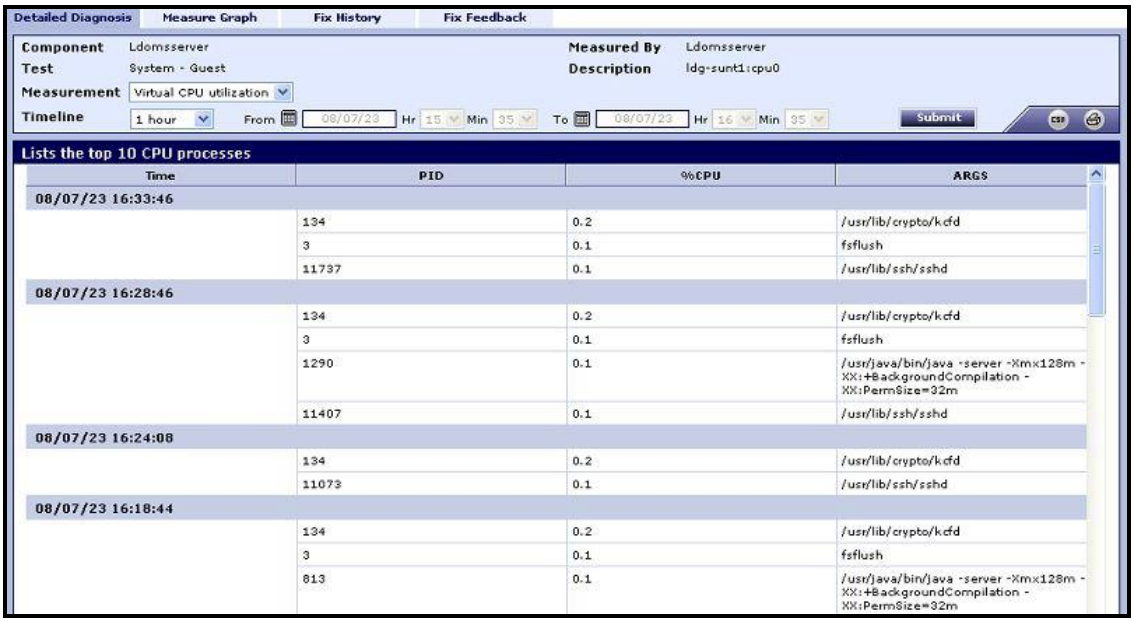

Figure 15: The top 10 CPU consuming processes

Note:

While instantaneous spikes in CPU utilization are captured by the eG agents and displayed in the Measures page, the detailed diagnosis will not capture/display such instantaneous spikes. Instead, detailed diagnosis will display only a consistent increase in CPU utilization observed over a period of time.

1.4.4.4 Network Traffic - VM Test

This is an internal test that monitors the incoming and outgoing traffic through each guest on a Solaris LDoms server.

Purpose	To measure the incoming and outgoing traffic through each guest on a Solaris LDoms server
Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TESTPERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL. 4. **DOMAIN** - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to *none*. 5. **ADMIN USER** - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against **ADMIN USER**, and specify his password against **ADMIN PASSWORD**. On the other hand, if the user credentials vary from one guest to another, then multiple **ADMIN USER**s and **ADMIN PASSWORD**s might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document. 6. **ADMIN PASSWORD** - The password of the **ADMIN USER** needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document. 7. **CONFIRM PASSWORD** - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.		
Outputs of the test	One set of results for every *virtual_guest:network_interface* combination		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Incoming network traffic: Indicates the rate of incoming traffic.	Pkts/Sec	An increase in traffic to the guest domain can indicate an increase in accesses to the domain (from users or from other applications) or that the domain is under an attack of some form.
	Outgoing network traffic: Represents the rate of outgoing traffic.	Pkts/Sec	An increase in traffic from the guest domain can indicate an increase in accesses to the domain (from users or from other applications).

1.4.4.5 TCP - VM Test

This test tracks various statistics pertaining to TCP connections to and from each guest of a Solaris LDoms server. The details of the test are provided below:

Purpose	To measure statistics pertaining to the TCP layer of a guest domain

Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TESTPERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL. 4. **DOMAIN** - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to *none*. 5. **ADMIN USER** - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against **ADMIN USER**, and specify his password against **ADMIN PASSWORD**. On the other hand, if the user credentials vary from one guest to another, then multiple **ADMIN USER**s and **ADMIN PASSWORD**s might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document. 6. **ADMIN PASSWORD** - The password of the **ADMIN USER** needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document. 7. **CONFIRM PASSWORD** - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.

Outputs of the test	One set of results for each powered-on guest on the Solaris LDoms server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Incoming connections to VM: Indicates the connections per second received by the guest.	Conns/Sec	A high value can indicate an increase in input load.
	Outgoing connections to VM: Indicates the connections per second initiated by the guest.	Conns/Sec	A high value can indicate that one or more of the applications executing on the guest have started using a number of TCP connections to some other guest.
	Current connections to VM: Indicates the currently established connections.	Number	A sudden increase in the number of connections established on a guest can indicate either an increase in load to one or more of the applications executing on the guest, or that one or more of the applications are experiencing a problem (e.g., a slow down).
	Connection drops on VM: Indicates the rate of established TCP connections dropped from the TCP listen queue.	Conns/Sec	This value should be 0 for most of the time. Any non-zero value implies that one or more applications on the guest are under overload.
	Connection failures on VM: Indicates the rate of half open TCP connections dropped from the listen queue.	Conns/Sec	This value should be 0 for most of the time. A prolonged non-zero value can indicate either that the guest is under SYN attack or that there is a problem with the network link to the guest that is resulting in connections being dropped without completion.

1.4.4.6 TCP Traffic - VM Test

Since most popular applications rely on the TCP protocol for their proper functioning, traffic monitoring at the TCP protocol layer can provide good indicators of the performance seen by the applications that use TCP. The most critical metric at the TCP protocol layer is the percentage of retransmissions. Since TCP uses an exponential back-off algorithm for its retransmissions, any retransmission of packets over the network (due to network congestion, noise, data link errors, etc.) can have a significant impact on the throughput seen by applications that use TCP. This test monitors the TCP protocol traffic to and from a guest, and particularly monitors retransmissions.

Purpose	Monitors the TCP protocol traffic to and from each guest domain of a Solaris LDoms server, and particularly measures the percentage of retransmission
Target of the test	A Solaris LDoms server
Agent deploying the	An internal agent

<table>
<tr><td>test</td><td colspan="3"></td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TESTPERIOD – How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT - The port at which the specified HOST listens. By default, this is NULL.
4. DOMAIN - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to none.
5. ADMIN USER - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against ADMIN USER, and specify his password against ADMIN PASSWORD. On the other hand, if the user credentials vary from one guest to another, then multiple ADMIN USERs and ADMIN PASSWORDs might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document.
6. ADMIN PASSWORD - The password of the ADMIN USER needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document.
7. CONFIRM PASSWORD - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each powered-on guest on the Solaris LDoms server being monitored</td></tr>
<tr><td rowspan="4">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Segments received by VM:
Indicates the rate at which segments are received by the guest.</td><td>Segments/Sec</td><td></td></tr>
<tr><td>Segments sent by VM:
Indicates the rate at which segments are sent to clients or other guests.</td><td>Segments/Sec</td><td></td></tr>
<tr><td>Retransmits by VM:
Indicates the rate at which segments are being retransmitted by the guest</td><td>Segments/Sec</td><td></td></tr>
</table>

	Retransmit ratio from VM: Indicates the ratio of the rate of data retransmissions to the rate of data being sent by the guest	Percent	Ideally, the retransmission ratio should be low (< 5%). Most often retransmissions at the TCP layer have significant impact on application performance. Very often a large number of retransmissions are caused by a congested network link, bottlenecks at a router causing buffer/queue overflows, or by lousy network links due to poor physical layer characteristics (e.g., low signal to noise ratio). By tracking the percentage of retransmissions at a guest, an administrator can quickly be alerted to problem situations in the network link(s) to the guest that may be impacting the service performance.

1.4.4.7 Uptime - VM Test

In most virtualized environments, it is essential to monitor the uptime of logical domains hosting critical server applications in the infrastructure. By tracking the uptime of each of the guest domains, administrators can determine what percentage of time a domain has been up. Comparing this value with service level targets, administrators can determine the most trouble-prone areas of the virtualized infrastructure.

In some environments, administrators may schedule periodic reboots of their guest domain. By knowing that a specific domain has been up for an unusually long time, an administrator may come to know that the scheduled reboot task is not working on a domain.

The Uptime - Guest test included in the eG agent monitors the uptime of critical logical domains on a Solaris LDoms server.

Purpose	To monitor the uptime of each guest on a Solaris LDoms server
Target of the test	A Solaris LDoms server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TESTPERIOD** – How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL. 4. **DOMAIN** - Specify the domain within which the virtual guests reside. Since the Solaris LDoms server supports only Solaris and Linux guests, this parameter should always be set to *none*. 5. **ADMIN USER** - This test connects to each virtual guest and collects status and resource usage statistics from the guest. In order to do so, the test must be configured with user privileges that allow a remote connection to the virtual guest from the Solaris host. If a single user has access to all the guest domains on the Solaris server, specify the name of that user against **ADMIN USER**, and specify his password against **ADMIN PASSWORD**. On the other hand, if the user credentials vary from one guest to another, then multiple **ADMIN USER**s and **ADMIN PASSWORD**s might have to be specified for every Solaris LDoms server being monitored. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 1.4.4.1.1 of this document. 6. **ADMIN PASSWORD** - The password of the **ADMIN USER** needs to be provided here. Here again, if multiple passwords need to be specified, then follow the procedure detailed in Section 1.4.4.1.1 of this document. 7. **CONFIRM PASSWORD** - Confirm the password by retyping it here. Here again, if multiple passwords need to be confirmed, then follow the procedure detailed in Section 1.4.4.1.1 of this document. 8. **REPORTMANAGERTIME** - By default, this flag is set to **Yes**, indicating that, by default, the detailed diagnosis of this test, if enabled, will report the shutdown and reboot times of the LDOMs in the manager's time zone. If this flag is set to **No**, then the shutdown and reboot times are shown in the time zone of the system where the agent is running. 9. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.		
Outputs of the test	One set of results for every guest on a Solaris LDoms server		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Has the VM been rebooted? : Indicates whether the logical domain has been rebooted during the last measurement period or not.	Boolean	If this measure shows 1, it means that the guest domain 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 domain was rebooted.
	Uptime of VM during the last measure period: Indicates the time period that the logical domain has been up since the last time this test ran.	Secs	If the guest 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 guest 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 guest 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 VM: Indicates the total time that the Ldom has been up since its last reboot.	Mins	Administrators may wish to be alerted if a guest 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 - Guest test, then, the **Uptime during the last measure period** measure will report the value 0 for Unix VMs (only) until the minute boundary is crossed. For instance, if you configure the Uptime - Guest test to run every 10 seconds, then, for the first 5 test execution cyles (i.e., 10 x 5 = 50 seconds), the **Uptime during the last measure period** measure will report the value 0 for Unix VMs; 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 same VMs. Thereafter, every sixth measurement period will report 60 seconds as the uptime of the Unix VMs. This is because Unix-based operating systems report uptime only in minutes and not in seconds.

As stated earlier, by default, clicking on the **Virtual Servers** layer of a managed *Solaris LDoms* server, leads you to a page displaying the current status of the virtual guests executing on that server. If you want to override this default setting - i.e., if you prefer to view the tests mapped to the **Virtual Servers** layer first, and then proceed to focus on individual guest performance, follow the steps given below:

- Edit the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory
- Set the **LAYERMODEL_LINK_TO_VIRTUAL** flag in the file to **false**; this is set to **true** by default.
- Save the **eg_ui.ini** file.

Doing so ensures that as soon as the **Virtual Servers** layer is clicked, the list of tests mapped to that layer appears, as depicted by Figure 16.

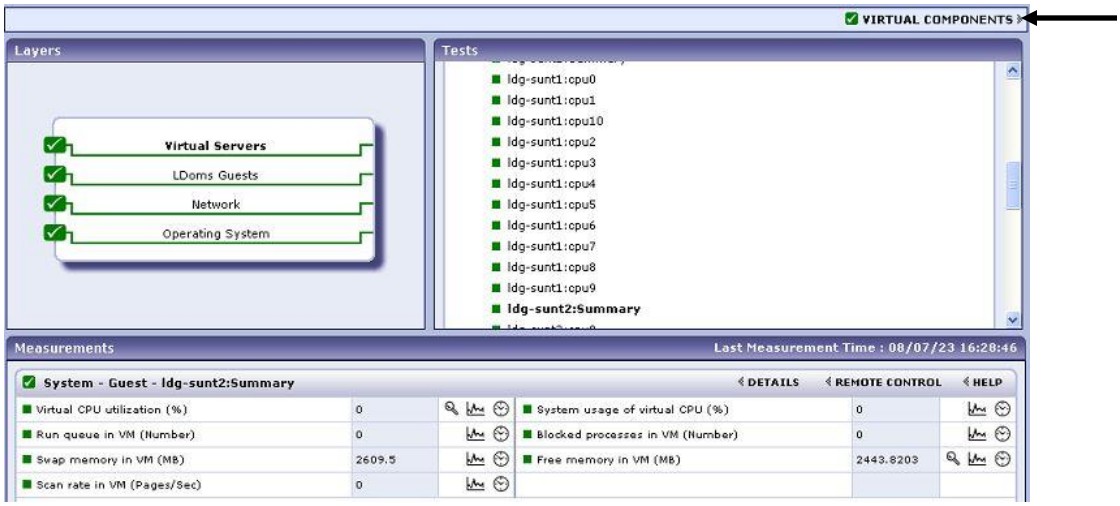

Figure 16: The tests mapped to the Virtual Servers layer

If you now want the **Server view** of Figure 9, simply click on the **Servers** link above the list of tests in Figure 16 (indicated by the arrow).

Clicking on any of the guests in the **Server view** leads you to Figure 17 that displays all the performance metrics extracted from that guest, in real-time. 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, you can view its graph by clicking on that measure in Figure 17.

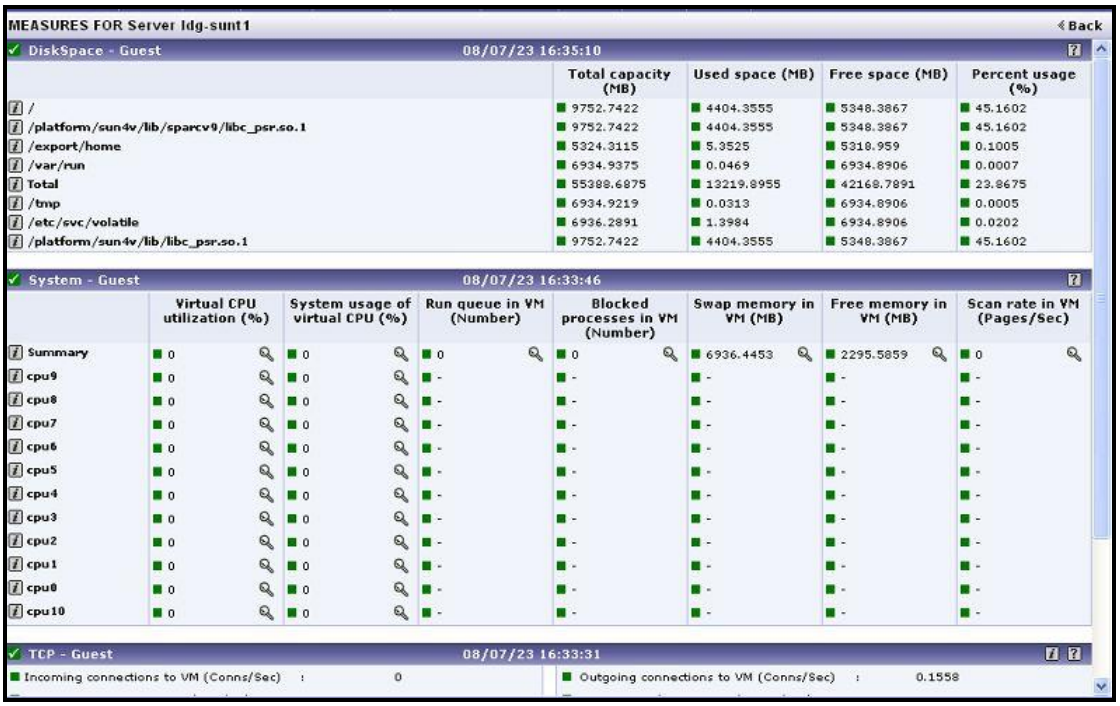

Figure 17: Measures pertaining to a chosen guest

To view real-time graphs of pre-configured measures (pertaining to the Solaris LDoms server and the guests operating on it), click on the **LIVE GRAPH** link in Figure 9. The graph display that appears subsequently (see Figure 18) has been organized in such a way that next to every host-pertinent measure graph, the closely related guest-specific measure graph appears. For instance, next to the graph of the 'Physical CPU utilization' measure of the *LdomsSystemDetails* test, you will find a graph of the 'Physical CPU utilization of VM' measure of the *LDomsGuestDetails* test. This way, you can easily compare and correlate how well the physical CPU resources are being utilized by both the Solaris host and the LDoms configured on it. 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 Solaris host? or is it the logical domain? If you access this page from the **LIVE GRAPH** link in Figure 9, then, by default, you will view live graphs pertaining to the *Solaris LDoms* server. However, you can select a different virtualized component-type and a different virtualized component using the **type** and **Component Name** lists (respectively) in Figure 18.

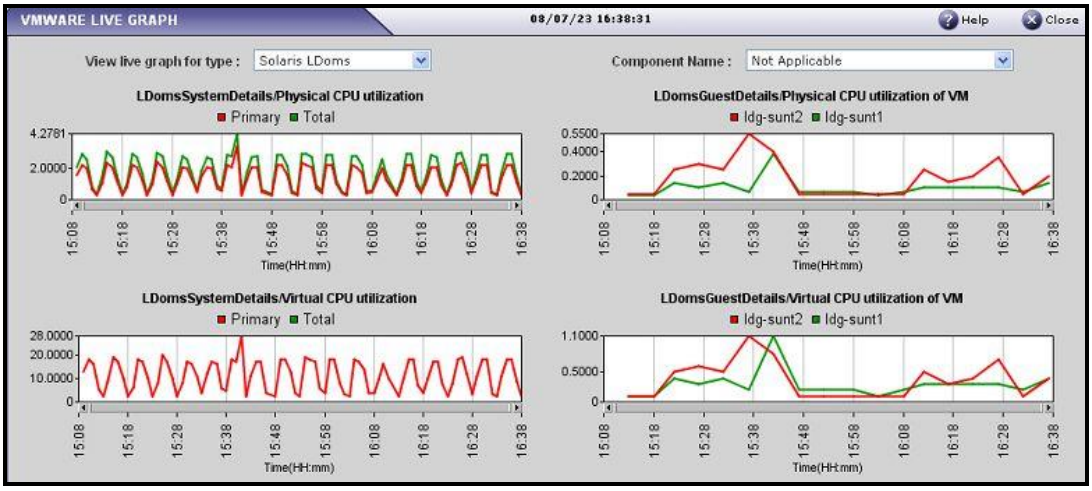

Figure 18: Live graph of the Solaris LDoms server

1.5 Correlation between Applications in a Solaris LDoms Virtualized Environment

Using the eG Enterprise administration console, administrators can add applications running on the logical domains for monitoring. To monitor these applications, agents can be installed in the guests, or an agentless monitoring approach can be used. To effectively monitor the applications running in a virtual environment, it is important to be able to determine which Solaris LDoms server an application is running on. This mapping of applications to virtual servers is important for root-cause diagnosis – for example, a problem with the virtual server (e.g., excessive disk slowdowns) can impact the performance of all the applications running on the server's virtual machines.

eG Enterprise is able to automatically determine the mapping of applications to virtual servers. Whether eG Enterprise automatically determines the mapping of applications to 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 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 LDoms* server is added, then, 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 LDoms* server.

The mapping of applications to virtual servers is used by eG Enterprise for correlation – e.g., since the application runs on the virtual server, it is most likely that a problem with the virtual server will impact the performance of the application running on one of the guests. To view this application-virtual server association, simply click on the **VIRTUAL COMPONENTS** link in the layer model page of the virtual server.

Note:

The **VIRTUAL COMPONENTS** link will also be available in the layer model page of those server applications that are executing on virtual guests.

Doing so reveals Figure 19 depicting the *Solaris LDoms* server and the server applications executing on it. By clicking on any of the components in Figure 19, the user can drill down into specific layers of this component for specific details on the performance of the component.

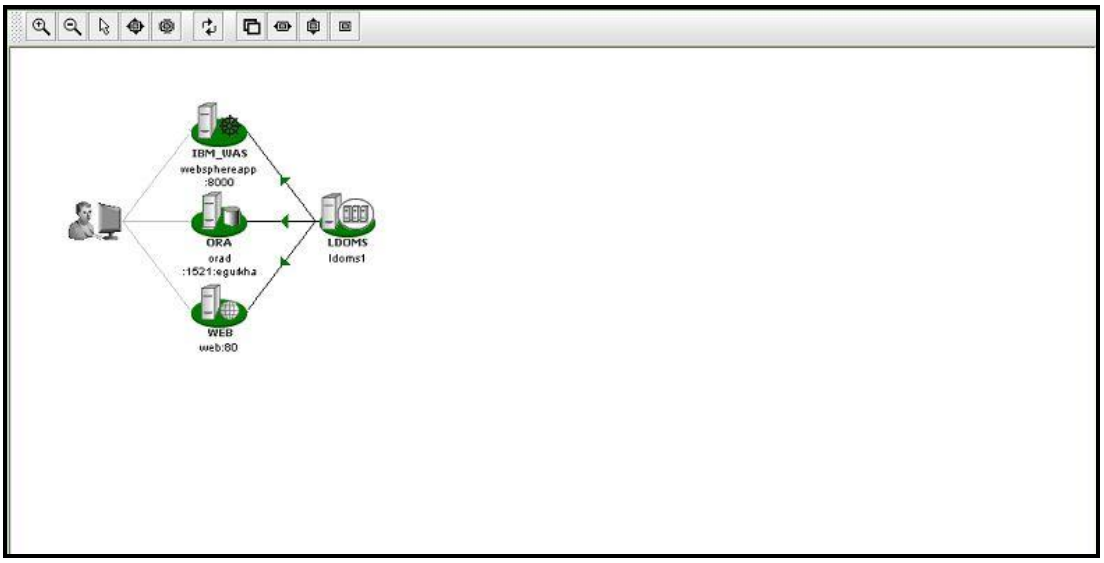

Figure 19: Depicts the applications that have been deployed on the guest OS of a virtual server

The arrows in Figure 19 depict the dependencies between the virtual server host and the applications running on it. Since the applications are hosted on one of the guests running on the host, they depend on the virtual server host – i.e., any unusual resource usage on the virtual server host impacts the applications running on any of the virtual guests. The dependency information between the virtual server host and the applications hosted on it is used by eG Enterprise for end-to-end correlation.

1.6 Troubleshooting Solaris LDoms Monitoring

If the tests mapped to the **Inside View of VMs** layer and the **LDoms Guests** layer are not reporting measures for a particular guest domain, it could be because the eG agent on the control domain is unable to connect to that guest for extracting performance data. Such an event could occur if a **telnet** connection to that guest from the host operating system is already open. To figure out if **telnet** connections pre-exist, do the following:

1. Login to the Solaris LDoms server as a **root** user.
2. First, determine the port number at which the guest listens by issuing the following command on the host operating system: **ldm ls**
3. This will list all the guest domains on the Solaris server and their corresponding port numbers. Figure out the port number of the problem guest domain from this list.
4. Then, from the host operating system, issue the following command:

   telnet localhost <port_number_problemguestdomain>
5. If this command, upon execution, returns the error message depicted by Figure 20, it is a clear indication that one/more **telnet** connections to the guest pre-exist.

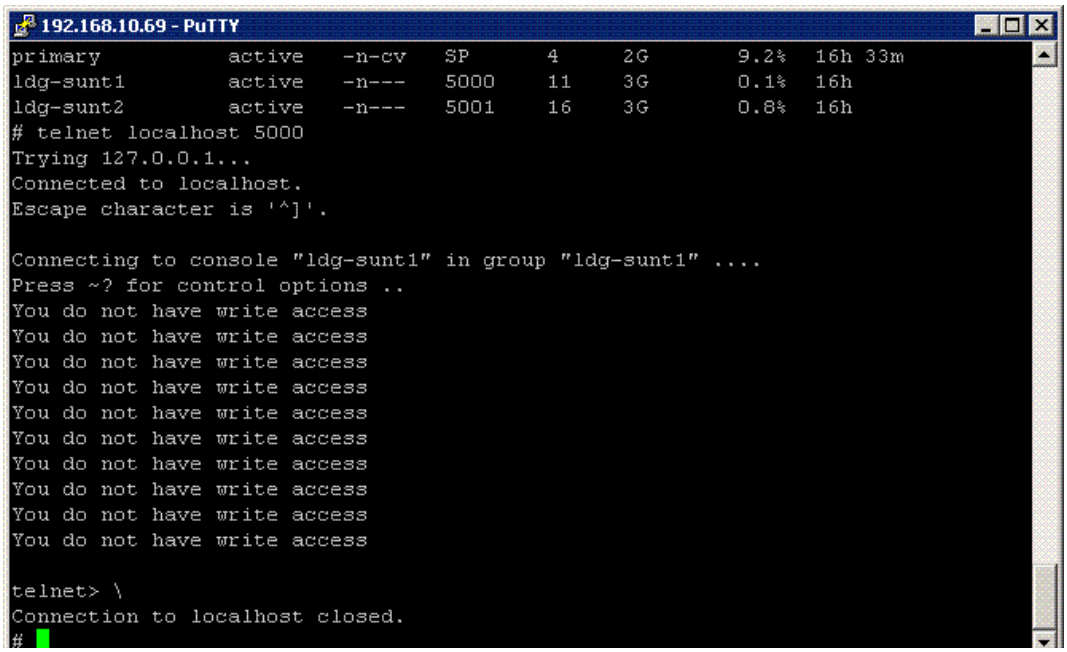

Figure 20: Error message

6. To ensure that the eG agent connects to the guest domain, close all the open **telnet** connections to the guest.

Chapter 2

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.