Monitoring DNS, LDAP, and FTP Servers

eG Enterprise v6

Restricted Rights Legend

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

Trademarks

Microsoft Windows, Windows NT, Windows 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

Table of Contents

Table of Figures

Chapter 1

Introduction

Although users do not directly interact with infrastructure servers such as DNS (Domain Name Service) and LDAP (Lightweight Directory Access Protocol) servers, the e-business applications rely on one or more of these infrastructure servers for their operation. While the DNS servers are used to translate host names to IP addresses and vice versa, LDAP servers are used to maintain access rights, policies, and other relevant information for an organization. The FTP servers, on the other hand, are used for sharing critical data across multiple geographies.

Consequently, the unavailability or slow down of any of these infrastructure servers can impact the performance of the e-business applications. To monitor these infrastructure servers, the eG Enterprise suite uses a combination of external and internal tests.

This document sheds light on these internal and external tests, and the metrics they collect.

Chapter 2

Monitoring the DNS Servers

Domain Name System (DNS) is the name resolution protocol for TCP/IP networks, such as the Internet. Client computers query a DNS server to resolve memorable, alphanumeric DNS names to the IP addresses that computers use to communicate with each other.

Imagine a situation where the DNS server is rendered temporarily unavailable. If a client computer attempts to send across a critical information request to a server at this time, the attempt is sure to fail due to the absence of the DNS server to translate the human-readable DNS name to a machine-readable IP address. In an environment where there is continuous exchange of data between components, such unplanned DNS server failures can result in total chaos!

In order to avoid such problem conditions, the performance of the DNS server should be constantly monitored.

eG Enterprise prescribes a specialized *DNS* server monitoring model (see Figure 2.1), which executes tests on the DNS server at pre-configured intervals to determine the following:

- Resource usage levels of the DNS host
- The TCP connection load on the host
- The health of the network traffic to and from the host
- The availability of the DNS server, and its responsiveness to user requests

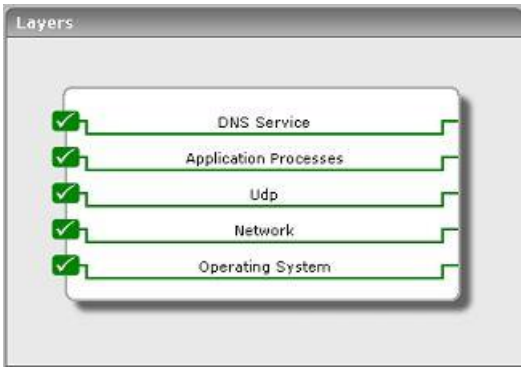

Figure 2.1: eG Enterprise's model of a DNS server

Since the DNS service is based on the UDP protocol, the layer model includes measures of the status of the UDP stack of a host. The **Application Processes** layer tracks the health of the processes corresponding to the DNS server. On Unix systems, the "named" process supports the DNS service.

Since the bottom 4 layers of Figure 2.1 have been extensively discussed in the *Monitoring Unix and Windows Servers* document, let us focus only on the **DNS Service** layer.

2.1 The DNS Service Layer

In the above figure, the **DNS Service** layer tracks the health of the DNS service. To measure the state of a DNS server, the eG agent uses a Dns test shown in Figure 2.2.

Figure 2.2: Tests mapping to the DNS Service layer

2.1.1 Name Resolutions Test

This test emulates a client accessing a DNS server to issue a query. The query can either request the DNS server to resolve a domain name to an IP address or vice versa. Based on the response reported by the server, measurements are made of the availability and responsiveness of the DNS server.

The DNS service is organized hierarchically, i.e., one DNS server can forward a client request to another server to resolve the client's query. To ensure that the results of the query reflect the state of a DNS server in isolation, a non-recursive query is issued by this test.

Purpose	To measure the state of a DNS server
Target of the test	A DNS server
Agent deploying the test	An external agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">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 on which the specified host is listening
4. TARGETS - The IP address or host name to be resolved during the test. Multiple TARGETS can be specified as a comma-separated list.
5. RECURSIVE - A DNS server supports two types of queries. For a non-recursive query, the DNS server attempts to respond to the request based on its local cache only. For a recursive query, a DNS server may use other DNS servers to respond to a request. The Recursive flag can be used to determine the type of queries to be issued to a DNS server.
6. USEEXE - In older versions of the eG Enterprise Suite, this test used native APIs to collect the desired metrics. To ensure backward compatability with older versions of the solution, this flag has been set to Yes by default. Set this flag to No if you want the test to use Java APIs instead to determine the availability and responsiveness of the DNS server. This flag is only relevant if the test is being executed by an external agent operating on a Windows host.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results per TARGET configured</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>DNS availability:
Whether a successful response is received from the DNS server in response to the emulated user request.</td><td>Percent</td><td>An availability problem can be caused by different factors – e.g., the server process may not be up, a network problem may exist, or there could be a configuration problem with the DNS server.</td></tr>
<tr><td>DNS response time:
Time taken (in seconds) by the server to respond to a request.</td><td>Secs</td><td>An increase in response time can be caused by several factors such as a server bottleneck, a configuration problem with the DNS server, a network problem, etc.</td></tr>
</table>

Chapter 3

Monitoring Windows DNS Servers

The eG Enterprise suite automatically discovers DNS servers running on Windows environments. For such servers, the eG Enterprise suite prescribes an exclusive *Windows DNS* monitoring model depicted by Figure 3.1 below:

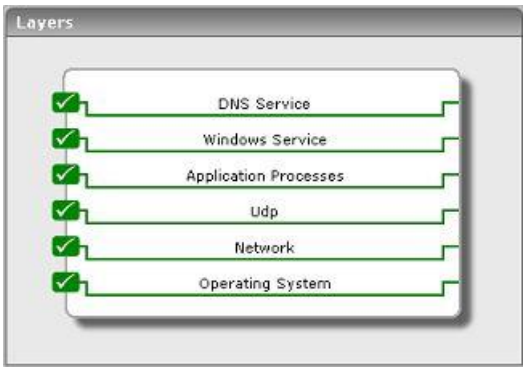

Figure 3.1: Layer model of the Windows DNS server

The additional **Windows Service** layer in Figure 3.1 reveals whether/not the critical DNS service is up and running on the Windows host. This layer and all the 4 layers below it have been discussed extensively in the *Monitoring Unix and Windows Servers* document. The section to come therefore talks only of the **DNS Service** layer.

3.1 The DNS Service Layer

Besides the **Dns** test that is common to both the DNS and Windows DNS servers, the **DNS Service** layer of a Windows DNS server is mapped to two additional tests – the WindowsDns test and the DNSEvt test (see Figure 3.2).

Figure 3.2: The DNS Service layer of a Windows DNS server

3.1.1 Windows Dns Test

This test reports various performance statistics pertaining to the DNS server running on Windows.

Purpose	Reports various statistics pertaining to the DNS server running on Windows		
Target of the test	A DNS server running on Windows		
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 –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every DNS server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total queries: The rate of queries received by the DNS server	Reqs/sec	Indicates the workload of the DNS server
	Total responses: The rate of responses from the DNS server to clients	Resp/sec	Ideally, the total responses should match the total queries. Significant differences between the two can indicate that the DNS server is not able to handle the current workload.
	Recursive queries: The rate of recursive queries successfully handled by the DNS server	Reqs/sec	The ratio of recursive queries to total queries indicates the number of queries that required the DNS server to communicate with other DNS servers to resolve the client requests.

	Recursive query failures: The rate of recursive queries that could not be resolved by the DNS server	Reqs/sec	Query failures can happen due to various reasons - e.g., requests from clients to invalid domain names/IP addresses, failure in the external network link thereby preventing a DNS server from communicating with other DNS servers on the Internet, failure of a specific DNS server to which a DNS server is forwarding all its requests, etc. A small percentage of failures is to be expected in any production environment. If a significant percentage of failures are happening, this could result in application failures due to DNS errors.
	Recursive timeouts: The rate of recursive queries that failed because of timeouts	Reqs/sec	Timeouts can happen because of a poor external link preventing a DNS server from communicating with others. In some cases, improper/invalid domain name resolution requests can also result in timeouts. DNS timeouts can adversely affect application performance and must be monitored continuously.
	Zone transfers received: The number of zone transfer requests received by a DNS	Reqs	Zone transfers are resource intensive. Moreover, zone transfers to unauthorized clients can make an IT environment vulnerable to security attacks. Hence, it is important to monitor the number of zone transfer requests and responses on a periodic basis.
	Zone transfers failed: The number of zone transfers that were not serviced by the DNS server in the last measurement period	Reqs	Zone transfers may fail either because the DNS server does not have resources, or the request is not valid, or the client requesting the transfer is not authorized to receive the results.

3.1.2 DNS Events Test

This test reports statistical information about the DNS Service events recorded in the DNS Service event log.

Purpose	Reports statistical information about the DNS Service events recorded in the DNS Service event log
Target of the test	A Windows DNS server
Agent deploying the	An internal agent

test	

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 –** Refers to the port used by the EventLog Service. Here it is null. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. The default value is *application*. 5. **POLICY BASED FILTER** - Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: ➢ Manually specify the event sources, IDs, and descriptions in the **FILTER** text area, or, ➢ Select a specification from the predefined filter policies listed in the **FILTER** box For explicit, manual specification of the filter conditions, select the **NO** option against the **POLICY BASED FILTER** field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **YES** option against the **POLICY BASED FILTER** field. 6. **FILTER** - If the **POLICY BASED FILTER** flag is set to **NO**, then a **FILTER** text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the **FILTER** text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here: ➢ *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; ➢ *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. ➢ Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. ➢ In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring.

- Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring.
- The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters.
- In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring.

By default, the **FILTER** parameter contains the value*: all:all:none:all:none:all:none*. Multiple filters are to be separated by semi-colons (;).

Note:

The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window.

On the other hand, if the **POLICY BASED FILTER** flag is set to **YES**, then a **FILTER** list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format:

{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}

To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the **FILTER** list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a **Click here** link appears just above the test configuration section, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the **Click here** link leads you to a page where you can modify the existing policies or create a new one. The changed policy or the new policy can then be associated with the test by selecting the policy name from the **FILTER** list box in this page.

<table>
<tr><td></td><td colspan="3">7. USEWMI - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the USEWMI flag is YES, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows 2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the USEWMI parameter value to NO.
8. 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.
9. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.
The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
• The eG manager license should allow the detailed diagnosis capability
• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the FILTER configured</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>DNS Errors:
This refers to the number of DNS Service events that were generated.</td><td>Number</td><td>A very low value (zero) indicates that the DNS Service is in a healthy state without any potential problems.
An increasing trend or high value indicates the existence of problems like loss of functionality or data.
The detailed diagnosis capability, if enabled, lists the description of specific events.
Please check the Application Logs in the Event Log Viewer for more details.</td></tr>
<tr><td>DNS information count:
This refers to the number of DNS Service information events generated when the test was last executed.</td><td>Number</td><td>A change in the value of this measure may indicate infrequent but successful operations performed by the DNS Service.
The detailed diagnosis capability, if enabled, lists the description of specific events.</td></tr>
</table>

	DNS Warnings: This refers to the number of warnings that were generated when the test was last executed.	Number	A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems in the DNS Service. The detailed diagnosis capability, if enabled, lists the description of specific events.

3.1.3 DNS Server Checks Test

If the DNS server is inaccessible or is unable to provide domain name resolution services, then users may be denied access to their mission-critical servers and applications. Under such circumstances, you may want to quickly check what is stalling the operations of your DNS server, so that the source of the issue can be isolated and eliminated. This test enables you to perform such a check, periodically. To perform this check, this test uses the **DCDIAG** utility that ships with Windows 2003 Support Tools and is built into Windows 2008 R2 and Windows 2008 Server. **DCDIAG** is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in a system. For validating DNS health, the **DCDIAG** utility runs six tests, each of which reports the current state of a critical performance aspect of the DNS server; these DNS tests are as follows:

a. **Authentication**: This test is run by default and checks the following:
    - Are domain controllers registered in DNS?
    - Can they be pinged?
    - Do they have Lightweight Directory Access Protocol/Remote Procedure Call (LDAP/RPC)?

b. **Basic**: Performs basic DNS tests, including network connectivity, DNS client configuration, service availability, and zone existence.

c. **Forwarders**: Performs the **Basic** tests, and also checks the configuration of forwarders

d. **Delegation**: Performs the **Basic** tests, and also checks for proper delegations

e. **Dynamic Update**: Performs the **Basic** tests, and also determines if dynamic update is enabled in the Active Directory zone

f. **Record Registration**: Performs the **Basic** tests, and also checks if the address (A), canonical name (CNAME) and well-known service (SRV) resource records are registered. In addition, creates an inventory report based on the test results.

The **DNS Server Checks** test uses the **DCDIAG.exe** to execute each of the above-mentioned tests at configured intervals, reports the output of each test, promptly captures current/potential DNS failures, and provides detailed diagnostics describing the reasons for the failure. This way, administrators are enabled to troubleshoot DNS-related issues quickly and efficiently.

Note:

For this test to run, the **DCDIAG.exe** should be available in the **<WINDOWS_INSTALL_DIR>\windows\system32** directory of the DNS server to be monitored. The **DCDIAG** utility ships with the Windows Server 2003 Support Tools and is built into Windows 2008 R2 and Windows Server 2008. This utility may hence not be available in older versions of the Windows operating sytem. When monitoring the AD server on such Windows hosts, this test will run only if the **DCDIAG.exe** is copied from the **<WINDOWS_INSTALL_DIR>\windows\system32** directory of any Windows 2003 (or higher) host in the environment to the same directory on the target host.

Purpose	Reports the current state of the server, promptly captures DNS failures, and provides detailed diagnostics describing the reasons for the failure		
Target of the test	A DNS server		
Agent deploying the test	An internal/remote agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port on which the specified host is listening 4. **DOMAIN**, **USERNAME**, **PASSWORD**, and **CONFIRM PASSWORD** - In order to execute the **DCDIAG** command, the eG agent has to be configured with *Enterprise Admin* privileges. Therefore, specify the domain name and login credentials of a user who has been assigned the *Enterprise Admin* account in the **DOMAIN**, **USERNAME** and **PASSWORD** text boxes. Confirm the **PASSWORD** you provide by retyping it in the **CONFIRM PASSWORD** text box. 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 every test that **DCDIAG** executes		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td>test</td>
<td>Status:
Reports the output returned by this test.</td>
<td></td>
<td>Each test that DCDIAG runs will report one of the following values as the output:
• Fail
• Pass
• Warning
The DNS Server Checks test will report the same output as the value of the Status measure.
The numeric values that correspond to these outputs are indicated below:
<table><tr><th>Output/Measure name</th><th>Numeric value</th></tr><tr><td>Fail</td><td>0</td></tr><tr><td>Pass</td><td>1</td></tr><tr><td>Warning</td><td>2</td></tr></table>Note:
By default, this measure reports the Output/Measure Values listed in the table above as values of the Status measure. In the graph of the Status measure however, these measure values are represented using their numeric equivalents only - i.e., 0 to 2.
You can use the detailed diagnosis of this measure to view detailed descriptions of failures (if any). This information will help in investigating the reasons for the failure and fixing them.</td>
</tr>
</table>

Chapter 4

Externally Monitoring DNS Servers

eG Enterprise offers the *DNS* server or the *Windows DNS* server model (discussed previously), which not only checks how well the DNS server performs host name resolutions, but also indicates how healthy the DNS server host is by reporting a wide variety of operating system-level metrics. However, some administrators might not have access to the DNS server for installing agents. To enable such administrators to deploy an eG agent on a remote host to monitor just the availability of the DNS server, and determine how quickly the server can resolve a host name to an IP address or vice-versa, eG Enterprise offers an *External DNS* server model (see Figure 4.1).

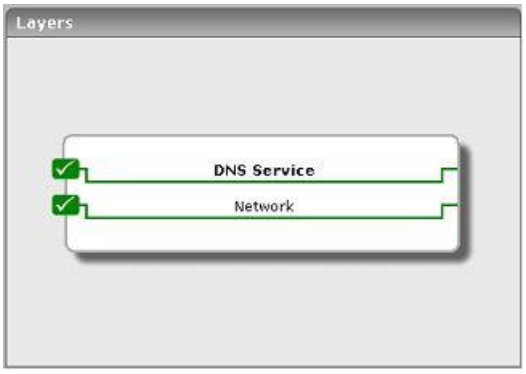

Figure 4.1: Layer model of the External DNS server

The **DNS Service** layer of this model uses an external agent to execute a Dns test, which emulates a user request to the DNS server to ascertain its availability and responsiveness. The **Network** test associated with the **Network** layer performs periodic network health checks.

Chapter 5

Monitoring LDAP Servers

LDAP, Lightweight Directory Access Protocol, is an Internet protocol that email and other programs use to look up information from a server. "LDAP-aware" client programs can ask LDAP servers to look up entries in a wide variety of ways. LDAP servers index all the data in their entries, and "filters" may be used to select just the person or group you want, and return just the information you want. LDAP is not limited to contact information, or even information about people. LDAP is used to look up encryption certificates, pointers to printers and other services on a network, and provide "single signon" where one password for a user is shared between many services. LDAP is appropriate for any kind of directory-like information, where fast lookups and less-frequent updates are the norm.

In environments providing mission-critical end-user services, the LDAP server, by enabling speedy retrieval of information, ensures that the quality of the user experience with the service is top-notch. The non-availability of the LDAP server in such infrastructures could therefore significantly slowdown service delivery, thereby impacting the overall service quality. To avoid such unpleasant consequences, it is imperative that the LDAP server's availability and responsiveness be continuously monitored.

eG Enterprise offers a 100%, web-based *LDAP* server monitoring model, which runs quick availability checks on the LDAP server at pre-set intervals, and in the process, also reports the responsiveness of the server and its overall performance.

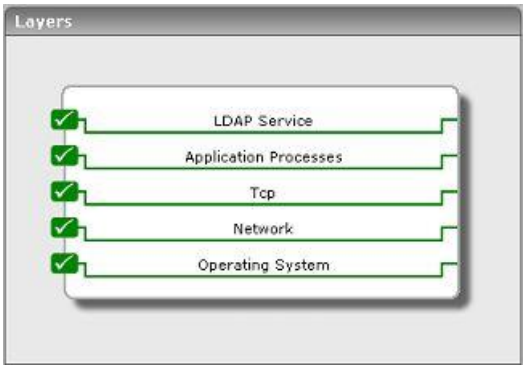

Figure 5.1: The layer model of the LDAP server

Figure 5.1 depicts the layer model of an LDAP server. Since the LDAP service is based on the TCP protocol, the layer model includes measures of the status of the UDP stack of a host. The **Application Processes** layer tracks the health of the processes corresponding to the LDAP server.

The sections to come discuss the **LDAP Service** layer only, as all other layers have already been dealt with in the *Monitoring Unix and Windows Servers* document.

5.1 The LDAP Service Layer

The **LDAP Service** layer tracks the health of the LDAP service. To measure the state of an LDAP server, the eG Enterprise suite uses an Ldap test shown in Figure 5.2.

Figure 5.2: Tests mapping to the LDAP Service layer

5.1.1 Ldap Test

This test emulates a client accessing an LDAP server to issue a query. Based on the response reported by the server, measurements are made of the availability and responsiveness of the LDAP server. Many LDAP servers have been designed to report a number of usage statistics if the query is the special string **"cn=monitor"**. For such LDAP servers, this test also reports a number of usage statistics.

Purpose	To measure the state of an LDAP server
Target of the test	An LDAP server
Agent deploying the test	An external agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">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 on which the specified host is listening
4. SEARCH - is the Distinguished Name to search for in the LDAP server. By default, this value is set to "cn=monitor". Many LDAP servers (iPlanet/Sun One LDAP, Open LDAP, etc.) expose performance metrics when this DN is used. However, for more recent LDAP versions, the distinguished name has to be represented as a sequence of relative distinguished names (RDN) connected by commas. For instance, the Search parameter can be configured as: cn=eguser,cn=Users,dc=citrix,dc=eGinnovations,dc=com
5. DISTINGUISHEDNAME – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required. Where authentication is required, the DistinguishedName has to be represented as a sequence of relative distinguished names (RDN) connected by commas. For instance, your specification can be: cn=ctxuser,cn=Users,dc=citrix,dc=eGinnovations,dc=com
6. PASSWORD – Password to be used for authenticating the request. The password is to be specified whenever the DISTINGUISHEDNAME is not "none".
7. CONFIRM PASSWORD – Confirm the PASSWORD (if specified) by retyping it here.
8. ISPASSIVE - If the value chosen against this parameter is YES, then the LDAP server under consideration is a passive server in an LDAP cluster. No alerts will be generated if the server is not running. Measures will be reported as “Not applicable” by the agent if the server is not up.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results per LDAP server 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>LDAP availability:
Whether a successful response is received from the LDAP server in response to the emulated user request.</td><td>Percent</td><td>An availability problem can be caused by different factors – e.g., the server process may not be up, a network problem may exist, or there could be a configuration problem with the LDAP server.</td></tr>
<tr><td>LDAP response time:
Time taken (in seconds) by the server to respond to a request.</td><td>Secs</td><td>An increase in response time can be caused by several factors such as a server bottleneck, a configuration problem with the LDAP server, a network problem, etc.</td></tr>
<tr><td>Current LDAP connections:
Number of connections currently being processed by the LDAP server.</td><td>Number</td><td>A high value could result whenever the server is experiencing a problem (due to overload, or because of application problems).</td></tr>
<tr><td>LDAP connection rate:
Quantifies the workload in terms of connections per second to the LDAP server.</td><td>Conns/Sec</td><td>This value directly represents the user workload.</td></tr>
</table>

	Operations outstanding: The number of outstanding requests waiting for processing by the LDAP server.	Number	A consistent non-zero value of this metric is indicative of a server bottleneck.
	Data transmit rate: Quantifies the traffic handled by the LDAP server in Kbytes/Sec.	KB/Sec	Typically, an increase or decrease in connection rate will result in a corresponding change in the data transmission rate. A deviation from this rule signifies a possible change in the characteristics of applications accessing the LDAP server, or a change in the organization of the LDAP schema.
	LDAP TCP port availability: This measure indicates whether the test managed to establish a TCP connection to the server.	Percent	1. Failure to establish a TCP connection may imply that either the web server process is not up, or that the process is not operating correctly. In some cases of extreme overload, the failure to establish a TCP connection may be a transient condition. As the load subsides, the server may start functioning properly again. 2. If this measure is 100% but the *LDAP availability* is 0, this could indicate a problem with the server configuration. Alternatively, the search string provided as a parameter to the LdapTest may not be supported by the corresponding LDAP server. In this case, configure the test based on the directory structure configured for the target LDAP server (i.e., change the "cn=monitor" value with an appropriate alternative).

Note:

The Processes test of LDAP servers takes an additional parameter named ispassive. If the value thosen against this parameter is **YES**, then the LDAP server under consideration is a passive server in an LDAP cluster. No alerts will be generated if the server is not running. Measures will be reported as "Not applicable" by the agent if the server is not up.

Chapter 6

Monitoring the Netscape Directory Server

Netscape Directory Server is an LDAP server that centralizes application settings, user profiles, group data, policies, and access control information into a network-based registry. Directory Server simplifies user management by eliminating data redundancy and automating data maintenance. It also improves security, enabling administrators to store policies and access control information in the directory for a single authentication source across enterprise or extranet applications.

To enable the uninterrupted use of critical applications, the Directory server that controls access to those applications has to be available and operating at peak capacity 24x7. If the Directory server experiences delays in authenticating access requests the user experience with the server and the application will suffer. To avoid this, the availability and overall performance of the Netscape Directory server should be periodically monitored.

eG Enterprise provides a specialized *Netscape Directory Server* model for monitoring the Netscape Directory server.

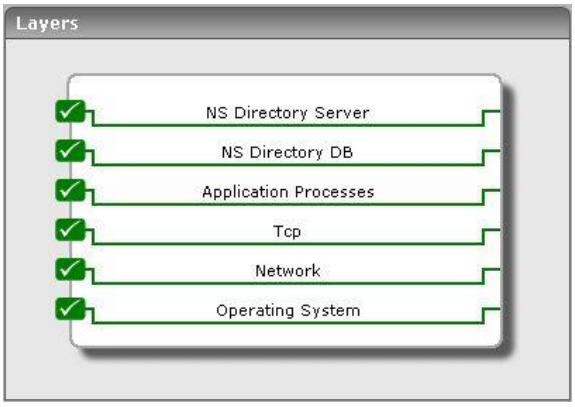

Figure 6.1: The layer model of the SunONE LDAP/Netscape Directory server

Using the model depicted by Figure 6.1, the following can be monitored:

- External monitoring of the directory service health including server availability and response time for a typical request

- Server usage metrics such as connection rate, data transfer rate, operations backlog, current connections to the server, etc

This section will discuss the tests associated with the **NS Directory DB** and **NS Directory Server** layers of Figure 6.1.

6.1 The NS Directory DB Layer

You can judge the effectiveness of the caching activity performed by the directory server using the tests associated with this layer.

Figure 6.2: The tests associated with the NS Directory DB layer

6.1.1 NsDbCache Test

The NsDbCache test reports measures pertaining to the caching activity performed by the Netscape Directory server.

Purpose	Reports measures pertaining to the caching activity performed by the Netscape Directory server
Target of the test	A SunONE LDAP / Netscape Directory 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 on which the specified host is listening 4. **DISTINGUISHEDNAME** – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required. 5. **PASSWORD** – Password to be used for authenticating the request. The password is to be specified whenever the **DISTINGUISHEDNAME** is not "none". 6. **CONFIRM PASSWORD** – Confirm the **PASSWORD** (if specified) by retyping it here.
Outputs of the test	One set of results per LDAP server monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Cache hit ratio: The ratio of database cache hits to database cache tries.	Percent	The closer this value is to 100%, the better. Whenever a directory operation attempts to find a portion of the database that is not resident in the database cache, the directory server has to perform a disk access to obtain the appropriate database page. Thus, as this ratio drops towards zero, the number of disk accesses increases and directory server performance drops.
	Pages read in: The rate of pages read from disk into the database cache during the last measurement period.	Pages/Sec	
	Pages written out: The rate of pages for this file written from cache to disk during the last measurement period.	Pages/Sec	A database page is written out to disk whenever a read-write page has been modified and then subsequently evicted from the cache. Pages are evicted from the database cache when the cache is full and a directory operation requires a database page that is not currently stored in cache.

6.1.2 NsDbFileCache Test

The NsDbFileCache test reports measures pertaining to each of the index files that make up the Netscape Directory server database.

Purpose	Reports measures pertaining to each of the index files that make up the Netscape Directory server database
Target of the test	A SunONE LDAP / Netscape Directory 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 on which the specified host is listening 4. **DISTINGUISHEDNAME** – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required. 5. **PASSWORD** – Password to be used for authenticating the request. The password is to be specified whenever the **DISTINGUISHEDNAME** is not "none". 6. **CONFIRM PASSWORD** – Confirm the **PASSWORD** (if specified) by retyping it here.		
Outputs of the test	One set of results per LDAP server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Cache hit ratio: The percentage of times a search result resulted in a cache hit.	Percent	
	Pages read in: The rate of pages brought to the cache from this file during the last measurement period.	Pages/Sec	
	Pages written out: The rate of pages for this file written from cache to disk file during the last measurement period.	Pages/Sec	

6.1.3 NsEntryCache Test

The NsEntryCache test reports measures pertaining to the entry caches of the Netscape Directory server.

Purpose	Reports measures pertaining to pertaining to the entry caches of the Netscape Directory server
Target of the test	A SunONE LDAP / Netscape Directory server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">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 on which the specified host is listening
4. DISTINGUISHEDNAME – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required.
5. PASSWORD – Password to be used for authenticating the request. The password is to be specified whenever the DISTINGUISHEDNAME is not "none".
6. CONFIRM PASSWORD – Confirm the PASSWORD (if specified) by retyping it here.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results per LDAP server 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>Cache hit ratio:
The ratio of the number of entry cache tries to successful entry cache lookups.</td><td>Percent</td><td>This number is based on the total lookups and hits since the last measurement period of the test. The closer this value is to 100% the better. Whenever a search operation attempts to find an entry that is not resident in the entry cache, the directory server has to perform a disk access to obtain the entry. Thus, as this ratio drops towards zero, the number of disk accesses increases and directory server search performance drops.</td></tr>
<tr><td>Entries present:
The number of directory entries currently resident in the entry cache.</td><td>Number</td><td></td></tr>
</table>

6.2 The NS Directory Server Layer

The availability, responsiveness, and the request processing capability of the directory server can be analyzed and ascertained using the measures reported by the tests associated with this layer.

Figure 6.3: The test associated with the NS Directory Server layer

6.2.1 NsDirectory Test

The NsDirectory test emulates a client accessing a Netscape Directory Server to issue a query. Based on the response reported by the server, measurements are made of the availability and responsiveness of the Directory server and also reports a number of usage statistics.

Purpose	Reports measures pertaining to pertaining to the entry caches of the Netscape Directory server		
Target of the test	A SunONE LDAP / Netscape Directory server		
Agent deploying the test	An external 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 on which the specified host is listening 4. **DISTINGUISHEDNAME** – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required. 5. **PASSWORD** – Password to be used for authenticating the request. The password is to be specified whenever the **DISTINGUISHEDNAME** is not "none". 6. **CONFIRM PASSWORD** – Confirm the **PASSWORD** (if specified) by retyping it here.		
Outputs of the test	One set of results per LDAP server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Availability: Whether a successful response is received from the Directory server in response to the emulated user request.	Percent	An availability problem can be caused by different factors - e.g., the server process may not be up, a network problem may exist, or there could be a configuration problem with the Directory server.

	Response time: Time taken (in seconds) by the server to respond to a request.	Secs	An increase in response time can be caused by several factors such as a server bottleneck, a configuration problem with the Directory server, a network problem, etc.
	Tcp connection availability: This measure indicates whether the test managed to establish a TCP connection to the server.	Percent	While the value 100 indicates that a TCP connection has been successfully established, 0 indicates that the connection attempt has failed.
	Data sent: The rate of data being transmitted by the server to clients during the last measurement period.	KB/Sec	
	Entries sent: The rate of entries being transmitted by the server to clients during the last measurement period	KB/Sec	
	Active threads: The current number of active threads used for handling requests. Additional threads may also be created by internal server tasks, such as replication, or writing to logs.	Number	
	Current connections: The number of connections currently in service by the directory server	Number	
	Connections handled: Quantifies the workload in terms of connections handled by the directory server per second.	Conns/Sec	This value directly represents the user workload.

	Ops initiated user: The rate of operations the server has initiated during the last measurement period. Operations include any client requests for server action, such as searches, adds, and modifies in the directory tree. It is likely that multiple operations will be initiated for each connection.	Operations/Sec	
	Ops completed rate: The rate of operations the server has completed during the last measurement period.	Operations/Sec	
	Outstanding operations: The number of outstanding operations waiting for processing by the Directory server.	Number	

Chapter 7

Monitoring the SunONE Directory Server

The Sun ONE Application suite offers a comprehensive list of products for Internet infrastructures, i.e., web server, middleware application server, LDAP server, messaging server, and identity server, that are used in many domains such as banking, trading, healthcare, and logistics to support mission-critical services. IT infrastructures based on the Sun ONE Application suite follow the popular multi-tier architecture wherein the web server functions as the front-end receiving client requests, the application server hosts the business logic components, the identity server manages user policies, the directory server handles access rights and other user information lookups, and the database server stores and retrieves application data.

Routine monitoring of the infrastructure including the network, system, and application is imperative to ensure that the infrastructure functions at peak performance at all times. Since each Sun ONE application performs a different, specialized function, the monitoring has to be specific to each application – e.g., is the mail server delivering emails? is the application server's heap effectively sized?. More importantly, since the different Sun ONE applications inter-operate to support the end-user service, it is critical that the monitoring system track the inter-dependencies between applications in order to pin-point the exact source of a performance bottleneck in the infrastructure.

The eG Sun ONE monitor offers extensive infrastructure monitoring capabilities for the Sun ONE application suite. Pre-built models for Sun ONE web, application, directory, and messaging servers dictate what metrics are to be collected by eG agents, what thresholds are to be applied to the metrics, and how the metrics are to be correlated in order to assist with problem diagnosis.

Figure 7.1 depicts the *SunONE Directory Server* monitoring model offered by the eG Enterprise Suite.

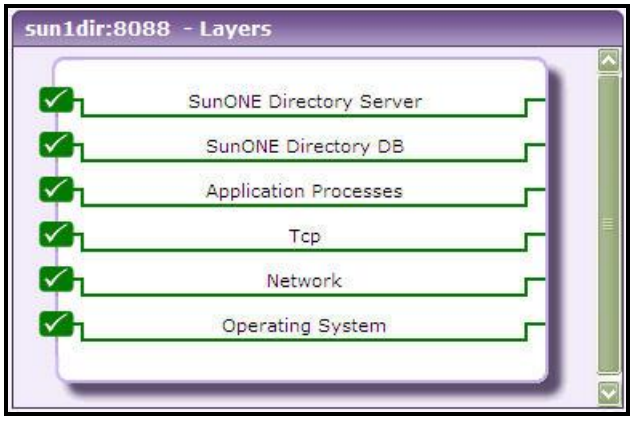

Figure 7.1: The layer model of the SunONE LDAP/Netscape Directory server

Using the model depicted by Figure 7.1, the following can be monitored:

- Is the directory server available? If so, how quickly is it reponding to user requests?
- Are the database and entry caches optimally utilized? Are too many requests to any database been fulfilled by direct disk accesses?
- Is the Directory server overloaded with connection requests?
- Are too many operations awaiting processing by the server?

This section will discuss the tests associated with the **SunONE Directory DB** and **SunONE Directory Server** layers of Figure 7.1.

7.1 The SunONE Directory DB Layer

You can judge the effectiveness of the caching activity performed by the directory server using the tests associated with this layer.

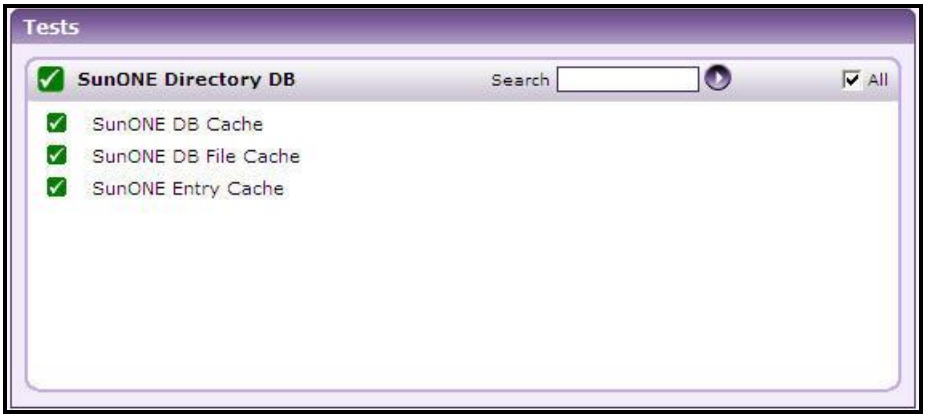

Figure 6.4: The tests associated with the SunONE Directory DB layer

7.1.1 SunONE DB Cache Test

Each Directory Server instance has one database cache. The database cache is a pool of memory that holds pages from the database containing indexes and entries.

The size of the database cache is configurable. The cache should be rightly sized in order to prevent/at least minimize expensive and potentially harmful direct disk accesses. With the help of this test, you can monitor how the cache is utiilized over time, promptly detect sizing irregularities, and correct them.

<table>
<tr><td>Purpose</td><td colspan="3">Monitors how the cache is utiilized over time, promptly detects sizing irregularities, and helps correct them</td></tr>
<tr><td>Target of the test</td><td colspan="3">A SunONE Directory 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. 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 on which the specified host is listening
4. DISTINGUISHEDNAME – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required.
5. PASSWORD – Password to be used for authenticating the request. The password is to be specified whenever the DISTINGUISHEDNAME is not "none".
6. CONFIRM PASSWORD – Confirm the PASSWORD (if specified) by retyping it here.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results per Directory server 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>Cache hits:
Indicates the number of times the Directory server has successfully processed a search request to retrieve data from the database cache.</td><td>Number</td><td></td></tr>
<tr><td>Cache tries:
Indicates the number of times the Directory server has been looked in to retrieve data from the database cache.</td><td>Number</td><td></td></tr>
</table>

	Cache hit ratio: Indicates the ratio of database cache hits to database cache tries.	Percent	The closer this value is to 100%, the better. Whenever a directory operation attempts to find a portion of the database that is not resident in the database cache, the directory server has to perform a disk access to obtain the appropriate database page. Thus, as this ratio drops towards zero, the number of disk accesses increases and directory server performance drops.
	Pages read in: Indicates the rate of pages read from disk into the database cache during the last measurement period.	Pages/Sec	
	Pages written out: Indicates the rate of pages for this file written from cache to disk during the last measurement period.	Pages/Sec	A database page is written out to disk whenever a read-write page has been modified and then subsequently evicted from the cache. Pages are evicted from the database cache when the cache is full and a directory operation requires a database page that is not currently stored in cache.

7.1.2 SunONE DB File Cache Test

While the **SunONE DB Cache** test reveals the inefficient use of the database cache as a whole, using the **SunONE DB File Cache** test, you can accurately point to the specific database(s) that is under-utilizing the cache. The latter monitors requests to each database that has been explicitly configured for monitoring, and reports the percentage of these requests that were serviced by the database cache.

Purpose	Monitors requests to each of the specified databases, and reports the percentage of these requests that were serviced by the database cache
Target of the test	A SunONE Directory server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">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 on which the specified host is listening
4. DISTINGUISHEDNAME – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required.
5. PASSWORD – Password to be used for authenticating the request. The password is to be specified whenever the DISTINGUISHEDNAME is not "none".
6. CONFIRM PASSWORD – Confirm the PASSWORD (if specified) by retyping it here.
7. DATABASENAME - Specify a comma-separated list of databases to be monitored.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results per database to be 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>Cache hit ratio:
Indicates the percentage of times a search request to this database resulted in a cache hit.</td><td>Percent</td><td>The closer this value is to 100%, the better. Whenever a directory operation attempts to find a portion of the database that is not resident in the database cache, the directory server has to perform a disk access to obtain the appropriate database page. Thus, as this ratio drops towards zero, the number of disk accesses increases and directory server performance drops.
Compare the value of this measure across databases to identify the database that is not utilizing the cache well.</td></tr>
<tr><td>Pages read in:
Indicates the rate of pages brought to the cache from this database during the last measurement period.</td><td>Pages/Sec</td><td></td></tr>
<tr><td>Pages written out:
The rate of pages for this database written from cache to disk file during the last measurement period.</td><td>Pages/Sec</td><td>A database page is written out to disk whenever a read-write page has been modified and then subsequently evicted from the cache. Pages are evicted from the database cache when the cache is full and a directory operation requires a database page that is not currently stored in cache.</td></tr>
</table>

7.1.3 SunONE Entry Cache Test

The entry cache is a mechanism that uses system memory for holding entries in a manner that may be quickly accessed so that it is not necessary to decode them from the database whenever they are needed. Entry caching

mechanisms are particularly effective when used with applications that access the same entry multiple times in a sequence of operations (for example, an application which first searches to find a user entry and then binds as that user to verify a password, which is a very common usage pattern).

The entry cache size and the maximum number of entries in the cache are configurable. If these values are not set prudently, then the entry cache may not be able to hold adequate entries to serve the search requests to the Directory server; this in turn will compel the Directory server to directly access the disk for fulfilling the requests. Direct disk accesses are resource-intensive operations, and should be avoided. To do so, you need to continuously monitor how the entry cache services requests, detect sizing irregularities quickly, and correct them.

Using the **SunONE Entry Cache** test, you can track the requests to each database on the Directory server, observe how the cache handles these requests, and ascertain whether the cache size needs to be fine-tuned.

<table>
<tr><td>Purpose</td><td colspan="3">You can track the requests to each database on the Directory server, observe how the cache handles these requests, and ascertain whether the cache size needs to be fine-tuned</td></tr>
<tr><td>Target of the test</td><td colspan="3">A SunONE Directory 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. 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 on which the specified host is listening
4. DISTINGUISHEDNAME – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required.
5. PASSWORD – Password to be used for authenticating the request. The password is to be specified whenever the DISTINGUISHEDNAME is not "none".
6. CONFIRM PASSWORD – Confirm the PASSWORD (if specified) by retyping it here.
7. DATABASENAME - Specify a comma-separated list of databases to be monitored.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results per database configured for monitoring</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>Cache hit ratio:
Indicates the ratio of the number of entry cache tries for this database to successful entry cache lookups.</td><td>Percent</td><td>This number is based on the total lookups and hits since the last measurement period of the test. The closer this value is to 100% the better. Whenever a search operation attempts to find an entry that is not resident in the entry cache, the directory server has to perform a disk access to obtain the entry. Thus, as this ratio drops towards zero, the number of disk accesses increases and directory server search performance drops.</td></tr>
</table>

	Entries present: Indicates the number of directory entries for this database currently resident in the entry cache.	Number	The maximum number of entries that can be held by the cache can be configured using the *nsslapd-cachesize* parameter. If the value of this measure grows dangerously close to the *nsslapd-cachesize* setting, it indicates that the entry cache will soon lose its ability to hold new entries, after which, it may not be able to service search requests. You may want to consider tuning the *nsslapd-cachesize* setting under such circumstances, so as to avoid direct disk accesses.

7.2 The SunONE Directory Server Layer

The availability, responsiveness, and the request processing capability of the directory server can be analyzed and ascertained using the measures reported by the tests associated with this layer.

Figure 6.5: The test associated with the NS Directory Server layer

7.2.1 SuONE Directory Service Test

This test emulates a client accessing a SunONE Directory Server to issue a query. Based on the response reported by the server, measurements are made of the availability and responsiveness of the Directory server and also reports a number of usage statistics.

Purpose	Emulates a client accessing a SunONE Directory Server to issue a query; based on the response reported by the server, measurements are made of the availability and responsiveness of the Directory server and also reports a number of usage statistics
Target of the test	A SunONE Directory server
Agent deploying the test	An external 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 on which the specified host is listening 4. **DISTINGUISHEDNAME** – Represents the server's Distinguished name. This value is to be specified when the server requires explicit authentication of requests. By default, this attribute is set to "none", implying that authentication is not required. 5. **PASSWORD** – Password to be used for authenticating the request. The password is to be specified whenever the **DISTINGUISHEDNAME** is not "none". 6. **CONFIRM PASSWORD** – Confirm the **PASSWORD** (if specified) by retyping it here.		
Outputs of the test	One set of results per server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Availability: Whether a successful response is received from the Directory server in response to the emulated user request.	Percent	An availability problem can be caused by different factors - e.g., the server process may not be up, a network problem may exist, or there could be a configuration problem with the Directory server.
	Response time: Indicates the time taken (in seconds) by the server to respond to a request.	Secs	An increase in response time can be caused by several factors such as a server bottleneck, a configuration problem with the Directory server, a network problem, etc.
	Tcp connection availability: This measure indicates whether the test managed to establish a TCP connection to the server.	Percent	While the value 100 indicates that a TCP connection has been successfully established, 0 indicates that the connection attempt has failed.
	Data sent: The rate of data being transmitted by the server to clients during the last measurement period.	KB/Sec	
	Entries sent: The rate of entries being transmitted by the server to clients during the last measurement period	KB/Sec	

	Active threads: The current number of active threads used for handling requests. Additional threads may also be created by internal server tasks, such as replication, or writing to logs.	Number	
	Current connections: The number of connections currently in service by the directory server	Number	
	Connections handled: Quantifies the workload in terms of connections handled by the directory server per second.	Conns/Sec	This value directly represents the user workload.
	Ops initiated user: The rate of operations the server has initiated during the last measurement period. Operations include any client requests for server action, such as searches, adds, and modifies in the directory tree. It is likely that multiple operations will be initiated for each connection.	Operations/Sec	
	Ops completed rate: The rate of operations the server has completed during the last measurement period.	Operations/Sec	
	Outstanding operations: The number of outstanding operations waiting for processing by the Directory server.	Number	

Chapter 8

Monitoring FTP Servers

Many IT infrastructures include one or more FTP servers from where users may download static content such as registration forms, product documentations, etc. Similarly, users may upload any document/file to the FTP server for sharing content across a wide area.

If the FTP servers in an environment play repository to critical information, then users are bound to be intolerant towards brief or prolonged delays in uploading data to or downloading data from the server. To ensure that users are always assured of swift FTP access, the FTP server's performance should be periodically monitored.

eG Enterprise prescribes an *FTP* monitoring model (see Figure 8.1), that verifies the availability and response time of the FTP service at frequent intervals.

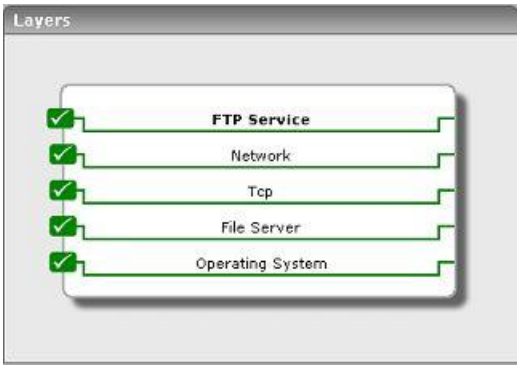

Figure 8.1: Layer model for a FTP server

The sections to come discuss the **FTP Service** layer only, as the remaining layers have been discussed elaborately

8.1 The FTP Service Layer

In the above figure, the **FTP Service** layer tracks the health of the FTP server. The status of the layer is determined by the results of an Ftp test that is shown in Figure 8.2. The The details of this test are provided below:

Figure 8.2: The Ftp test tracks the health of the FTP Service layer

8.1.1 Ftp Test

This test emulates a user connecting to an FTP server and performing an operation on the server. The operation may either be a GET or a PUT. For the specified operation, this test measures the availability of the FTP server and its response time.

Purpose	To measure statistics pertaining to an FTP server
Target of the test	An FTP server
Agent deploying the test	An external agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST – The hostname (or IP address) of the FTP server
3. PORT – The port number on which the FTP server is listening
4. USER – The user name used for connecting to the FTP server
5. PASSWORD –password corresponding to the user
6. CONFIRM PASSWORD – Confirm the PASSWORD by retyping it here.
7. REMOTEFILE – The remote file that is downloaded in the case of GET operation; In the case of a PUT operation, the remote file represents the file to which data is uploaded.
8. LOCALFILE – The local file that is written to in the case of a GET operation. In the case of a PUT operation, this string represents the name of the file that is uploaded to the FTP server. This value can be "none" if the test is not required to write the downloaded data to a local file.
9. CMD – Signifies the command to be executed by the test whether GET or PUT
10. TIMEOUT – The maximum time (in seconds) that the client will wait for a response from the FTP server</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each FTP server 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>Availability:
This measurement indicates whether the server was able to respond successfully to the query made by the test.</td><td>Percent</td><td>Availability failures could be caused by several factors such as the FTP server being down, the FTP server being misconfigured, authentication problems, file access permission problems, network failures, etc. Temporary unavailability may also occur if the FTP server is overloaded.</td></tr>
<tr><td>Total FTP response time:
This measurement indicates the total time taken by the server to respond to the requests it receives. This time includes the TCP connection time, user authentication time and the data transfer time.</td><td>Secs</td><td>An increase in the total response time can occur because there are too many simultaneous requests or because of a bottleneck with any of the applications executing on the server.</td></tr>
<tr><td>Tcp connection availability to FTP port:
This measure indicates whether the test was able to successfully establish a TCP connection to the FTP server.</td><td>Percent</td><td>Availability failures could be caused due to a network failure. Another possibility is that the FTP application server is not running.</td></tr>
<tr><td>Tcp connection time to FTP port:
The time taken for the TCP connection establishment to complete.</td><td>Secs</td><td>A high value indicates a bottleneck and could be due to the reasons that the server is being overloaded or there has been a network performance degradation.</td></tr>
</table>

	FTP authentication status: Indicates whether the test was able to successfully log in to the FTP server using the specified user account.	Percent	A low value indicates a problem logging in to the server.
	FTP authentication time: Time taken for user authentication.	Secs	This value gives an idea of where the performance bottleneck with the FTP server could be.

8.1.2 Secure FTP Test

In computing, the **SSH File Transfer Protocol** (also **Secret File Transfer Protocol**, **Secure FTP**, or **SFTP**) is a network protocol that provides file access, file transfer, and file management functionalities over any reliable data stream. This protocol assumes that it is run over a secure channel, such as SSH, that the server has already authenticated the client, and that the identity of the client user is available to the protocol.

This test emulates a user connecting to an SFTP server (on Windows/Unix) and performing an operation on the server. The operation may either be a GET or a PUT. For the specified operation, this test measures the availability of the SFTP server and its response time.

Purpose	To measure statistics pertaining to an SFTP server		
Target of the test	An SFTP server		
Agent deploying the test	An external agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The hostname (or IP address) of the SFTP server 3. **PORT** – The port number on which the SFTP server is listening 4. **USER** – The user name used for connecting to the SFTP server 5. **PASSWORD** –password corresponding to the user 6. **CONFIRM PASSWORD** – Confirm the **PASSWORD** by retyping it here. 7. **REMOTEFILE** – The remote file that is downloaded in the case of GET operation; In the case of a PUT operation, the remote file represents the file to which data is uploaded. This value can be *none* in the case of a PUT operation. 8. **REMOTEFOLDER** - The **REMOTE FOLDER** indicates the remote SFTP folder where the specified **REMOTEFILE** exists. For example, this can be */* or */sftpRoot* or */mysftpFolder*, in the case of a GET operation. In the case of a PUT operation, the **REMOTE FOLDER** represents the remote destination folder. 9. **LOCAL** – The local folder that is written to in the case of a GET operation. In the case of a PUT operation, this string represents the name of the file that is uploaded to the SFTP server. If this value is "none", then the test will write the downloaded data to the eG agent's **logs** folder. 10. **CMD** – Signifies the command to be executed by the test whether GET or PUT 11. **TIMEOUT** – The maximum time (in seconds) that the client will wait for a response from the SFTP server		
Outputs of the test	One set of results for each SFTP server monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Availability:** This measurement indicates whether the server was able to respond successfully to the query made by the test.	Percent	Availability failures could be caused by several factors such as the SFTP server being down, the SFTP server being misconfigured, authentication problems, file access permission problems, network failures, etc. Temporary unavailability may also occur if the SFTP server is overloaded.
	Total response time: This measurement indicates the total time taken by the server to respond to the requests it receives, including the data transfer time.	Secs	An increase in the total response time can occur because there are too many simultaneous requests or because of a bottleneck with any of the applications executing on the server.

Chapter 9

Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to **DNS**, **LDAP**, **and FTP servers**. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

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