



Monitoring Microsoft Applications
eG Enterprise v6

Restricted Rights Legend

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

Trademarks

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

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

Copyright

©2014 eG Innovations Inc. All rights reserved.

Table of Contents

| | |
|--|-----------|
| INTRODUCTION | 1 |
| MONITORING MICROSOFT RDS SERVERS | 2 |
| 2.1 The Windows Service Layer | 3 |
| 2.1.1 App-V Client Admin Log Test | 4 |
| 2.1.2 App-V Client Operational Log Test | 9 |
| 2.1.3 App-V Client Virtual Application Log Test | 14 |
| 2.2 The Terminal Server Layer | 19 |
| 2.2.1 Session Login Status Test | 19 |
| 2.2.2 Terminal Connection Test | 20 |
| 2.2.3 Terminal Authentication Test | 21 |
| 2.2.4 Redirector Test | 23 |
| 2.2.5 User Profile Test | 25 |
| 2.2.6 User Environment Test | 27 |
| 2.2.7 Terminal Server CALs Test | 30 |
| 2.2.8 GDI Objects Test | 32 |
| 2.3 The Terminal Applications Layer | 34 |
| 2.3.1 Terminal Applications Test | 35 |
| 2.3.2 App-V Applications Test | 38 |
| 2.4 The Terminal Users Layer | 43 |
| 2.4.1 Terminal Sessions Test | 44 |
| 2.4.2 Terminal Logins Test | 47 |
| 2.4.3 Terminal Clients Test | 49 |
| 2.4.4 Terminal Users Test | 51 |
| 2.4.5 Terminal Disconnects Test | 56 |
| 2.4.6 Rdp Client Access Test | 58 |
| 2.4.7 RemoteFX User Experience Test | 61 |
| 2.4.8 ICA/RDP Listeners Test | 66 |
| MONITORING ACTIVE DIRECTORY SERVERS | 68 |
| 3.1 The Operating System Layer | 70 |
| 3.1.1 Net Logon Test | 70 |
| 3.2 The AD Server Layer | 72 |
| 3.2.1 Asynchronous Thread Queue Test | 72 |
| 3.2.2 ADAM Access Details Test | 74 |
| 3.2.3 ADAM Database Test | 81 |

| | | |
|--------|---|-----|
| 3.2.4 | Active Directory Access Test | 82 |
| 3.2.5 | Windows Access Test | 83 |
| 3.2.6 | Windows Sessions Test..... | 85 |
| 3.2.7 | FSMO Roles Test | 86 |
| 3.2.8 | Directory System Agent Logs Test | 89 |
| 3.2.9 | Domain Controller Summary | 90 |
| 3.2.10 | Security Accounts Manager Test..... | 91 |
| 3.2.11 | Trust Relation Test | 93 |
| 3.3 | The DNS/DHCP Layer..... | 94 |
| 3.3.1 | Active Directory Checks Test..... | 95 |
| 3.3.2 | AD Checks Test..... | 97 |
| 3.3.3 | DNS Server Health Test..... | 99 |
| 3.3.4 | Name Resolutions Test | 105 |
| 3.3.5 | Windows DNS Test..... | 106 |
| 3.4 | The AD Replication Service Layer..... | 108 |
| 3.4.1 | File Replication Connections Test | 108 |
| 3.4.2 | File Replication Events Test | 110 |
| 3.4.3 | File Replication Set Test | 114 |
| 3.4.4 | Replication Performance Test..... | 117 |
| 3.4.5 | Replication Traffic from Other Sites Test | 120 |
| 3.4.6 | Replication Traffic to Other Sites Test | 121 |
| 3.4.7 | Replication Queue Test | 123 |
| 3.4.8 | Lingering Objects Test..... | 124 |
| 3.4.9 | Replication Status Test..... | 126 |
| 3.4.10 | Inter-Site Replication Test | 128 |
| 3.4.11 | Intra-Site Replication Test | 129 |
| 3.4.12 | Replication Test..... | 130 |
| 3.4.13 | AD Replications Test | 132 |
| 3.5 | The AD Service Layer..... | 134 |
| 3.5.1 | Orphaned Objects Test..... | 134 |
| 3.5.2 | Active Directory Status Test..... | 135 |
| 3.5.3 | Directory Service Events Test | 138 |
| 3.5.4 | User Account Lockouts Test | 143 |
| 3.5.5 | Active Directory Lost and Found Test..... | 146 |
| 3.5.6 | Global Catalog Search Test | 147 |
| 3.5.7 | Address Book Details Test | 148 |

| | | |
|--|--|------------|
| 3.5.8 | ADAM LDAP Performance Test | 149 |
| 3.5.9 | Authentication Performance Test | 151 |
| 3.5.10 | ADAM Binding Test | 154 |
| 3.5.11 | Global Catalogs Test | 155 |
| 3.5.12 | Active Directory Users | 156 |
| 3.5.13 | Account Management Events Test | 157 |
| 3.5.14 | Active Directory Computers Test | 166 |
| MONITORING THE BIZTALK SERVER..... | | 168 |
| 4.1 | Monitoring the BizTalk Server 2000 | 168 |
| 4.1.1 | The BTS Transport Layer | 169 |
| 4.1.2 | The BTS Documents Layer | 173 |
| 4.2 | Monitoring the BizTalk Server 2010 | 176 |
| 4.2.1 | The Messaging Engine Layer | 178 |
| 4.2.2 | The Message Box Layer | 217 |
| 4.2.3 | The Orchestration Engine Layer | 222 |
| MONITORING DHCP SERVERS | | 232 |
| 5.1 | The DHCP Services Layer | 233 |
| 5.1.1 | DHCP Performance Test | 233 |
| 5.1.2 | DHCP Utilization Test | 235 |
| MONITORING THE WINDOWS INTERNET NAME SERVICE (WINS) | | 237 |
| 6.1 | The WINS Server Layer | 238 |
| 6.1.1 | Wins Test | 238 |
| MONITORING MS PRINT SERVERS | | 240 |
| 7.1 | The MS Print Service Layer | 240 |
| 7.1.1 | Print Server Test | 241 |
| MONITORING MS PROXY SERVERS | | 243 |
| 8.1 | The Proxy Service Layer | 244 |
| 8.1.1 | Win Sock Test | 244 |
| 8.1.2 | Proxy Server Test | 246 |
| 8.1.3 | Proxy Cache Test | 248 |
| 8.1.4 | Proxy Svc Test | 249 |
| MONITORING WINDOWS DOMAIN CONTROLLERS | | 253 |
| 9.1 | The Windows Server Layer | 253 |
| 9.1.1 | Windows Access Test | 254 |
| 9.1.2 | Windows Sessions Test | 255 |
| 9.1.3 | Window Authentication Test | 256 |

| | |
|---|------------|
| MONITORING MS FILE SERVERS | 259 |
| 10.1 The Windows Server Layer | 260 |
| 10.1.1 Windows Access Test | 260 |
| 10.1.2 Windows Sessions Test..... | 261 |
| 10.2 The File Server Layer..... | 263 |
| 10.2.1 MS File Stats Test | 263 |
| 10.2.2 Windows Usage Test..... | 264 |
| MONITORING ISA PROXY SERVERS | 266 |
| 11.1 The Firewall Service Layer..... | 267 |
| 11.1.1 ISA Cache Test..... | 267 |
| 11.1.2 ISA Firewall Test | 268 |
| 11.1.3 ISA Web Proxy Test | 269 |
| 11.1.4 Packet Engine Test | 270 |
| 11.1.5 Proxy Server Test | 271 |
| 11.1.6 Tests that are Disabled by Default | 273 |
| MONITORING MICROSOFT RADIUS SERVERS..... | 278 |
| 12.1 The MS Radius Layer..... | 279 |
| 12.1.1 IAS Acc Server Test..... | 280 |
| 12.1.2 IAS Acc Client Test | 282 |
| 12.1.3 IAS Auth Server Test | 283 |
| 12.1.4 IAS Auth Client Test..... | 285 |
| MONITORING THE MICROSOFT RAS SERVER..... | 288 |
| 13.1 The MS RAS Service Layer | 289 |
| 13.1.1 Microsoft RAS Port Test | 289 |
| 13.1.2 Microsoft RAS Test | 291 |
| 13.1.3 Windows Telephony Test | 292 |
| MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS) | 295 |
| 14.1 The SMS Site Server Layer | 296 |
| 14.1.1 Data Discovery Test | 296 |
| 14.1.2 Inv Load Test..... | 297 |
| 14.1.3 Memory Queue Test..... | 298 |
| 14.1.4 SMS Status Messages Test | 299 |
| 14.1.5 SMS Threads Test | 300 |
| 14.1.6 Software Inventory Proc Test | 301 |
| 14.1.7 Software Metering Test..... | 302 |
| 14.2 The SMS Mgmt Point Layer..... | 304 |
| 14.2.1 Management Point Data Loader Test..... | 304 |

| | | |
|---|--|------------|
| 14.2.2 | MgmtPointHwInv Test | 305 |
| 14.2.3 | Management Point Policy Manager Test | 306 |
| 14.2.4 | Management Point Policy Test..... | 306 |
| 14.2.5 | Management Point Status Manager Test | 307 |
| 14.2.6 | Management Point Software Inventory Test | 308 |
| EXTERNALLY MONITORING THE ACTIVE DIRECTORY SERVER..... | | 309 |
| 15.1 | The Network Layer..... | 309 |
| 15.2 | The Application Processes Layer | 310 |
| 15.3 | The DC Server Layer | 310 |
| MONITORING THE AD CLUSTER SERVICE | | 312 |
| 16.1 | The DC Server Layer | 313 |
| MONITORING WINDOWS CLUSTERS | | 314 |
| 17.1 | The Windows Service Layer..... | 315 |
| MONITORING MICROSOFT SHAREPOINT | | 322 |
| 18.1 | Monitoring Sharepoint 2007..... | 322 |
| 18.1.1 | The Sharepoint Services Layer..... | 324 |
| 18.2 | Monitoring Sharepoint 2010/2013 | 341 |
| 18.2.1 | The Sharepoint Documents Layer..... | 342 |
| 18.2.2 | The Sharepoint Objects Layer | 347 |
| 18.2.3 | Sharepoint Web Applications Test..... | 363 |
| 18.2.4 | The Sharepoint Services Layer..... | 368 |
| 18.2.5 | Sharepoint Search Content Feed Layer..... | 388 |
| MONITORING MICROSOFT DYNAMICS AX..... | | 399 |
| 19.1 | Dynamics AOS Service..... | 400 |
| 19.1.1 | AX Object Statistics Test..... | 400 |
| 19.1.2 | AX Portal Statistics Test..... | 402 |
| MONITORING THE MICROSOFT RDS LICENSE SERVER | | 404 |
| 20.1 | RD License Manager Layer | 405 |
| 20.1.1 | TS CAL Licenses Utilization Test | 405 |
| CONCLUSION | | 411 |

Table of Figures

| | |
|---|-----|
| Figure 2.1: Layer model of a Microsoft RDS server | 2 |
| Figure 2.1: The tests mapped to the Windows Service layer | 4 |
| Figure 2.2: Tests associated with the Terminal Server layer | 19 |
| Figure 2.3: The detailed diagnosis of the <i>Total GDI objects</i> measure | 34 |
| Figure 2.4: Tests associated with the Terminal Applications layer | 34 |
| Figure 2.5: The detailed diagnosis of the Processes running measure | 37 |
| Figure 2.6: Tests associated with the Terminal Users layer | 44 |
| Figure 2.7: The detailed diagnosis of the Active sessions measure | 47 |
| Figure 2.8: The detailed diagnosis of the Sessions logging out measure | 49 |
| Figure 2.9: The detailed diagnosis of the User sessions measure | 56 |
| Figure 2.10: The detailed diagnosis of the New disconnects measure | 58 |
| Figure 2.11: The detailed diagnosis of the Quick reconnects measure | 58 |
| Figure 3.1: Layer model for Active Directory | 69 |
| Figure 3.2: The tests associated with the AD Server layer | 72 |
| Figure 3.3: The tests mapped to the DNS/DHCP layer | 95 |
| Figure 3.4: The tests mapped to the AD Replication Service layer | 108 |
| Figure 3.5: Tests mapping to the DC Service layer | 134 |
| Figure 3.6: The details of orphaned objects | 135 |
| Figure 4.1: Layer model of a BizTalk server | 169 |
| Figure 4.2: Tests mapping to the BTS Transport layer | 170 |
| Figure 4.3: Tests mapping to the BTS Documents layer | 173 |
| Figure 4.4: The major components of a BizTalk server | 176 |
| Figure 4.5: The layer model of the BizTalk Server 2010 | 177 |
| Figure 4.6: Messaging architecture | 179 |
| Figure 4.7: The tests mapped to the Messaging Engine layer | 180 |
| Figure 4.8: The tests mapped to the Message Box layer | 218 |
| Figure 4.9: The tests mapped to the Orchestration Engine layer | 223 |
| Figure 4.10: How does BAM work? | 228 |
| Figure 5.1: Layer model of a DHCP server | 233 |
| Figure 5.2: Tests associated with the DHCP Services layer | 233 |
| Figure 6.1: Layer model of a WINS server | 238 |
| Figure 6.2: Test associated with the WINS server layer | 238 |
| Figure 7.1: Layer model of an MS Print server | 240 |
| Figure 7.2: Tests associated with the MS Print Service layer | 241 |
| Figure 8.1: Layer model of an MS Proxy server | 243 |
| Figure 8.2: Tests associated with the Proxy Service layer | 244 |
| Figure 9.1: Layer model of a Windows Domain Controller | 253 |
| Figure 9.2: Tests associated with the Windows Server layer | 254 |
| Figure 10.1: Layer model of an MS File server | 259 |
| Figure 10.2: Tests associated with the Windows Server layer | 260 |
| Figure 10.3: Tests associated with the File server layer | 263 |
| Figure 11.1: Layer model of an ISA Proxy server | 266 |
| Figure 11.2: The tests associated with the Firewall Service layer | 267 |
| Figure 12.1: The layer model of the MS Radius server | 279 |
| Figure 12.2: The tests associated with the MS Radius layer | 279 |
| Figure 13.1: Layer model of the MS RAS server | 288 |
| Figure 13.2: The tests associated with the MSRAS_SERVICE layer | 289 |
| Figure 14.1: The layer model of Microsoft SMS | 295 |
| Figure 14.2: The tests associated with the SMS Site Server layer | 296 |
| Figure 14.3: The tests associated with the SMS Mgmt Point layer | 304 |
| Figure 15.1: Layer model of the External AD server | 309 |
| Figure 15.2: The test associated with the Network layer | 310 |
| Figure 15.3: The tests associated with the Application Processes layer | 310 |
| Figure 15.4: The tests associated with the DC Server layer | 311 |
| Figure 16.1: Layer model of the AD cluster service | 312 |
| Figure 16.2: The tests associated with the DC_SERVER layer | 313 |
| Figure 17.1: The layer model of the IIS web server with the Windows Cluster layer | 314 |
| Figure 17.2: The tests mapped to the Windows Service layer | 315 |
| Figure 18.1: The layer model of Sharepoint | 323 |
| Figure 18.2: The tests mapped to the Sharepoint Services layer | 324 |
| Figure 18.3: Excel services architecture | 328 |
| Figure 18.4: The layer model of Microsoft Sharepoint 2010 | 341 |

| | |
|--|-----|
| Figure 19.1: The tests mapped to the Sharepoint Documents Layer | 342 |
| Figure 19.2: The detailed diagnosis of the Number of document libraries measure..... | 345 |
| Figure 19.3: The detailed diagnosis of the Lists count measure | 345 |
| Figure 19.4: The tests mapped to the Sharepoint Objects layer..... | 348 |
| Figure 18.5: Site Collections and Sites | 356 |
| Figure 19.5: The detailed diagnosis of the Least active site collections measure | 362 |
| Figure 19.6: The detailed diagnosis of the Least active sites measure..... | 363 |
| Figure 18.6: The tests mapped to the Sharepoint Services layer | 368 |
| Figure 18.7: How Search works in Sharepoint 2010? | 371 |
| Figure 19.7: How search works in Sharepoint 2013? | 389 |
| Figure 19.8: Flows and operators in CPC | 396 |
| Figure 19.9: The layer model of the Microsoft Dynamics AX solution | 399 |
| Figure 19.10: The tests mapped to the Dynamics AOS Service | 400 |
| Figure 20.1: Layer model of the Microsoft RDS License server..... | 404 |
| Figure 20.2: The tests mapped to the TS CAL Licenses Utilization test..... | 405 |
| Figure 20.3: The detailed diagnosis of the CAL type measure..... | 410 |
| Figure 20.4: The detailed diagnosis of the Licenses in use measure..... | 410 |

Introduction

Microsoft applications are common-place in IT infrastructures today. From web interfaces to domain controllers to authentication servers to Microsoft RDS servers to simple browsers, a wide range of Windows-based applications are being increasingly utilized by infrastructure operators to keep the IT environment afloat and easily accessible to end-users.

This means that even a slight slowdown in the performance of one of these applications, if not resolved soon, can prove to be fatal to the critical end-user service riding on it. This is reason enough for bringing Microsoft applications under the purview of '24x7 monitoring'.

eG Enterprise provides 100% web-based monitoring models to continuously monitor and report on the status of critical Microsoft applications such as Active Directory servers, Microsoft RDS servers, Windows Domain Controllers, etc.

This document describes the monitoring model that eG Enterprise prescribes for every Microsoft application, and the performance metrics each model collects.

Monitoring Microsoft RDS Servers

The Microsoft RDS Server is a server program that provides the graphical user interface (GUI) of the Windows desktop to user terminals that don't have this capability themselves. The latter include the relatively low-cost NetPC or "thin client" that some companies are purchasing as alternatives to the autonomous and more expensive PC with its own operating system and applications.

Typically, Microsoft RDS server environments involve multiple tiers of software. Domain servers in the target infrastructure handle authentication of users. Authenticated requests are passed to the Microsoft RDS servers that host a number of applications. In turn, the applications may use backend databases, printers, etc., for different functionalities. Owing to the multi-tier nature of Microsoft RDS server environments, a slow-down in one tier (e.g., the authentication server) can cause a slow-down of the entire service. When a slow-down occurs, an administrator of the server farm has to quickly determine what the source of the problem could be - i.e., Is it the network? Or the authentication server? Or the Microsoft RDS server? Or the backend database? Or the application? Accurate, fast diagnosis of problems helps reduce downtime and improve customer satisfaction.

The eG Enterprise suite offers 100% web-based monitoring of Microsoft RDS server farms. The suite includes an extensive, pre-defined, customized *Microsoft RDS* model for this server (see Figure 2.1), which defines the key performance metrics that need to be tracked to determine the service level achieved by the server/server farm.

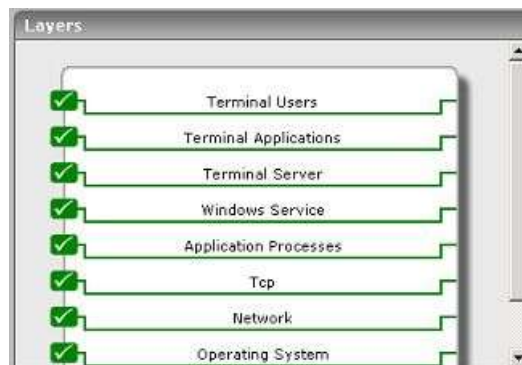


Figure 2.1: Layer model of a Microsoft RDS server

Using the metrics reported by each of the layers depicted by Figure 2.1, administrators can find answers to persistent

MONITORING MICROSOFT RDS SERVERS

performance-related queries discussed hereunder:

| | |
|---|---|
| Microsoft RDS server Monitoring | <p>Are the Microsoft RDS servers available to service user requests?</p> <p>Are there sporadic disconnects from the Microsoft RDS server?</p> <p>At what times do peak usage of the servers happen and is the server capacity adequate?</p> <p>Is the user load being balanced across all the servers?</p> <p>Is the data store available?</p> |
| User Monitoring | <p>What is the average response time that a user sees when connecting to a Microsoft RDS server?</p> <p>How many users are logged in to each server in the Microsoft RDS server farm?</p> <p>What is the resource usage (CPU and memory) for each user?</p> <p>What is the I/O activity generated by every user?</p> <p>How much network bandwidth is consumed by every user?</p> <p>Are too many page faults occurring in the processes executed on a server?</p> <p>If so, what are those processes, and who are the users executing them?</p> <p>Which user is using a lot of handles?</p> |
| Operating System Monitoring | <p>What is the average CPU and memory usage on all the servers in the farm?</p> <p>Is any unusual memory scanning/paging activity happening on the systems?</p> <p>Are the critical Microsoft RDS server processes up?</p> <p>What is their resource consumption?</p> |
| Hosted Application Monitoring | <p>What are the applications hosted on a Microsoft RDS server?</p> <p>Who is using each application?</p> <p>What is the resource usage for each published application?</p> |
| Infrastructure Services Monitoring | <p>Are the backend databases working?</p> <p>What is the resource usage of the databases?</p> <p>Are users able to login to the server farm? How long is the login process taking?</p> <p>What is the usage of the Microsoft Windows Domain Controller?</p> |

Since the 4 layers at the bottom of Figure 2.1 have already been discussed in the *Monitoring Unix and Windows Servers* document, the sections to come will discuss the top 4 layers only.

2.1 The Windows Service Layer

This layer represents the different services of the corresponding Windows components in the environment. An eG agent uses **Windows Services** test to track the health of this layer. In addition, the layer also periodically monitors the application, security, and system-related events that occur on the target Windows host. Since most of the tests

MONITORING MICROSOFT RDS SERVERS

of this layer have already been dealt in the Monitoring Unix and Windows servers document, let us now discuss the tests that are exclusive for the Microsoft RDS Servers alone.



Figure 2.1: The tests mapped to the Windows Service layer

2.1.1 App-V Client Admin Log Test

This test reports the statistical information about the admin events generated by the target system.



Note

This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

| | |
|---------------------------------|---|
| Purpose | Reports the statistical information about the admin events generated by the target system |
| Target of the test | An App-V Client on the target Microsoft RDS Server |
| Agent deploying the test | An internal agent |

| | |
|---|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Specify the port at which the specified HOST listens to. By default, this is 8080. 4. LOGTYPE – Refers to the type of event logs to be monitored. The default value is <i>Microsoft-AppV-Client/Admin</i>. 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: <ul style="list-style-type: none"> ➤ 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 <p>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.</p> 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: <i>{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}</i>. For example, assume that the FILTER text area takes the value, <i>OS_events:all:Browse,Print:all:none:all:none</i>. Here: <ul style="list-style-type: none"> ➤ <i>OS_events</i> is the display name that will appear as a descriptor of the test in the monitor UI; ➤ <i>all</i> 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 <i>none</i>. ➤ Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, <i>Browse</i> and <i>Print</i> have been excluded from monitoring. Alternatively, you can use <i>all</i> to indicate that all the event sources have to be excluded from monitoring, or <i>none</i> 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 <i>all</i> in our example represents that all the event IDs need to be considered while monitoring. ➤ Similarly, the <i>none</i> (following <i>all</i> 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 <i>all</i> 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*. Multiple filters are to be separated by semi-colons (;).



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  icon appears near the **FILTER** list box, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the  icon 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.

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 NT/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**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the USEWMI flag should always be set to 'Yes'**.
8. **STATELESS ALERTS** - Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **CRITICAL** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **CRITICAL**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the **STATELESS ALERTS** flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
9. **DDFORINFORMATION** – eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the **DDFORINFORMATION** and **DDFORWARNING** flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the **DDFORINFORMATION** flag to **No**.
10. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.

| | | | |
|--------------------------------------|--|-------------------------|---|
| | <p>11. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DDFREQ.</p> <p>12. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>1. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the App-V Client that is to be monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Information messages: Indicates the number of App-V Client admin information events generated when the test was last executed.</p> | Number | <p>A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications.</p> <p>Please check the App-V Client admin logs in the Event Log Viewer for more details.</p> |
| | <p>Warnings: Indicates the number of App-V Client admin warnings that were generated when the test was last executed.</p> | Number | <p>A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications.</p> <p>Please check the App-V Client admin logs in the Event Log Viewer for more details.</p> |
| | <p>Error messages: Indicates the number of App-V Client admin error events that were generated during the last measurement period.</p> | Number | <p>A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems.</p> <p>An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications.</p> <p>Please check the App-V Client admin logs in the Event Log Viewer for more details.</p> |

| | | | |
|--|--|---------------|--|
| | <p>Critical messages: Indicates the number of App-V Client admin critical error events that were generated when the test was last executed.</p> | <p>Number</p> | <p>A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems.</p> <p>An increasing trend or high value indicates the existence of fatal/irreparable problems in one or more applications.</p> <p>Please check the App-V Client admin logs in the Event Log Viewer for more details.</p> |
| | <p>Verbose messages: Indicates the number of App-V Client admin verbose events that were generated when the test was last executed.</p> | <p>Number</p> | <p>The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.</p> <p>Please check the App-V Client admin logs in the Event Log Viewer for more details.</p> |

2.1.2 App-V Client Operational Log Test

This test reports the statistical information about the operation events generated by the target system.



This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

| | |
|---------------------------------|---|
| Purpose | Reports the statistical information about the operation events generated by the target system |
| Target of the test | An App-V Client on the target Microsoft RDS Server |
| Agent deploying the test | An internal agent |

| | |
|--|---|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 2. TEST PERIOD - How often should the test be executed 3. HOST - The host for which the test is to be configured 4. PORT – Specify the port at which the specified HOST listens to. By default, this is 8080. 5. LOGTYPE – Refers to the type of event logs to be monitored. The default value is <i>Microsoft-AppV-Client/Operational</i>. 6. 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: <ul style="list-style-type: none"> ➤ 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 <p>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.</p> 7. 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: <i>{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}</i>. For example, assume that the FILTER text area takes the value, <i>OS_events:all:Browse,Print:all:none:all:none</i>. Here: <ul style="list-style-type: none"> ➤ <i>OS_events</i> is the display name that will appear as a descriptor of the test in the monitor UI; ➤ <i>all</i> 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 <i>none</i>. ➤ Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, <i>Browse</i> and <i>Print</i> have been excluded from monitoring. Alternatively, you can use <i>all</i> to indicate that all the event sources have to be excluded from monitoring, or <i>none</i> 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 <i>all</i> in our example represents that all the event IDs need to be considered while monitoring. ➤ Similarly, the <i>none</i> (following <i>all</i> 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 <i>all</i> 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*. Multiple filters are to be separated by semi-colons (;).



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  icon appears near the **FILTER** list box, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the  icon 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.

8. **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 NT/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**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the USEWMI flag should always be set to 'Yes'**.
9. **STATELESS ALERTS** - Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **CRITICAL** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **CRITICAL**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the **STATELESS ALERTS** flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
10. **DDFORINFORMATION** – eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the **DDFORINFORMATION** and **DDFORWARNING** flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the **DDFORINFORMATION** flag to **No**.
11. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.
12. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DDREQ**.
13. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
 - o The eG manager license should allow the detailed diagnosis capability
 - o Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|---|--|
| Outputs of the test | One set of results for the App-V Client that is to be monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Information messages: Indicates the number of App-V Client operational information events generated when the test was last executed. | Number | A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details. |
| | Warnings: Indicates the number of App-V Client operational warnings that were generated when the test was last executed. | Number | A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details. |
| | Error messages: Indicates the number of App-V Client operational error events that were generated during the last measurement period. | Number | A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details. |
| | Critical messages: Indicates the number of App-V Client operational critical error events that were generated when the test was last executed. | Number | A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irreparable problems in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details. |
| Verbose messages: Indicates the number of App-V Client operational verbose events that were generated when the test was last executed. | Number | The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App-V Client Operational logs in the Event Log Viewer for more details. | |

2.1.3 App-V Client Virtual Application Log Test

This test reports the statistical information about the virtual application events generated by the target system.



This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

| | |
|---------------------------------|---|
| Purpose | Reports the statistical information about the virtual application events generated by the target system |
| Target of the test | An App-V Client on the target Microsoft RDS Server |
| Agent deploying the test | An internal agent |

| | |
|--|--|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Specify the port at which the specified HOST listens to. By default, this is 8080. 4. LOGTYPE – Refers to the type of event logs to be monitored. The default value is <i>Microsoft-AppV-Client/Virtual Applications</i>. 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: <ul style="list-style-type: none"> ➤ 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 <p>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.</p> 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: <i>{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}</i>. For example, assume that the FILTER text area takes the value, <i>OS_events:all:Browse,Print:all:none:all:none</i>. Here: <ul style="list-style-type: none"> ➤ <i>OS_events</i> is the display name that will appear as a descriptor of the test in the monitor UI; ➤ <i>all</i> 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 <i>none</i>. ➤ Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, <i>Browse</i> and <i>Print</i> have been excluded from monitoring. Alternatively, you can use <i>all</i> to indicate that all the event sources have to be excluded from monitoring, or <i>none</i> 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 <i>all</i> in our example represents that all the event IDs need to be considered while monitoring. ➤ Similarly, the <i>none</i> (following <i>all</i> 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 <i>all</i> 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*. Multiple filters are to be separated by semi-colons (;).



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  icon appears near the **FILTER** list box, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the  icon 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.

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 NT/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**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the USEWMI flag should always be set to 'Yes'.**
8. **STATELESS ALERTS** - Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **CRITICAL** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **CRITICAL**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the **STATELESS ALERTS** flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
9. **DDFORINFORMATION** – eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the **DDFORINFORMATION** and **DDFORWARNING** flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the **DDFORINFORMATION** flag to **No**.
10. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.
11. **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**.
12. **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.
 1. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:
 - o The eG manager license should allow the detailed diagnosis capability
 - o 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 App-V Client that is to be monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Information messages: Indicates the number of App-V Client virtual application informational events that were generated when the test was last executed.</p> | Number | <p>A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.</p> |
| | <p>Warnings: Indicates the number of App-V Client virtual application warnings that were generated when the test was last executed.</p> | Number | <p>A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.</p> |
| | <p>Error messages: Indicates the number of App-V Client virtual application error events that were generated during the last measurement period.</p> | Number | <p>A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.</p> |
| | <p>Critical messages: Indicates the number of App-V Client virtual applications critical error events that were generated when the test was last executed.</p> | Number | <p>A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irreparable problems in one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.</p> |

| | | | |
|--|---|--------|--|
| | <p>Verbose messages:</p> <p>Indicates the number of App-V Client virtual application verbose events that were generated when the test was last executed.</p> | Number | <p>The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.</p> <p>Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.</p> |
|--|---|--------|--|

2.2 The Terminal Server Layer

The tests associated with this layer (see Figure 2.2) enable administrators to measure the health of the client to server connectivity, using metrics such as the following:

- The availability of the Microsoft RDS server and its responsiveness to client requests
- Login time to the server
- The status of file serving as seen by a Microsoft RDS client




Figure 2.2: Tests associated with the Terminal Server layer

2.2.1 Session Login Status Test

Administrators typically use the *Change logon* command line tool to enable / disable logons from client sessions to the Citrix / Microsoft RDS server. Disabling client logons will deny all users access to the server. Whenever users complaint of login failures, administrators might first want to check the status of the client logons to determine whether it has been disabled or not. This test periodically reports the status of logons from client sessions to the Citrix / Microsoft RDS server.

| | |
|---------------------------|---|
| Purpose | Periodically reports the status of logons from client sessions to the Citrix / Microsoft RDS server |
| Target of the test | A Microsoft RDS server |
| Agent | An internal agent |

| | | | |
|--------------------------------------|--|-------------------------|--|
| deploying the test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 2. TEST PERIOD - How often should the test be executed 3. HOST - Host name of the server for which the test is to be configured 4. PORT - Enter the port to which the HOST listens | | |
| Outputs of the test | One set of results the Microsoft RDS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Session login status: Indicates whether the client sessions to the server are currently enabled or not. | Percent | If the value for this measure is 100, it indicates all client logons are enabled. If the value of this measure is 0, it indicates that client logons are disabled. |

2.2.2 Terminal Connection Test

This test tracks various statistics pertaining to Microsoft RDS server connections to and from a host, from an external perspective.

| | | | |
|--------------------------------------|--|-------------------------|-----------------------|
| Purpose | Tracks various statistics pertaining to Microsoft RDS server connections to and from a host, from an external perspective | | |
| Target of the test | A Microsoft RDS server | | |
| Agent deploying the test | An external agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - Host name of the server for which the test is to be configured 3. PORT - Enter the port to which the specified TARGETHOST listens 4. TARGETPORTS – Specify a comma-separated list of port numbers that are to be tested (eg., 80,7077,1521). By default, the default terminal sever port, 3389, will be displayed here. | | |
| Outputs of the test | One set of results for every port being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | | | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|---------|--|
| | Connection availability: Whether the Microsoft RDS server connection is available | 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 DNS server. |
| | Connection 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 DNS server, a network problem, etc. |

2.2.3 Terminal Authentication Test

This test emulates the user login process at the system level on a Microsoft RDS server and reports whether the login succeeded and how long it took.

| | |
|---------------------------------|--|
| Purpose | Emulates the user login process at the system level on a Microsoft RDS server and reports whether the login succeeded and how long it took |
| Target of the test | A Microsoft RDS server |
| Agent deploying the test | An internal agent |

| | | | |
|--|---|--|--|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. USERNAME - This test emulates the user login process at the system level on a Microsoft RDS server. Therefore, specify the login name of a user with both interactive logon and logon locally privileges. 5. PASSWORD - Enter the password that corresponds to the specified USERNAME. 6. CONFIRM PASSWORD – Confirm the password by retyping it here. 7. DOMAIN - Specify the name of the domain to which the test will try to login. If the test is to login to a local host, specify 'none' here. <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p>Note:</p> <p>If users are spread across multiple domains, then, you can configure this test with multiple DOMAIN specifications; in this case, for every DOMAIN, a USER-PASSWORD pair might also have to be configured. Sometimes, you might want the test to login as specific users from the same domain, to check how long each user login takes. Both these scenarios require the configuration of multiple DOMAINS and/or multiple USER names and PASSWORDS. In order to enable users to specify these details with ease, eG Enterprise provides a special page; to access this page, click on the Click here hyperlink at the top of the parameters in the test configuration page. To know how to use this page, refer to the Configuring Multiple Users for the Citrix Authentication Test section in the <i>Monitoring Citrix Environments</i> document.</p> </div> <ol style="list-style-type: none"> 8. REPORT BY DOMAIN - By default, this flag is set to Yes. This implies that by default, this test will report metrics for every <i>domainname\username</i> configured for this test. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the <i>username</i> alone, then set this flag to No. | | |
| <p>Outputs of the test</p> | <p>One set of results for every user account being checked</p> | | |
| <p>Measurements made by the test</p> | <p style="text-align: center;">Measurement</p> | <p style="text-align: center;">Measurement Unit</p> | <p style="text-align: center;">Interpretation</p> |
| | <p>Authentication status: Indicates whether the login was successful or not</p> | <p>Percent</p> | <p>A value of 100 % indicates that the login has succeeded. The value 0 is indicative of a failed login.</p> |
| | <p>Authentication time: Indicates the time it took to login</p> | <p>Secs</p> | <p>If this value is very high then it could be owing to a configuration issue (i.e. the domain might not be configured properly) or a slow-down/unavailability of the primary domain server.</p> |

2.2.4 Redirector Test

File serving very often is a much underestimated part of Citrix and Microsoft RDS server environments. Improperly configured file serving components can wreak havoc on a server farm’s performance.

File serving in Citrix and Microsoft RDS server environments is used at different times. For instance, every time a user logs on or off, profile data may be copied back and forth between the file server and terminal or Citrix server. Another example involves multiple applications accessing configurations stored in files from a remote file server. Folder redirection, if used, is another form of file retrievals from file servers.

File serving problems can have a detrimental impact on the performance of Citrix/Microsoft RDS server environments. Often, these problems may manifest in many ways. For example, users may see very slow access to their home directory, or folders. Even with a small profile, logging on and off could take a long time. Random application crashes can also happen, especially for applications that rely on file servers to store their configuration files remotely. Such file serving problems are often the most difficult to diagnose.

The Redirector component of the Microsoft Windows operating system handles file serving at the client end, and the Redirector test monitors this component’s activity, and tracks the status of file serving as seen by a file server’s client (i.e., the Citrix or Microsoft RDS server).

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the activity of redirector component of the Microsoft windows operating system and tracks the status of the file serving as seen by a file server’s client. | | |
| Target of the test | Any Microsoft RDS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server | | |
| Outputs of the test | One set of results for the Microsoft RDS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Data received: This metric shows the rate of data that were received by the local server from the network. This includes all the application data as well as network protocol information. | MB/Sec | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|------------|---|
| | <p>Data sent:</p> <p>This metric represents the rate at which data is leaving the Redirector to the network. This includes all the application data as well as network protocol information.</p> | MB/sec | |
| | <p>Current commands:</p> <p>This metric indicates the number of requests to the Redirector that are currently queued for service.</p> | Number | <p>The Current Commands measure indicates the number of pending commands from the local computer to all destination servers. This means that if one of the destination servers does not respond in a timely manner, the number of current commands on the local computer may increase.</p> <p>If the local computer is serving many sessions, a high number of current commands does not necessarily indicate a problem or a bottleneck. However, if the Current Commands measure shows a high number and the local computer is idle, this may indicate a network-related problem or a redirector bottleneck on the local computer. For example, there may be a network-related problem or a local bottleneck if the computer is idle overnight but the counter shows a high number during that period.</p> |
| | <p>Network errors:</p> <p>This metric denotes the rate at which serious unexpected errors are occurring during file system access from a remote server.</p> | Errors/sec | <p>Such errors generally indicate that the Redirector and one or more Servers are having serious communication difficulties. For example an SMB (Server Manager Block) protocol error is a Network Error. An entry is written to the System Event Log and provides details.</p> |
| | <p>Reads denied :</p> <p>This metric denotes the rate at which the server is unable to accommodate requests for raw read operations.</p> | Reads/sec | <p>When a read is much larger than the server's negotiated buffer size, the Redirector requests a Raw Read which, if granted, would permit the transfer of the data without lots of protocol overhead on each packet. To accomplish this, the server must lock out other requests, so the request is denied if the server is really busy.</p> |

| | | | |
|--|--|------------|--|
| | <p>Hung server sessions:</p> <p>This metric shows the number of active sessions that are timed out and unable to proceed due to a lack of response from the remote file server.</p> | Number | |
| | <p>Writes denied:</p> <p>This metric denotes the rate at which the server is unable to accommodate requests for raw write operations</p> | Writes/sec | When a write is much larger than the server's negotiated buffer size, the Redirector requests a Raw Write which, if granted, would permit the transfer of the data without lots of protocol overhead on each packet. To accomplish this, the server must lock out other requests, so the request is denied if the server is really busy. |

2.2.5 User Profile Test

User profiles are the heart of the Microsoft RDS server environment. User profiles contain the configuration settings, which bring the user desktop alive. One of the major problems in a server-based computing environment like the Microsoft RDS server is that the user's login process takes more time to open the user's desktop. This happens if the user profile size is huge. The UserProfile test monitors the size of the Microsoft RDS server user profiles and raises an alarm if the profile size exceeds the profile quota size.

| | |
|---------------------------------|--|
| Purpose | Monitors the size of the Microsoft RDS server user profiles and raises an alarm if the profile size exceeds the profile quota size |
| Target of the test | Any Microsoft RDS server |
| Agent deploying the test | An internal agent |

| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. PROFILESIZELIMIT - Specify the profile quota size (in MB). The default value is 50 MB. 5. EXCLUDE - Provide a comma-separated list of users who need to be excluded from the analysis. By default, this parameter is set to <i>All_Users</i>, indicating that, by default, the test will not monitor the <i>All_Users</i> profile. 6. CURRENTUSERONLY - If this is set to true, then the profile sizes of only those users who are currently logged into the server will be monitored. If this is set to false, eG Enterprise will perform profile monitoring for all the users to the server. 7. FILESIZELIMIT - Takes the file quota size (in KB). The default size is 10000 KB. 8. REPORT BY DOMAIN - By default, this flag is set to Yes. This implies that by default, this test will report metrics for every <i>domainname\username</i> to the server. This way, administrators will be able to quickly determine which user belongs to which domain. If you want the test to report metrics for every <i>username</i> alone, then set this flag to No. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | | | | | | | | | | |
|---|---|--|------------------|----------------|---|----------------|--|--|-----------|--|--|--|
| <p>Outputs of the test</p> | <p>One set of results for every user profile on the Microsoft RDS server monitored</p> | | | | | | | | | | | |
| <p>Measurements made by the test</p> | <table border="1"> <thead> <tr> <th data-bbox="386 1329 719 1392">Measurement</th> <th data-bbox="727 1329 919 1392">Measurement Unit</th> <th data-bbox="919 1329 1416 1392">Interpretation</th> </tr> </thead> <tbody> <tr> <td data-bbox="386 1392 719 1717"> <p>Is user profile exceeding quota?: Indicates whether the profile size exceeds the profile quota size by comparing the current profile size with the configured PROFILESIZELIMIT parameter.</p> </td> <td data-bbox="727 1392 919 1717"> <p>Boolean</p> </td> <td data-bbox="919 1392 1416 1717"> <p>If this measure shows 0, it indicates that the current profile size has not exceeded the quota size. The value 1 indicates that the current profile size has exceeded the quota size.</p> </td> </tr> <tr> <td data-bbox="386 1717 719 1831"> <p>Current profile size: Indicates the current profile size.</p> </td> <td data-bbox="727 1717 919 1831"> <p>MB</p> </td> <td data-bbox="919 1717 1416 1831"></td> </tr> </tbody> </table> | Measurement | Measurement Unit | Interpretation | <p>Is user profile exceeding quota?: Indicates whether the profile size exceeds the profile quota size by comparing the current profile size with the configured PROFILESIZELIMIT parameter.</p> | <p>Boolean</p> | <p>If this measure shows 0, it indicates that the current profile size has not exceeded the quota size. The value 1 indicates that the current profile size has exceeded the quota size.</p> | <p>Current profile size: Indicates the current profile size.</p> | <p>MB</p> | | | |
| Measurement | Measurement Unit | Interpretation | | | | | | | | | | |
| <p>Is user profile exceeding quota?: Indicates whether the profile size exceeds the profile quota size by comparing the current profile size with the configured PROFILESIZELIMIT parameter.</p> | <p>Boolean</p> | <p>If this measure shows 0, it indicates that the current profile size has not exceeded the quota size. The value 1 indicates that the current profile size has exceeded the quota size.</p> | | | | | | | | | | |
| <p>Current profile size: Indicates the current profile size.</p> | <p>MB</p> | | | | | | | | | | | |

| | | | |
|--|---|--------|--|
| | <p>Number of files in user's profile:</p> <p>Indicates the number of files available in the user profile.</p> | Number | |
| | <p>Large files in user's profile:</p> <p>The number of files in the user profile, which exceed the allowable FILESIZELIMIT parameter.</p> | Number | The detailed diagnosis of this measure, if enabled, lists all the files that have exceeded the configured FILESIZELIMIT . |

2.2.6 User Environment Test

The process of a user logging into a Citrix or Microsoft RDS server is fairly complex. First, the profile corresponding to a user has to be located, and the appropriate profile files copied over from a profile server (in the case of a roaming profile). Second, additional processing is often necessary after copying the profile locally. Processing for instance may involve creating new printers for the user logging in. Proper monitoring of profile loading and processing times is key because the login process is handled exclusively by Microsoft Windows. Hence, if a specific user profile takes a lot of time to load (e.g., because the profile is very big), or if specific processing for a user is taking time, this could delay logins for subsequent users who are trying to access the server at the same time. The typical process for monitoring the Windows login process is to use the user environment debugging mechanism. To enable this, the following steps are required. To set the logging level associated with the userenv.log file, perform the following steps:

- Start a registry editor (e.g., regedit.exe).
- Navigate to the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon registry subkey.
- From the Edit menu, select New, DWORD Value.
- Enter the name UserEnvDebugLevel, then press Enter.
- Double-click the new value, set it to 65538 (decimal) - which corresponds to the debugger output.

Once these changes are enabled, details about the Windows login process are logged into the file %systemroot%\debug\usermode\userenv.log. If the Userenv.log file is larger than 300 KB, the file is renamed Userenv.bak, and a new Userenv.log file is created. This action occurs when a user logs on locally or by using Terminal Services, and the Winlogon process starts. However, because the size check only occurs when a user logs on, the Userenv.log file may grow beyond the 300 KB limit. The 300 KB limit cannot be modified.

The UserEnvironment test periodically checks the userenv log file to monitor the user login and profile loading process. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Microsoft Terminal* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | |
|---------------------------|--|
| Purpose | Periodically checks the userenv log file to monitor the user login and profile loading process |
| Target of the test | Any Microsoft RDS server |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|---|---|-------------------------|--|
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Microsoft RDS server 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every Microsoft RDS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Profile load starts: Indicates the number of profile loads in the last measurement period. | Number | This metric gives an idea of the rate at which users are logging in to the server. |
| | Profile load successes: Indicates the number of successful profile loads in the last measurement period. | Number | |
| | Profile loading failures: Indicates the number of profile load failures in the last measurement period. | Number | An unusual increase in number of profile loading failures is a cause for concern. The userenv.log file will have details of what profile loads failed and why. |
| | Profile load failures percent: Indicates the percentage of profile loads that failed in the last measurement period. | Percent | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|--------|--|
| | <p>Avg user profile load time:</p> <p>Indicates the average time it took to load a profile successfully in the last measurement period.</p> | Secs | |
| | <p>Max profile load time:</p> <p>Indicates the maximum time it took to load a profile during the last measurement period.</p> | Secs | |
| | <p>System policy starts:</p> <p>Indicates the number of system policy applications started in the last measurement period.</p> | Number | |
| | <p>System policy completes:</p> <p>Indicates the number of system policy completions in the last measurement period.</p> | Number | Compare the total number of starts to completions. If there is a significant discrepancy, this denotes a bottleneck in system policy application. Check the userenv.log file for more details. |
| | <p>Avg system policy processing time:</p> <p>Indicates the average time taken for applying system policies in the last measurement period.</p> | Secs | If the system policy times are long, check the detailed diagnosis to view if the policy handling is taking time for all users. Analyze the userenv.log to determine the reason for any slowdown. |
| | <p>Max system policy time:</p> <p>Indicates the maximum time for applying system policies in the last measurement period.</p> | Secs | |
| | <p>Group policy starts:</p> <p>Indicates the number of group policy applications started in the last measurement period.</p> | Number | |
| | <p>Group policy completes:</p> <p>Indicates the number of group policy applications completed in the last measurement period.</p> | Number | |
| | <p>Avg group policy processing time:</p> <p>Indicates the average time taken for applying group policies.</p> | Secs | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|---------|--|
| | Max group policy time: Indicates the average time taken for applying group policies. | Secs | |
| | Profile unload starts: Indicates the number of profile unloads started during the last measurement period. | Number | |
| | Profile unload successes: Indicates the number of successful profile unloads during the last measurement period. | Number | |
| | Profile unload failures: Indicates the number of unsuccessful profile unloads during the last measurement period. | Number | |
| | Profile unload failures percent: Indicates the profile unload failures as a percentage of the total profile unloads. | Percent | |
| | Avg user profile unload time: Indicates the average time for unloading a profile during the last measurement period. | Secs | |
| | Max profile unload time: Indicates the maximum time for unloading a profile during the last measurement period. | Secs | |

2.2.7 Terminal Server CALs Test

This test reports the usage statistics pertaining to a Microsoft RDS server's client access licenses. To ensure that the test functions smoothly, the Terminal Services Licensing Reporter tool (**lsreport.exe**) needs to be available on the eG agent host. **lsreport.exe** is a command-line utility that you can use to display information about the licenses that are issued by Microsoft RDS License servers. **lsreport.exe** connects to Microsoft RDS License servers and logs information about the license key packs that are installed on the servers. In order to make sure that this utility is available to the eG Enterprise suite, do the following:

- Download the **lsreport.exe** from the Microsoft Windows 2000 Server Resource Kit.

MONITORING MICROSOFT RDS SERVERS

- Copy **Isreport.exe** to the **{EG_INSTALL_DIR}\bin** directory.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Microsoft Terminal* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | | | |
|---|---|-------------------------|--|
| Purpose | Reports the usage statistics pertaining to a Microsoft RDS server's client access licenses | | |
| Target of the test | A Microsoft RDS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the Microsoft RDS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active licenses: Represents number of active client access licenses that were currently consumed from the Microsoft RDS server license server. | Number | The detailed diagnosis of this provides the complete details of the active access licenses, which include critical session information such as the user who initiated the session, the start and end date/time of the session, the type of license issued to the user, the license ID, the issue type, the target server, the client from which the session was instantiated, etc. |

| | | | |
|--|---|--------|---|
| | <p>Temporary licenses:</p> <p>Indicates the number of temporary client access licenses that were currently consumed from the Microsoft RDS server license.</p> | Number | The detailed diagnosis of this provides the complete details of the temporary access licenses, which include critical session information such as the user who initiated the session, the start and end date/time of the session, the type of license issued to the user, the license ID, the issue type, the target server, the client from which the session was instantiated, etc. |
|--|---|--------|---|

2.2.8 GDI Objects Test

An object is a data structure that represents a system resource, such as a file, thread, or graphic image. An application cannot directly access object data or the system resource that an object represents. Instead, an application must obtain an object handle, which it can use to examine or modify the system resource. There are three categories of objects: user, GDI, and kernel. GDI objects support graphics. Here is a list of the GDI objects used in Windows:

- Bitmap
- Brush
- Device Context (DC)
- Enhanced Metafile
- Enhanced-metafile DC
- Font
- Memory DC
- Metafile
- Metafile DC
- Palette
- Pen/extended pen
- Region

GDI objects support only one handle per object, and only the process that created the object can use the object handle.

If an application creates a lot of these objects, without properly destroying references to the object (by closing the associated handle), then there will be multiple GDI objects occupying memory on the system for each object created. If this GDI leak is really bad, this can eventually bring a server to its knees, and cause all types of problems (slow logons, registry issues, system hangs, and so on).

MONITORING MICROSOFT RDS SERVERS

If such fatalities are to be avoided, administrators should closely monitor the number of GDI object handles created by every user to the Microsoft RDS server and proactively detect potential GDI leaks. This is where the **GDI Objects** test helps. This test periodically checks the GDI object handles created by each user to the Microsoft RDS server, reports the total number of handles created per user, and promptly notifies administrators if any user is creating more GDI handles than permitted. This way, the test brings probable GDI leaks to the attention of administrators. In addition, administrators can use the detailed diagnosis of the test to know which process is responsible for the GDI leak (if any).

| | | | |
|---|---|-------------------------|--|
| Purpose | Periodically checks the GDI object handles created by each user to the Microsoft RDS server, reports the total number of handles created per user, and promptly notifies administrators if any user is creating more GDI handles than permitted | | |
| Target of the test | A Microsoft RDS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Microsoft RDS server GDILIMIT – Specify the maximum number of GDI object handles that a user to the Microsoft RDS server can create. By default, this value is <i>10000</i>. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for each user to the Microsoft RDS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total GDI objects: Indicates the total number of GDI handles that this user has created. | Number | The detailed diagnosis of this measure, if enabled, provides the process-wise breakup of the GDI handles created by the user. In the event of a GDI leak, this information will enable you to figure out which process initiated by the user spawned the maximum number of GDI handles, and is hence responsible for the GDI leak. |

| | | | |
|--|---|---------|--|
| | <p>Percentage of GDI objects:</p> <p>Indicates what percentage of the configured GDILIMIT is the total number of GDI object handles created by this user's processes.</p> | Percent | <p>This value is calculated using the following formula:</p> $Total\ GDI\ objects / GDILimit * 100$ <p>A value close to 100% is a cause for concern, as it indicates that the count of GDI handles for the user is fast-approaching the permitted GDILIMIT. This hints at a potential GDI leak. You can then use the detailed diagnosis of the <i>Total GDI objects</i> measure to identify which process initiated by the user is spawning the maximum GDI handles and is hence contributing to the leak, and probe further.</p> |
|--|---|---------|--|

The detailed diagnosis of the *Total GDI objects* measure, if enabled, provides the process-wise breakup of the GDI handles created by the user. In the event of a GDI leak, this information will enable you to figure out which process initiated by the user spawned the maximum number of GDI handles, and is hence responsible for the GDI leak.

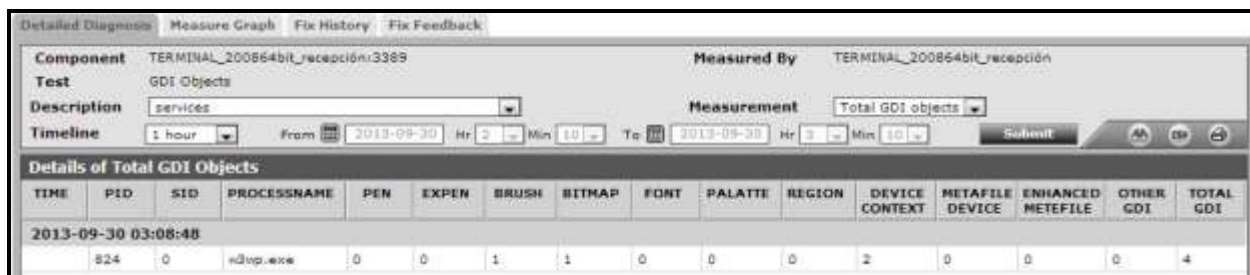


Figure 2.3: The detailed diagnosis of the *Total GDI objects* measure

2.3 The Terminal Applications Layer

The health of a Microsoft RDS server depends upon the health of the applications it hosts. The Terminal Applications test associated with this layer monitors application health.

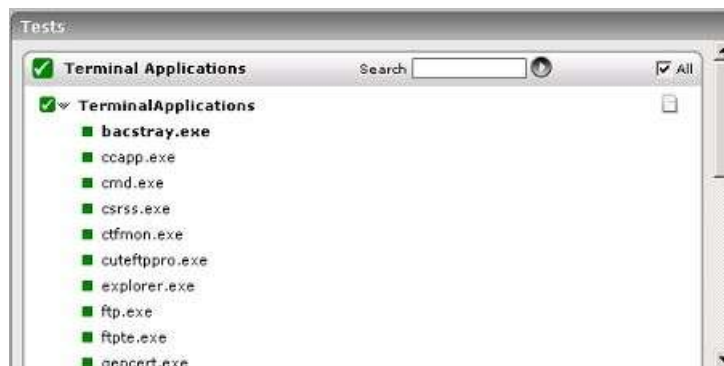


Figure 2.4: Tests associated with the Terminal Applications layer

2.3.1 Terminal Applications Test

This test reports statistics pertaining to the different applications deployed within the Microsoft RDS server and their usage by its clients.

| | |
|---------------------------------|---|
| Purpose | Returns the performance measures pertaining to the applications published on the Microsoft RDS server |
| Target of the test | A Microsoft RDS server |
| Agent deploying the test | An internal agent |

| | |
|--|---|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. APPS - By default, all is displayed here, which will auto-discover and monitor all the applications that are running from the Microsoft RDS server client. To monitor specific applications instead, you have to enter a comma separated list of processName:processPattern pairs which identify the applications published on the server being considered. processName is a string that will be used for display purposes only. processPattern is an expression of the form - *expr* or expr or *expr or expr* or *expr1*expr2*... or expr1*expr2, etc. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. The pattern(s) used vary from one application to another and must be configured per application. For example, if a Microsoft Word application has been published on the Microsoft RDS server, then the PROCESS to be specified is: Word:*winword*, where Word is the string to be displayed in the monitor interface, and *winword* is the application's executable. Other special characters such as slashes (\) can also be used while defining the process pattern. For example, if a server's root directory is /home/egurkha/apache and the server executable named httpd exists in the bin directory, then, the process pattern is "*home/egurkha/apache/bin/httpd*". The test will rediscover the applications every 6th time the test runs. 5. REPORT BY DOMAIN NAME – By default, this flag is set to Yes. This implies that by default, the detailed diagnosis of this test will display the <i>domainname\username</i> of each user who accessed an application on the server. This way, administrators will be able to quickly determine which user logged into the server from which domain. If you want the detailed diagnosis to display only the <i>username</i> of these users, set this flag to No. 6. ENABLE BROWSER MONITORING – By default, this flag is set to No, indicating that the eG agent does not monitor browser activity on the Microsoft RDS server. If this flag is set to Yes, then, whenever one/more IE (Internet Explorer) browser instances on the RDS server are accessed, the detailed diagnosis of the <i>Processes running</i> measure will additionally reveal the URL being accessed via each IE instance and the resources consumed by every URL. Armed with this information, administrators can identify the web sites that are responsible for excessive resource usage by an IE instance. 7. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
| <p>Outputs of the test</p> | <p>One set of results is reported for each application</p> |

MONITORING MICROSOFT RDS SERVERS

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|--|------------------|--|
| | Processes running: Number of instances of the published application currently executing on the Microsoft RDS server | Number | This value indicates if too many or too few instances corresponding to an application are executing on the host. The detailed diagnosis of this measure, if enabled, displays the complete list of processes executing, the users executing them, and their individual resource utilization. |
| | Cpu usage: Percentage of CPU used by the published application | Percent | A very high value could indicate that the specified application is consuming excessive CPU resources. |
| | Memory usage: This value represents the ratio of the resident set size of the memory utilized by the application to the physical memory of the host system, expressed as a percentage. | Percent | A sudden increase in memory utilization for an application may be indicative of memory leaks in the application. |

The detailed diagnosis of the *Processes running* measure, if enabled, provides the list of processes currently executing, the users executing them, and their CPU and memory usage. Using these details, you can quickly detect resource-intensive instances and the user executing them.

| Shows the User and their corresponding PID CPU% MEM% | | | | | |
|--|----------|------|-------|-------|--|
| Time | Username | PID | % CPU | % MEM | |
| 2008/1/9 11:28:12 | egtest | 6036 | 0 | .0191 | |
| 2008/1/9 11:17:59 | egtest | 6036 | 0 | .0191 | |
| 2008/1/9 11:07:39 | egtest | 6036 | 0 | .0233 | |
| 2008/1/9 10:57:53 | egtest | 6036 | 0 | .0233 | |
| 2008/1/9 10:47:49 | egtest | 6036 | 0 | .0233 | |
| 2008/1/9 10:37:33 | egtest | 6036 | 0 | .0233 | |
| 2008/1/9 10:26:43 | egtest | 6036 | 0 | .0233 | |
| 2008/1/9 10:16:24 | egtest | 6036 | 0 | .0516 | |
| 2008/1/9 10:06:48 | egtest | 6036 | 0 | .0516 | |
| 2008/1/9 09:56:20 | egtest | 6036 | 0 | .0516 | |
| 2008/1/9 09:46:24 | egtest | 6036 | 0 | .0516 | |
| 2008/1/9 09:35:43 | egtest | 6036 | 0 | .0516 | |

Figure 2.5: The detailed diagnosis of the Processes running measure

Moreover, if one or more browser instances are found to consume excessive CPU, memory and disk I/O resources on a server or a desktop, then for each such browser instance, administrators can now see a mapping of browser process to URL being accessed, as well as the resources used by each browser process in the detailed diagnosis. Armed with this information, administrators can determine the steps required to avoid excessive resource usage by browser instances – e.g., whether specific web sites are responsible for this, whether users are accessing

web sites (e.g., youtube, facebook, etc.) that they should not be accessing from a corporate network, etc.



Note

- The eG agent will perform browser activity monitoring only if the **ENABLE BROWSER MONITORING** flag is set to **Yes**.
- The eG agent will monitor browser activity only of the browser being accessed is **Internet Explorer**.

2.3.2 App-V Applications Test

This test reports statistics pertaining to the different applications executing on an App-V client and their usage. In addition, this test also reports the statistics pertaining to the processes running on the APP-V client.



Note

This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

| | |
|---------------------------------|---|
| Purpose | Reports statistics pertaining to the different applications executing on an App-V client and their usage. In addition, this test also reports the statistics pertaining to the processes running on the APP-V client. |
| Target of the test | An App-V Client on the target Mincorsoft RDS server |
| Agent deploying the test | An internal agent |

| | | | |
|--|--|--------------------------------|---|
| <p>Configurable parameters for the test</p> | <p>8. TEST PERIOD - How often should the test be executed</p> <p>9. HOST – The host for which the test is to be configured</p> <p>10. PORT – The port at which the specified HOST listens. By default, this is <i>NULL</i>.</p> <p>11. REPORT BY DOMAIN NAME – By default, this flag is set to No. This means that, by default, the test will report metrics for each <i>username</i> only. You can set this flag to Yes, to ensure that the test reports metrics for each <i>domainname username</i>.</p> <p>12. EXTENDED STATISTICS – By default, this test provides you with detailed measures on the resource utilization of each application. If you wish to obtain only the CPU and memory related measures, then set the EXTENDED STATISTICS flag to No.</p> <p>13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>14. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. <p>15.</p> | | |
| <p>Outputs of the test</p> | <p>One set of results for each application of the target App-V Client that is to be monitored</p> | | |
| <p>Measurements made by the test</p> | <p>Measurement</p> | <p>Measurement Unit</p> | <p>Interpretation</p> |
| | <p>Total size: Indicates the total size of this virtual application package.</p> | <p>MB</p> | <p>The detailed diagnosis of this measure lists the Version of the application, Application ID, Version ID of the applicaiton and the application path.</p> |

MONITORING MICROSOFT RDS SERVERS

| | <p>Is loading?: Indicates whether this application is currently loading or not on the App-V client.</p> | | <p>This measure reports a value <i>True</i> if the application is currently being loaded and a value <i>False</i> if otherwise.</p> <p>These measure values and their corresponding numeric values are listed in the table below:</p> <table border="1" data-bbox="933 447 1416 594"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>True</td> <td>1</td> </tr> <tr> <td>False</td> <td>0</td> </tr> </tbody> </table> <p>Note: By default, this measure reports the values <i>Yes</i> or <i>No</i> to indicate whether this application is currently being loaded on the client or not. The graph of this measure however is represented using the numeric equivalents - <i>0</i> or <i>1</i>.</p> | Measure Value | Numeric Value | True | 1 | False | 0 |
|---------------|---|---------|--|---------------|---------------|------|---|-------|---|
| Measure Value | Numeric Value | | | | | | | | |
| True | 1 | | | | | | | | |
| False | 0 | | | | | | | | |
| | <p>Loaded percentage: Indicates the percentage of this application that is currently being loaded on the App-V client.</p> | Percent | | | | | | | |
| | <p>In use?: Indicates whether this application is currently in use or not.</p> | | <p>This measure reports a value <i>True</i> if the application is currently in use and a value <i>False</i> if otherwise.</p> <p>These measure values and their corresponding numeric values are listed in the table below:</p> <table border="1" data-bbox="933 1310 1416 1457"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>True</td> <td>1</td> </tr> <tr> <td>False</td> <td>0</td> </tr> </tbody> </table> <p>Note: By default, this measure reports the values <i>Yes</i> or <i>No</i> to indicate whether this application is currently in use. The graph of this measure however is represented using the numeric equivalents - <i>0</i> or <i>1</i>.</p> | Measure Value | Numeric Value | True | 1 | False | 0 |
| Measure Value | Numeric Value | | | | | | | | |
| True | 1 | | | | | | | | |
| False | 0 | | | | | | | | |

| | <p>Any user based pending tasks available?</p> <p>Indicates whether any tasks are pending for the user using this application.</p> | | <p>This measure reports a value <i>Yes</i> if any tasks are pending for the user using the application and a value <i>No</i> if otherwise.</p> <p>These measure values and their corresponding numeric values are listed in the table below:</p> <table border="1" data-bbox="933 447 1414 594"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Yes</td> <td>1</td> </tr> <tr> <td>No</td> <td>0</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, this measure reports the values <i>Yes</i> or <i>No</i> to indicate whether any tasks are currently pending for the user using this application. The graph of this measure however is represented using the numeric equivalents - <i>0</i> or <i>1</i>.</p> | Measure Value | Numeric Value | Yes | 1 | No | 0 |
|---------------|---|---------------|--|---------------|---------------|-----|---|----|---|
| Measure Value | Numeric Value | | | | | | | | |
| Yes | 1 | | | | | | | | |
| No | 0 | | | | | | | | |
| | <p>Any global based pending tasks available:</p> <p>Indicates whether any global tasks are pending for this application.</p> | | <p>This measure reports a value <i>Yes</i> if any tasks are pending for the user using the application and a value <i>No</i> if otherwise.</p> <p>These measure values and their corresponding numeric values are listed in the table below:</p> <table border="1" data-bbox="933 1115 1414 1262"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Yes</td> <td>1</td> </tr> <tr> <td>No</td> <td>0</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, this measure reports the values <i>Yes</i> or <i>No</i> to indicate whether any tasks are currently pending for the user using this application. The graph of this measure however is represented using the numeric equivalents - <i>0</i> or <i>1</i>.</p> | Measure Value | Numeric Value | Yes | 1 | No | 0 |
| Measure Value | Numeric Value | | | | | | | | |
| Yes | 1 | | | | | | | | |
| No | 0 | | | | | | | | |
| | <p>Processes running:</p> <p>Indicates the number of instances of this application currently executing.</p> | <p>Number</p> | <p>This value indicates if too many or too few instances corresponding to an application are executing on the host. The detailed diagnosis of this measure, if enabled, displays the complete list of processes executing, the users executing them, and their individual resource utilization.</p> | | | | | | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|----------------|--|
| | <p>CPU utilization: Indicates the percentage of CPU used by this application.</p> | Percent | A very high value could indicate that the specified application is consuming excessive CPU resources. |
| | <p>Memory utilization: This value represents the ratio of the resident set size of the memory utilized by the application to the physical memory of the host system, expressed as a percentage.</p> | Percent | A sudden increase in memory utilization for an application may be indicative of memory leaks in the application. |
| | <p>Handle count: Indicates the number of handles opened by this application.</p> | Number | An increasing trend in this measure is indicative of a memory leak in the process. |
| | <p>I/O data rate: Indicates the rate at which processes are reading and writing bytes in I/O operations.</p> | Kbytes/Sec | This value counts all I/O activity generated by each process and includes file, network and device I/Os. |
| | <p>I/O data operations: Indicates the rate at which this application process is issuing read and write data to file, network and device I/O operations.</p> | Operations/Sec | |
| | <p>I/O read data rate: Indicates the rate at which the process is reading data from file, network and device I/O operations.</p> | Kbytes/Sec | |
| | <p>I/O write data rate: Indicates the rate at which the process is writing data to file, network and device I/O operations.</p> | Kbytes/Sec | |
| | <p>Number of threads: Indicates the number of threads that are used by this application.</p> | Number | |

| | | | |
|--|---|-------------------|--|
| | <p>Page fault rate: Indicates the total rate at which page faults are occurring for the threads of all matching application processes.</p> | <p>Faults/Sec</p> | <p>A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.</p> |
| | <p>Virtual memory used: Indicates the amount of virtual memory that is being used by the application.</p> | <p>MB</p> | |
| | <p>Memory working set: Indicates the current size of the working set of a process.</p> | <p>MB</p> | <p>The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use.</p> <p>When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed they will then be soft-faulted back into the Working Set before leaving main memory. If a process pattern matches multiple processes, the memory working set reported is the sum of the working sets for the processes that match the specified pattern. Detailed diagnosis for this test provides details of the individual processes and their individual working sets.</p> <p>Comparing the working set across processes indicates which process(es) are taking up excessive memory. By tracking the working set of a process over time, you can determine if the application has a memory leak or not.</p> |

2.4 The Terminal Users Layer

By continuously monitoring the user behavior on a Microsoft RDS server, administrators can accurately gauge resource usage per user, and derive guidelines for upgrading server capacity and imposing stricter access rules. The tests associated with this layer (see Figure 2.6) facilitate such user-related analysis.

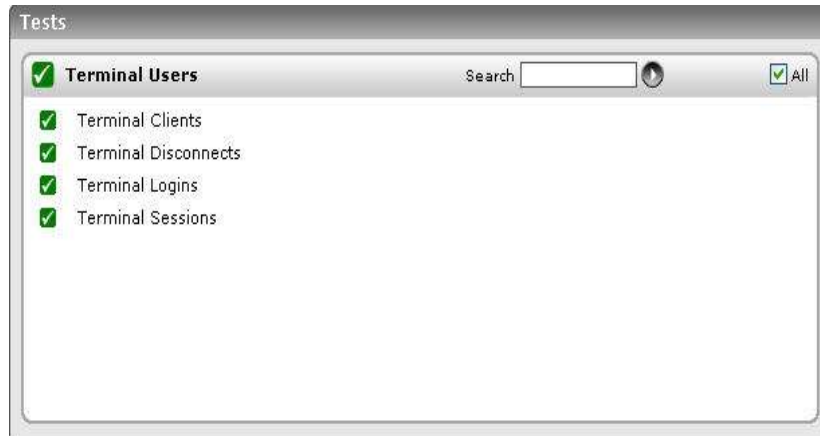


Figure 2.6: Tests associated with the Terminal Users layer

2.4.1 Terminal Sessions Test

This test reports performance statistics related to Microsoft RDS server user sessions.

| | |
|---------------------------------|--|
| Purpose | Reports performance statistics related to Microsoft RDS server user sessions |
| Target of the test | A Microsoft RDS server |
| Agent deploying the test | An internal agent |

| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. IGNORE DOWN SESSION IDS - By default, this parameter is set to <i>65536,65537,65538</i> – these are nothing but the default ports at which the listener component listens. If any of these ports go down, then by default, this test will not count any of the sessions that failed when attempting to connect to that port as a Down session. You can override this default setting by adding more ports or by removing one/more existing ports. 5. REPORTUSINGMANAGERTIME - By default, this flag is set to Yes. This indicates that the user login time displayed in the DETAILED DIAGNOSIS page for this test and in the Thin Client reports will be based on the eG manager's time zone by default. Set this flag to No if you want the login times displayed in the DETAILED DIAGNOSIS page for this test and in the Thin Client reports to be based on the Microsoft RDS server's local time. 6. REPORT BY DOMAIN NAME – By default, this flag is set to Yes. This implies that by default, the detailed diagnosis of this test will display the <i>domainname\username</i> of each user who logged into the Microsoft RDS server. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the <i>username</i> alone, then set this flag to No. 7. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | | | | | | | |
|--|--|---|------------------|----------------|--|---------------|---|--|--|
| <p>Outputs of the test</p> | <p>One set of results for every server being monitored</p> | | | | | | | | |
| <p>Measurements made by the test</p> | <table border="1"> <thead> <tr> <th data-bbox="386 1392 719 1455">Measurement</th> <th data-bbox="727 1392 922 1455">Measurement Unit</th> <th data-bbox="922 1392 1412 1455">Interpretation</th> </tr> </thead> <tbody> <tr> <td data-bbox="386 1455 719 1749"> <p>Active sessions: Indicates the number of active terminal services sessions currently on the server.</p> </td> <td data-bbox="727 1455 922 1749"> <p>Number</p> </td> <td data-bbox="922 1455 1412 1749"> <p>This measure gives an idea of the server workload in terms of active sessions. Tracking the number of active sessions with time, a Microsoft RDS server administrator can obtain information that can help him/her plan the capacity of their Microsoft RDS server farms. The detailed diagnosis capability, if enabled, lists the active and inactive sessions on the Microsoft RDS server.</p> </td> </tr> </tbody> </table> | Measurement | Measurement Unit | Interpretation | <p>Active sessions: Indicates the number of active terminal services sessions currently on the server.</p> | <p>Number</p> | <p>This measure gives an idea of the server workload in terms of active sessions. Tracking the number of active sessions with time, a Microsoft RDS server administrator can obtain information that can help him/her plan the capacity of their Microsoft RDS server farms. The detailed diagnosis capability, if enabled, lists the active and inactive sessions on the Microsoft RDS server.</p> | | |
| Measurement | Measurement Unit | Interpretation | | | | | | | |
| <p>Active sessions: Indicates the number of active terminal services sessions currently on the server.</p> | <p>Number</p> | <p>This measure gives an idea of the server workload in terms of active sessions. Tracking the number of active sessions with time, a Microsoft RDS server administrator can obtain information that can help him/her plan the capacity of their Microsoft RDS server farms. The detailed diagnosis capability, if enabled, lists the active and inactive sessions on the Microsoft RDS server.</p> | | | | | | | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|--------|---|
| | <p>Idle sessions:</p> <p>Indicates the number of sessions that are initialized and are currently ready to accept connections.</p> | Number | To optimize the performance of a server, two default (idle) sessions are initialized before any client connections are made. For performance reasons, the number of idle sessions should be less than ten. Note that this test does not differentiate between RDP and ICA sessions. |
| | <p>Connected sessions:</p> <p>Indicates the current number of sessions that are connected, but no user has logged on to the server.</p> | Number | A consistent increase in the value of this measure could indicate that users are having trouble logging in. Further investigation may hence be required. Note that this test does not differentiate between RDP and ICA sessions. |
| | <p>Connecting sessions:</p> <p>Indicates the number of sessions that are in the process of connecting.</p> | Number | A very high value for this measure indicates a problem with the session or connection. Note that this test does not differentiate between RDP and ICA sessions. |
| | <p>Disconnected sessions:</p> <p>Indicates the number of sessions from which users have disconnected, but which are still active and can be reconnected.</p> | Number | Too many disconnected sessions running indefinitely on a Microsoft RDS server cause excessive consumption of the server resources. To avoid this, a session limit is typically configured for disconnected sessions on the Microsoft RDS server. When a session limit is reached for a disconnected session, the session ends, which permanently deletes it from the server. Note that this test does not differentiate between RDP and ICA sessions. |
| | <p>Listen sessions:</p> <p>Indicates the current number of sessions that are ready to accept connections.</p> | Number | Note that this test does not differentiate between RDP and ICA sessions. |
| | <p>Shadow sessions:</p> <p>Indicates the current number of sessions that are remotely controlling other sessions.</p> | Number | A non-zero value for this measure indicates the existence of shadow sessions that are allowed to view and control the user activity on another session. Such sessions help in troubleshooting/resolving problems with other sessions under their control. |
| | <p>Down sessions:</p> <p>Indicates the current number of sessions that could not be initialized or terminated.</p> | Number | <p>Ideally, the value of this measure should be 0.</p> <p>By default, if sessions to any of these ports – 65536, 65537, 65538 – could not be initialized or terminated, they will not be counted as a ‘down session’.</p> |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|--------|---|
| | Init sessions: Indicates the current number of sessions that are initializing. | Number | A high value for this measure could indicate that many sessions are currently experiencing initialization problems. |
|--|--|--------|---|

The detailed diagnosis capability of the *Active sessions* measure, if enabled, lists the active and inactive sessions on the Microsoft RDS server, and provides details such as the user who initiated the sessions, the session login time, the duration for which the session was idle, etc.

| Time | Username | Sessionname | ID | State | Idle time | Logon time |
|-------------------|----------|-------------|----|--------|-----------|------------------|
| 2008/1/9 11:26:48 | egtest | rdp-tcp#6 | 1 | Active | 11:57 | 1/8/2008 3:15 PM |
| 2008/1/9 11:16:50 | egtest | rdp-tcp#6 | 1 | Active | 11:47 | 1/8/2008 3:15 PM |
| 2008/1/9 11:06:32 | egtest | rdp-tcp#6 | 1 | Active | 11:36 | 1/8/2008 3:15 PM |
| 2008/1/9 10:56:56 | egtest | rdp-tcp#6 | 1 | Active | 11:27 | 1/8/2008 3:15 PM |
| 2008/1/9 10:47:28 | egtest | rdp-tcp#6 | 1 | Active | 11:17 | 1/8/2008 3:15 PM |
| 2008/1/9 10:37:23 | egtest | rdp-tcp#6 | 1 | Active | 11:07 | 1/8/2008 3:15 PM |
| 2008/1/9 10:27:27 | egtest | rdp-tcp#6 | 1 | Active | 10:57 | 1/8/2008 3:15 PM |
| 2008/1/9 10:17:26 | | | | | | |

Figure 2.7: The detailed diagnosis of the Active sessions measure

2.4.2 Terminal Logins Test

This test monitors the new logins to the Microsoft RDS server.

| | |
|---------------------------------|---|
| Purpose | Monitors the new logins to the Microsoft RDS server |
| Target of the test | Any Microsoft RDS server |
| Agent deploying the test | An internal agent |

| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. REPORTUSINGMANAGERTIME - By default, this flag is set to Yes. This indicates that the user login time displayed in the DETAILED DIAGNOSIS page for this test and in the Thin Client reports will be based on the eG manager's time zone by default. Set this flag to No if you want the login times displayed in the DETAILED DIAGNOSIS page for this test and in the Thin Client reports to be based on the Microsoft RDS server's local time. 5. REPORT BY DOMAIN NAME – By default, this flag is set to Yes. This implies that by default, the detailed diagnosis of this test will display the <i>domainname\username</i> of each user session that logged out. This default setting ensures that administrators are able to quickly determine the domains to which the users who logged out belonged. You can set this flag to No if you want detailed diagnosis to display only the <i>username</i> of the users who logged out. 6. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 7. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | | | | | | | |
|--|--|---|------------------|----------------|--|---------------|---|--|--|
| <p>Outputs of the test</p> | <p>One set of results is reported for each Microsoft RDS server being monitored</p> | | | | | | | | |
| <p>Measurements made by the test</p> | <table border="1"> <thead> <tr> <th data-bbox="386 1455 719 1528">Measurement</th> <th data-bbox="727 1455 922 1528">Measurement Unit</th> <th data-bbox="922 1455 1412 1528">Interpretation</th> </tr> </thead> <tbody> <tr> <td data-bbox="386 1528 719 1701"> <p>New logins: Indicates the number of new logins to the Microsoft RDS server in the last measurement period.</p> </td> <td data-bbox="727 1528 922 1701"> <p>Number</p> </td> <td data-bbox="922 1528 1412 1701"> <p>A consistent zero value could indicate a connection issue.</p> </td> </tr> </tbody> </table> | Measurement | Measurement Unit | Interpretation | <p>New logins: Indicates the number of new logins to the Microsoft RDS server in the last measurement period.</p> | <p>Number</p> | <p>A consistent zero value could indicate a connection issue.</p> | | |
| Measurement | Measurement Unit | Interpretation | | | | | | | |
| <p>New logins: Indicates the number of new logins to the Microsoft RDS server in the last measurement period.</p> | <p>Number</p> | <p>A consistent zero value could indicate a connection issue.</p> | | | | | | | |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|---------|---|
| | <p>Percent new logins:</p> <p>Indicates the percentage of current sessions that logged in during the last measurement period.</p> | Percent | |
| | <p>Sessions logging out:</p> <p>Indicates the number of sessions that logged out.</p> | Number | If all the current sessions suddenly log out, it indicates a problem condition that requires investigation. |

The detailed diagnosis of the *Sessions logging out* measure lists the sessions that logged out.

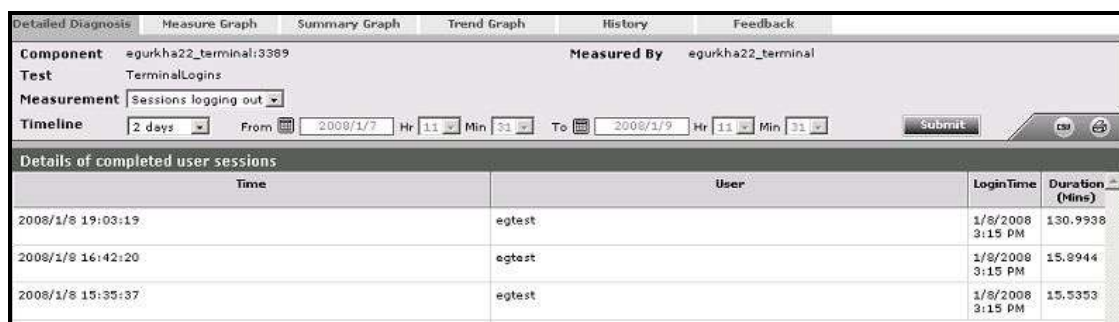


Figure 2.8: The detailed diagnosis of the Sessions logging out measure

2.4.3 Terminal Clients Test

This test measures the client connections to and from a Microsoft RDS server. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Microsoft Terminal* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | |
|---------------------------------|--|
| Purpose | To monitor the client connections to and from a Microsoft RDS server |
| Target of the test | A Microsoft RDS server |
| Agent deploying the test | Internal agent |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|---|--|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Microsoft RDS server SERVERIP - By default, the SERVERIP field will display the IP address of the Microsoft RDS server. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Current connections: The number of TCP connections currently established by clients to the Microsoft RDS server | Number | This measure directly indicates the loading on the Microsoft RDS server from clients. Typically one connection is established per active session to the Microsoft RDS server. |
| | New connections: The number of new TCP connections initiated by clients to the Microsoft RDS server during the last measurement period | Number | Tracking the new connections over time can provide an indication of when clients login to the Microsoft RDS server. A spurt of connections and disconnections may be indicative of sporadic failures of the Microsoft RDS server. |
| | Old connections removed: The number of TCP connections that were removed because the clients may have disconnected from the Microsoft RDS server during the last measurement period | Number | A large number of sudden connection drops may be early warning indicators of problems with the Microsoft RDS server. |

| | | | |
|--|--|------|---|
| | <p>Avg connection duration:</p> <p>The average time from when a connection is established to when the corresponding connection is disconnected. The duration of a connection is measured from its start time to the current time. The accuracy of this measurement is limited by the frequency at which this test is run.</p> | Secs | <p>This value can provide an indicator of how long clients stay connected to a Microsoft RDS server. This information together with the number of simultaneous clients can be useful for capacity planning in Microsoft RDS server environments (i.e., how to size the Microsoft RDS server).</p> |
|--|--|------|---|

2.4.4 Terminal Users Test

A Microsoft RDS server environment is a shared environment in which multiple users connect to a server/server farm and access a wide variety of applications. When server resources are shared, excessive resource utilization by a single user could impact the performance for other users. Therefore, continuous monitoring of the activities of each and every user on the server is critical. Towards this end, the TerminalUsers test assesses the traffic between the user terminal and the server, and also monitors the resources taken up by a user's session on the server. The results of this test can be used in troubleshooting and proactive monitoring. For example, when a user reports a performance problem, an administrator can quickly check the bandwidth usage of the user's session, the CPU/memory/disk usage of this user's session as well as the resource usage of other user sessions. The admin also has access to details on what processes/applications the user is accessing and their individual resource usage. This information can be used to spot any offending processes/ applications.

| | |
|---------------------------------|---|
| Purpose | Tracks every user connection from the Microsoft RDS client to the server, and monitors the resource utilization of every user on the Microsoft RDS server |
| Target of the test | A Microsoft RDS server |
| Agent deploying the test | An internal agent |

| | | | |
|--|---|--|---|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. USERNAMES - Specify the name of the user whose performance statistics need to be generated. Multiple user names can be specified as a comma-separated list. <i>all</i> is used to indicate that all users of the Microsoft RDS server are to be monitored. 5. REPORT BY DOMAIN NAME – By default, this flag is set to Yes. This implies that by default, this test will report metrics for every <i>domainname\username</i>. This way, administrators will know which user logged in from which domain. If you want the test to report metrics for every <i>username</i> only, then set this flag to No. 6. ENABLE BROWSER MONITORING – By default, this flag is set to No, indicating that the eG agent does not monitor browser activity on the Microsoft RDS server. If this flag is set to Yes, then, whenever one/more IE (Internet Explorer) browser instances on the RDS server are accessed, the detailed diagnosis of the <i>User sessions</i> measure will additionally reveal the URL being accessed via each IE instance and the resources consumed by every URL. Armed with this information, administrators can identify the web sites that are responsible for excessive resource usage by an IE instance. 7. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| <p>Outputs of the test</p> | <p>One set of results for every user logged into the Microsoft RDS server</p> | | |
| <p>Measurements made by the test</p> | <p style="text-align: center;">Measurement</p> | <p style="text-align: center;">Measurement Unit</p> | <p style="text-align: center;">Interpretation</p> |
| | <p>User sessions: Represents the current number of sessions for a particular user</p> | <p>Number</p> | <p>A value of 0 indicates that the user is not currently connected to the Microsoft RDS server.</p> |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|------------|--|
| | <p>CPU usage of user's processes:</p> <p>The cpu utilization for a session is the percentage of time that all of the threads/processes of a user session used the processor to execute instructions. If a user is connected via multiple sessions, the value reported is the sum of all cpu utilizations across all the sessions.</p> | Percent | <p>This value indicates the percentage of Cpu resources that are used by applications run by this user. Excessive CPU usage by a user can impact performance for other users. Check the detailed diagnosis to view the offending processes/applications.</p> |
| | <p>Memory usage of user's processes:</p> <p>This value represents the ratio of the resident set size of the memory utilized by the user to the physical memory of the host system, expressed as a percentage. If a user is connected via multiple sessions, the value reported is the sum of all memory utilizations across all the sessions.</p> | Percent | <p>This value indicates the percentage of memory resources that are used up by a specific user. By comparing this value across users, an administrator can identify the most heavy users of the Microsoft RDS server. Check the detailed diagnosis to view the offending processes/applications.</p> |
| | <p>Input bandwidth:</p> <p>Indicates the average bandwidth used for client to server communications for all the sessions of a user</p> | KB/Sec | <p>This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).</p> |
| | <p>Input errors:</p> <p>The average number of input errors of all types for all the sessions of a user. Example: Lost ACK's, badly formed packets, etc.</p> | Errors/Sec | <p>This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).</p> |
| | <p>Output bandwidth:</p> <p>Indicates the average bandwidth used for server to client communications for all the sessions of a user</p> | KB/Sec | <p>This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).</p> |
| | <p>Output errors:</p> <p>The average number of output errors of all types for all the sessions of a user. Example: Lost ACK's, badly formed packets, etc.</p> | Errors/Sec | <p>This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).</p> |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|------------|--|
| | <p>I/O read rate for user's processes:</p> <p>Indicates the rate of I/O reads done by all processes being run by a user.</p> | KBps | <p>These metrics measure the collective I/O activity (which includes file, network and device I/O's) generated by all the processes being executed by a user. When viewed along with the system I/O metrics reported by the DiskActivityTest, these measures help you determine the network I/O. Comparison across users helps identify the user who is running the most I/O-intensive processes. Check the detailed diagnosis for the offending processes/applications.</p> |
| | <p>I/O write rate for user's processes:</p> <p>Indicates the rate of I/O writes done by all processes being run by a user.</p> | KBps | |
| | <p>Faults for user's processes:</p> <p>Indicates the rate of page faults seen by all processes being run by a user.</p> | Faults/Sec | |
| | <p>Virtual memory of user's processes:</p> <p>Indicates the total virtual memory being used by all processes being run by a user.</p> | MB | <p>Comparison across users reveals the user who is being a drain on the virtual memory space.</p> |
| | <p>Handles used by user's processes:</p> <p>Indicates the total number of handles being currently held by all processes of a user.</p> | Number | <p>A consistent increase in the handle count over a period of time is indicative of malfunctioning of programs. Compare this value across users to see which user is using a lot of handles. Check detailed diagnosis for further information.</p> |

| | | | |
|--|---|----------------|---|
| | <p>CPU time used by user's sessions:</p> <p>Indicates the percentage of time, across all processors, this user hogged the CPU.</p> | <p>Percent</p> | <p>The CPU usage for user's processes measure averages out the total CPU usage of a user on the basis of the number of processors. For instance, if your Microsoft RDS server is using an 8-core processor and the total CPU usage of a user across all his/her sessions amounts to 80%, then the value of the CPU usage for user's processes measure for that user will be 10 % (80/8 processors = 10). This accurately denotes the extent of CPU usage in an environment where load is uniformly balanced across multiple processors. However, in environments where load is not well-balanced, the CPU usage for user's processes measure may not be an accurate indicator of CPU usage per user. For instance, if a single processor is used nearly 80% of the time by a user, and other 7 processors in the 8-core processor environment are idle, the CPU usage for user's processes measure will still report CPU usage as 10%. This may cause administrators to miss out on the fact that the user is actually hogging a particular processor! In such environments therefore, its best to use the CPU time used by user's sessions measure! By reporting the total CPU usage of a user across all his/her sessions and across all the processors the target Microsoft RDS server supports, this measure serves as the true indicator of the level of CPU usage by a user in dynamic environments. For instance, in the example above, the CPU time used by user's sessions of the user will be 80% (and not 10%, as in the case of the CPU usage for user's processes measure). A high value or a consistent increase in the value of this measure is hence serious and demands immediate attention. In such situations, use the detailed diagnosis of the CPU usage for user's processes measure to know what CPU-intensive activities are being performed by the user.</p> |
|--|---|----------------|---|

The detailed diagnosis of the *User sessions*, *CPU usage of user's processes*, and *Memory usage of user's processes* measures lists the processes executed by a user on the Microsoft RDS server, and reports the resource usage of each process (see Figure 2.9).

MONITORING MICROSOFT RDS SERVERS

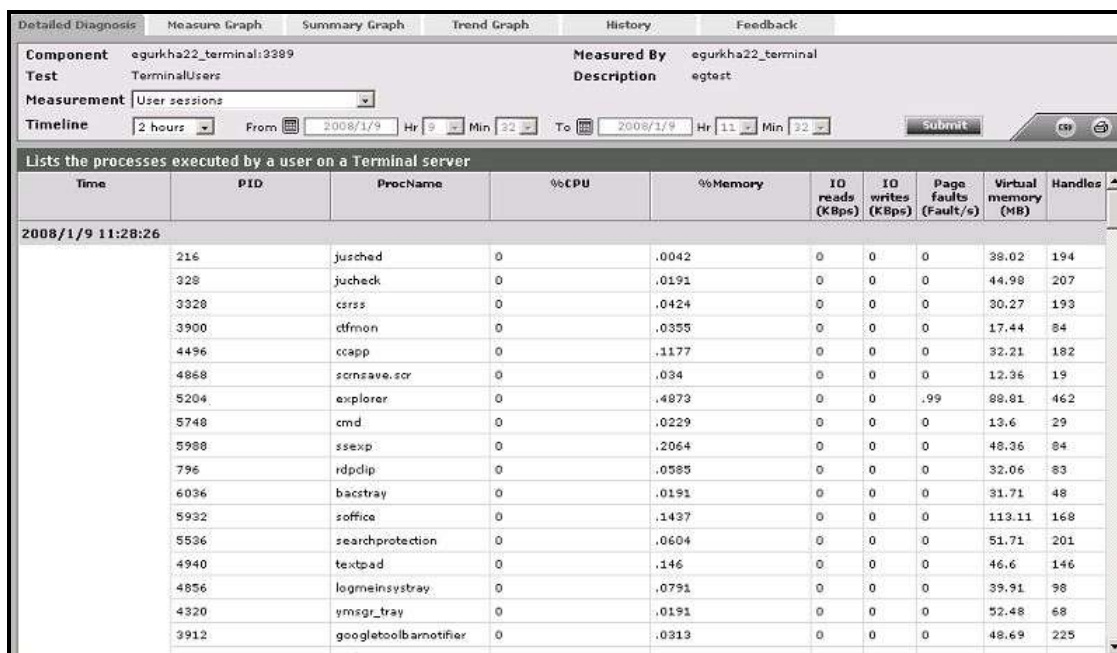


Figure 2.9: The detailed diagnosis of the User sessions measure

Where one or more instances of the Internet Explorer browser are running, the detailed diagnosis additionally displays the website URL accessed using each IE instance, the domain of every URL, and the website title. In the event of excessive resource usage by an IE instance, this information will shed light on the resource-intensive web site that was being accessed.



Note

- The eG agent will perform browser activity monitoring only if the **ENABLE BROWSER MONITORING** flag is set to **Yes**.
- The eG agent will monitor browser activity only of the browser being accessed is **Internet Explorer**.

2.4.5 Terminal Disconnects Test

A user session is terminated when a user logs off from the Citrix/Microsoft RDS server or when the session is abruptly interrupted (e.g., due to server, network, or application errors). When a user logs off, all the applications started by the user are terminated. However, when a user disconnects, the applications started by the user will keep running on the server consuming resources. Hence, the number of disconnected sessions on a Citrix/Microsoft RDS server should be kept to a minimum. Abrupt disconnects can significantly impact the end user experience, and hence, it is important to monitor the number of disconnected sessions at any point of time.

| | |
|----------------------------|---|
| Purpose | Measures the number of disconnected Microsoft RDS server sessions |
| Target of the test | Any Microsoft RDS server |
| Agent deploying the | An internal agent |

| | | | |
|--------------------------------------|--|-------------------------|--|
| test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. RECONNECTPERIOD - This parameter is used by the test while computing the value for the Quick reconnects measure. This measure counts all the users who reconnected to the Microsoft RDS server within the short period of time (in minutes) specified against RECONNECTPERIOD. 5. REPORT BY DOMAIN NAME - By default, this flag is set to Yes. This implies that by default, the detailed diagnosis of this test will display the <i>domainname\username</i> of each user who disconnected from the server recently. This way, administrators will be able to quickly determine which user belongs to which domain. If you want the detailed diagnosis to display the <i>username</i> alone, then set this flag to No. 6. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 7. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results is reported for each Microsoft RDS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total disconnected sessions: Indicates the total number of sessions that are in the disconnected state. | Number | |
| | New disconnects: Indicates the number of sessions that were disconnected in the last measurement period | Number | The detailed diagnosis of this measure, if enabled lists the users who have recently disconnected. |

| | | | |
|--|---|--------|--|
| | <p>Quick reconnects:</p> <p>Indicates the number of users who reconnected soon after a disconnect.</p> | Number | The detailed diagnosis of this measure, if enabled lists the users who have reconnected quickly. |
|--|---|--------|--|

The detailed diagnosis for the *New disconnects* measurement indicates the user, session ID, and client type for each newly disconnected session. This information can be used to track whether specific users are being disconnected often (see Figure 2.10).

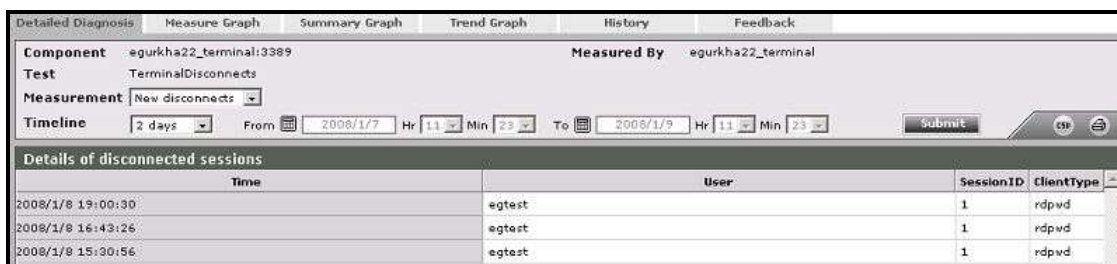


Figure 2.10: The detailed diagnosis of the New disconnects measure

The detailed diagnosis for the *Quick reconnects* measurement indicates the user, session ID, client type, the exact time at which the session disconnected, and duration of the disconnect, for each session that quickly reconnected. This information can be used to track whether specific users are being disconnected often (see Figure 2.11).

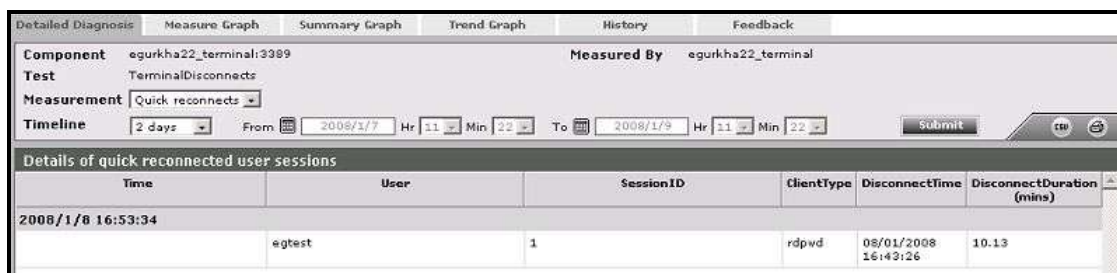


Figure 2.11: The detailed diagnosis of the Quick reconnects measure

2.4.6 Rdp Client Access Test

A Microsoft RDS server environment is a shared environment in which multiple users connect to a server from remote terminals using the Remote Desktop Protocol (RDP). One of the key factors influencing user experience in such an environment is the latency seen by the users when connecting to the server. High network latencies or packet losses during transmission can cause significant slow-downs in request processing by the server. Hence, monitoring latencies between the server and individual client terminals is important.

The Rdp Client Access test is executed by the eG agent on a Microsoft RDS server. This test auto-discovers the users who are currently logged on to the server and the IP address from which they are connecting to the Microsoft RDS server. For each user, the test monitors the quality of the link between the client and the Microsoft RDS server.

MONITORING MICROSOFT RDS SERVERS

Using this test, an administrator can identify user sessions that are being impacted by high latencies or by excessive packet drops. In some cases, a Microsoft RDS server may regard a user session as active, even though the network link connecting the user terminal to the Microsoft RDS server has failed. The Rdp Client Access test alerts administrators to such situations.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Microsoft Terminal* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | | | |
|--|--|-------------------------|---|
| Purpose | Reports on the latencies seen by users connecting to a Microsoft RDS server | | |
| Target | A Microsoft RDS server | | |
| Agent deploying this test | Internal agent | | |
| Configurable parameters for this test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured. 3. PORT - The port at which the HOST listens 4. DISPLAYDOMAIN - By default, the DISPLAYDOMAIN flag is set to Yes; this indicates that the Terminal to Desktop Connection test, by default, will report metrics for every <i>domainname username</i> who is currently connected to the server. This way, administrators can quickly figure out which user is connecting to the server from which domain. You can set this flag to No to ensure that this test reports metrics for each <i>username</i> only. 5. PACKETSIZE - The size of packets used for the test (in bytes) 6. PACKETCOUNT – The number of packets exchanged between the Microsoft RDS server and the user terminal during the test 7. TIMEOUT - How long after transmission should a packet be deemed lost (in seconds) 8. PACKETINTERVAL - Represents the interval (in milliseconds) between successive packet transmissions during the execution of this test. 9. REPORTUNAVAILABILITY – By default, this flag is set to No. This implies that, by default, the test will not report the unavailability of network connection between a user terminal and the Microsoft RDS server. In other words, if the <i>Packet loss</i> measure of this test registers the value <i>100%</i> for any user, then, by default, this test will not report any measure for that user; under such circumstances, the corresponding user name will not appear as a descriptor of this test. You can set this flag to Yes, if you want the test to report and alert you to the unavailability of the network connection between a user terminal and the Microsoft RDS server. | | |
| Outputs of the test | One set of outputs for every user currently connected to the Microsoft RDS server | | |
| Measurements of the test | Measurement | Measurement Unit | Interpretation |
| | Number of sessions: Indicates the current number of sessions for a particular user | Number | The value 0 indicates that the user is not currently connected to the Microsoft RDS server. |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|--|---------|--|
| | Average delay: Indicates the average delay between transmission of a request by the agent on a Microsoft RDS server and receipt of the response back from the user terminal. | Secs | Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing higher latency when connecting to a Microsoft RDS server. |
| | Minimum delay: Indicates the minimum delay between transmission of a request by the agent on a Microsoft RDS server and receipt of the response back from the user terminal. | Secs | A significant increase in the minimum round-trip time is often a sure sign of a poor link between the server and a user's terminal. |
| | Packet loss: Indicates the percentage of packets lost during data exchange between the Microsoft RDS server and the user terminal. | Percent | Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing slowdowns because of poor performance on the network links between their terminals and the Microsoft RDS server. |

Note:

- If the same user is connecting to the Microsoft RDS server from multiple client terminals, the value of the *Number of sessions*, *Avg delay*, and *Packet loss* measures will be averaged across all the sessions of that user. The *Minimum delay* measure, on the other hand, will display the least value reported for *Minimum delay* across all the sessions of that user.
- When a user logs out, the number of sessions will be reduced by 1. If the number of user sessions becomes 0, the corresponding entry for that user in the eG user interface will be removed after a short period of time.
- By default, the Rdp Client Access test spawns a maximum of one thread at a time for monitoring each of the RDP connections to a Microsoft RDS server. Accordingly, the **MaxRdpClientThreads** parameter in the **eg_tests.ini** file (in the <EG_INSTALL_DIR>\manager\config directory) is set to 1 by default. In large Microsoft RDS server farms however, numerous concurrent users attempt to connect to the Microsoft RDS server from multiple remote client terminals. To enhance the efficiency of the test and to make sure that it scales to monitor the large number of RDP connections to the Microsoft RDS server, you might want to consider increasing the value of the **MaxRdpClientThreads** parameter. If this parameter is set to say, 15, then, it implies that the test will spawn a maximum of 15 threads at one shot, thus monitoring 15 RDP connections to the Microsoft RDS server, simultaneously.

2.4.7 RemoteFX User Experience Test

Microsoft® RemoteFX™ enables the delivery of a full Windows user experience to a range of client devices including rich clients, thin clients, and ultrathin clients. RemoteFX delivers a rich user experience for Virtual Desktop Infrastructure (VDI) by providing a 3D virtual adapter, intelligent codecs, and the ability to redirect USB devices in virtual machines. RemoteFX is integrated with the RDP protocol, which enables shared encryption, authentication, management, and device support. RemoteFX also delivers a rich user experience for session-based desktops and RemoteApp programs to a broad range of client devices.

If a remote user’s experience with a RemoteFX-enabled Microsoft RDS server is poor, then administrators should be able to quickly figure out what is causing the quality of the UX to suffer – is it poor frame quality? or severe packet loss? or bad picture output owing to a high compression ratio? or bottleneck in TCP/UDP connectivity? The **RemoteFX User Experience** test helps answer this question. For each remote user connecting to a RemoteFX-enabled Microsoft RDS server, this test measures user experience and reports abnormalities (if any). This way, users who are experiencing a poor visual experience can be isolated and the reason for the same can be ascertained. In addition, the test points you to RemoteFX features that may have to be tweaked in order to improve overall performance.

This test works only on Windows 2008 Service Pack 1 (or above).

| | |
|---------------------------------|---|
| Purpose | For each remote user connecting to a RemoteFX-enabled Microsoft RDS server, this test measures user experience and reports abnormalities (if any) |
| Target of the test | A Microsoft RDS server |
| Agent deploying the test | An internal agent |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|---|---|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Microsoft RDS server 4. REPORT BY DOMAIN NAME – By default, this flag is set to Yes. This implies that by default, this test will report metrics for every <i>domainname\username</i>. This way, administrators will know which user logged in from which domain. If you want the test to report metrics for every <i>username</i> only, then set this flag to No. | | |
| Outputs of the test | One set of results for every user logged into the Microsoft RDS server | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | User sessions: Represents the current number of sessions for a particular user. | Number | A value of 0 indicates that the user is not currently connected to the Microsoft RDS server. |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|------------|---|
| | <p>Average frames encoding time:</p> <p>Indicates the average time taken for encoding the frames of this user.</p> | Secs | <p>Compare the value of this measure across users to know for which user frames encoding took too long.</p> |
| | <p>Frame quality:</p> <p>Indicates the quality of the output frame expressed as a percentage of the quality of the source frame for this user.</p> | Percent | <p>High frame rates produce a smooth representation of frames for the particular user, while low frame rates may cause rough or choppy representation of frames for the particular user. A high value is hence desired for this measure.</p> <p>Compare the value of this measure across users to know which user received the poorest frame quality.</p> |
| | <p>Frames skipped due to insufficient client resources:</p> <p>Indicates the rate at which frames were skipped for this user due to insufficient client resources.</p> | Frames/Sec | <p>A low value is desired for this measure. Compare the value of this measure across users to know which user is connecting from a client sized with inadequate resources.</p> |
| | <p>Frames skipped due to insufficient network resources:</p> <p>Indicates the rate at which frames were skipped for this user due to insufficient network resources.</p> | Frames/Sec | <p>A low value is desired for this measure. Compare the value of this measure across users to know which user is connecting via a network that is sized with inadequate resources.</p> |
| | <p>Frames skipped due to insufficient server resources:</p> <p>Indicates the rate at which frames were skipped for this user due to insufficient server resources.</p> | Frames/Sec | <p>A low value is desired for this measure. Compare the value of this measure across users to know which user was unable to receive frames due to the lack of enough resources on the Microsoft RDS server.</p> |
| | <p>Graphics compression ratio:</p> <p>Indicates the ratio of the number of bytes encoded to the number of bytes input for this user.</p> | Percent | <p>The compression ratio typically affects the quality of the picture. Generally, the higher the compression ratio, the poorer the quality of the resulting picture. Ideally therefore, the value of this measure should be 0. You can compare the value of this measure across users to identify that user whose picture output was very poor owing to high compression.</p> |

MONITORING MICROSOFT RDS SERVERS

| | | | |
|--|---|------------|--|
| | <p>Input frames:</p> <p>Indicates the number of source frames provided per second as input to the RemoteFx graphics for this user.</p> | Frames/Sec | |
| | <p>Output Frames:</p> <p>Indicates the number of source frames sent per second to this user as output of RemoteFx graphics.</p> | Frames/Sec | |
| | <p>Source frames:</p> <p>Indicates number of frames per second composed at the source for this user.</p> | Frames/Sec | |
| | <p>Base TCP round trip time:</p> <p>Indicates the time between initiating a network request and receiving a response over TCP for this user.</p> | Secs | A high value for this measure could indicate a bottleneck in TCP connectivity between the user terminal and the server. |
| | <p>Base UDP round trip time:</p> <p>Indicates the time between initiating a network request and receiving a response over UDP for this user.</p> | Secs | A high value for this measure could indicate a bottleneck in UDP connectivity between the user terminal and the server. |
| | <p>Current TCP bandwidth:</p> <p>Indicates the amount of data that is currently carried from one point to another over TCP for this user.</p> | Kbps | A consistent rise in the value of this measure could indicate that TCP traffic to/from the user is consuming bandwidth excessively. Compare the value of this measure across users to identify that user who is performing bandwidth-intensive operations on the Microsoft RDS server. |

| | | | |
|--|--|---------|---|
| | <p>Current TCP round trip time:</p> <p>Indicates the average time between initiating a network request and receiving a response over TCP for this user.</p> | Secs | A high value could indicate a current problem with TCP connectivity between the user terminal and the server. |
| | <p>Current UDP bandwidth:</p> <p>Indicates the amount of data that is currently carried from one point to another over UDP for this user.</p> | Kbps | A consistent rise in the value of this measure could indicate that UDP traffic to/from the user is consuming bandwidth excessively. Compare the value of this measure across users to identify that user who is performing bandwidth-intensive operations on the Microsoft RDS server. |
| | <p>Current UDP round trip time:</p> <p>Indicates the average time between initiating a network request and receiving a response over UDP for this user.</p> | Secs | A high value could indicate a current problem with UDP connectivity between the user terminal and the server. |
| | <p>Forward error correction rate:</p> <p>Indicates the percentage of forward error corrections performed for this user.</p> | Percent | <p>RemoteFX UDP transport uses Forward Error Correction (FEC) to recover from the lost data packets. In the cases where such packets can be recovered, the transport doesn't need to wait for the data to be retransmitted, which allows immediate delivery of data and prevents Head of Line Blocking. Preventing this stall results in an overall improved responsiveness.</p> <p>A high value is hence desired for this measure.</p> |
| | <p>Loss:</p> <p>Indicates the percentage of packets lost when being transmitted to this user.</p> | Percent | A high value indicates that a large number of packets were lost without being retransmitted. By comparing the value of this measure across users, you can find that user who has suffered the maximum data loss. This could be owing to a bad network connection between the remote user terminal and the server. |

| | | | |
|--|---|---------|--|
| | <p>Retransmission:</p> <p>Indicates the percentage of packets that have been retransmitted to this user.</p> | Percent | Retransmissions should only occur when it is certain that a packet to be retransmitted was actually lost. Redundant retransmissions can also occur because of lost acknowledgments, coarse feedback, and bad retransmissions. Retransmission rates over 5% can indicate degraded network performance on a LAN. The internet may vary between 5 and 15 percent depending upon traffic conditions. Any value above 25 percent indicates an excessive number of retransmissions that will significantly increase the time for the file transfer and annoy the user. |
| | <p>TCP received rate:</p> <p>Indicates the rate at which the data is received over TCP for this user.</p> | Kbps | A high value is desired for these measures as it indicates high TCP throughput. |
| | <p>TCP sent rate:</p> <p>Indicates the rate at which the data is sent over TCP for this user.</p> | Kbps | |
| | <p>UDP received rate:</p> <p>Indicates the rate at which the data is received over UDP for this user.</p> | Kbps | A high value is desired for these measures as it indicates high UDP throughput. |
| | <p>UDP sent rate:</p> <p>Indicates the rate at which the data is sent over UDP for this user.</p> | Kbps | |

Note:

Optionally, you can enable an **EventLog** test for the Microsoft RDS server to closely monitor the system and application events on the server. This test is disabled by default. To enable the test, open the **ENABLE / DISABLE TESTS** page using the Agents -> Tests -> Enable/Disable menu sequence, select **Microsoft Terminal** as the component-type, **Performance** as the *Test type*, select the test from the **DISABLED TESTS** list, and click on >> to move it to the **ENABLED TESTS** list. Finally, click on the **Update** button. This test is mapped to the **Windows Service** layer of the Microsoft RDS server component.

2.4.8 ICA/RDP Listeners Test

The listener component runs on the XenApp/Microsoft RDS server and is responsible for listening for and accepting new ICA/RDP client connections, thereby allowing users to establish new sessions on the XenApp/Microsoft RDS server. If this listener component is down, users may not be able to establish a connection with the XenApp server!

MONITORING MICROSOFT RDS SERVERS

This is why, if a user to the Microsoft RDS server complains of the inaccessibility of the server, administrators should first check whether the listener component is up and running or not. The **ICA/RDP Listeners** test helps administrators perform this check. This test tracks the status of the default listener ports and reports whether any of the ports is down.

| Purpose | Tracks the status of the default listener ports and reports whether any of the ports is down | | | | | | | |
|--|---|-------------------------|--|---------------|---------------|-----|---|----|
| Target | A Microsoft RDS server | | | | | | | |
| Agent deploying this test | Internal agent | | | | | | | |
| Configurable parameters for this test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT - The port at which the HOST listens SESSION IDS – The default listener ports - <i>65536,65537,65538</i> – will be displayed here by default. You can override this default specification by adding more ports or by removing one/more existing ports. | | | | | | | |
| Outputs of the test | One set of outputs for every listener port configured | | | | | | | |
| Measurements of the test | Measurement | Measurement Unit | Interpretation | | | | | |
| | <p>Is listener down?: Indicates whether/not this listener port is down.</p> | | <p>This measure reports the value <i>Yes</i> if the listener port is down and <i>No</i> if the port is up and running. The numeric values that correspond to these measure values are as follows:</p> <table border="1" data-bbox="1008 1167 1378 1383"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Yes</td> <td>0</td> </tr> <tr> <td>No</td> <td>1</td> </tr> </tbody> </table> <p>Note: By default, this measure reports the above-mentioned Measure Values to indicate the status of a listener port. However, the graph of this measure will represent the same using the numeric equivalents only.</p> | Measure Value | Numeric Value | Yes | 0 | No |
| Measure Value | Numeric Value | | | | | | | |
| Yes | 0 | | | | | | | |
| No | 1 | | | | | | | |

Monitoring Active Directory Servers

A directory service consists of both a directory storage system called the “directory store” and a mechanism that is used to locate and retrieve information from the system. The primary functions of the directory service are managed by the Directory System Agent (DSA), which is a process that runs on each domain controller (abbreviated as DC). Active Directory is the directory service that is included with Microsoft Windows. It stores objects that provide information about the real entities that exist in an organization’s network like printers, applications, databases, users etc. Active Directory is a part of the domain controller. It is associated with one or more domains. It stores information about users, specific groups of users like the Administrator, computers, applications, services, files, and distribution lists etc. Active Directory then makes this information available to the users and applications throughout the organization.

Active Directory is an important component of the Windows environment. Like any other Windows applications, its performance can affect the rest of the target environment. Active Directory consumes resources and the administrator needs to be aware of how much of the system’s resources are being consumed over a long term. This helps the administrators to plan for future upgrades. Gathering performance data gives the administrators a good way to see the effects of any optimization efforts that he/she might attempt, and provides a great way for diagnosing problems when they occur. Most of the Windows servers and components are dependent on Active Directory either directly or indirectly. So monitoring the Active Directory server’s performance regularly is necessary to make sure that the target environment is meeting your business and networking goals.

The eG Enterprise suite provides extensive monitoring support to the Active Directory (AD) server operating on Windows 2000, 2003, and 2008/2012. The specialized monitoring model that the eG Enterprise offers (see Figure 3.1) periodically executes a number of tests on the AD server to extract a wide gamut of metrics indicating the availability, responsiveness, and overall health of the AD server and its underlying operating system. Using this model, Active Directory servers can be monitored in an agent-based or an agentless manner.

MONITORING ACTIVE DIRECTORY SERVERS

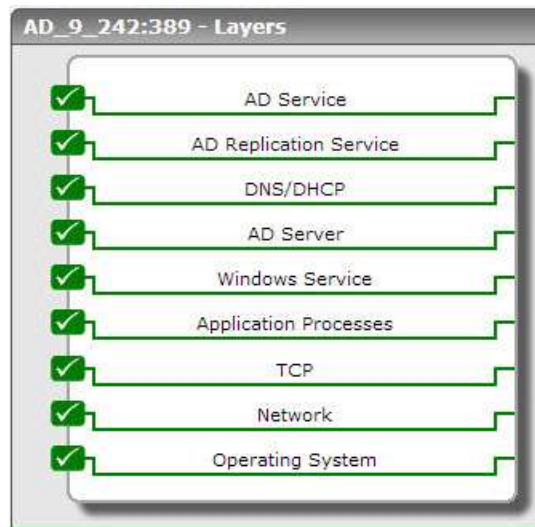


Figure 3.1: Layer model for Active Directory

Using these metrics, an administrator can find quick answers to the following performance queries:

- Is the AD server available?
- How quickly is the server responding to user requests?
- Are there adequate work items to service blocking requests, or are too many requests getting rejected?
- Have any internal server errors been reported recently?
- Have too many login attempts failed?
- Did session timeouts occur too frequently?
- Is the schema cache effectively utilized, or is disk read/write activity high?
- Is the server currently overloaded? Are sufficient domain controllers available in the environment to handle the load?
- Are all changes to the AD server getting replicated across and within sites?
- How many directory synchronizations are in queue? Is the number high enough to force a replication?

The last 5 layers of Figure 3.1 have been discussed in the *Monitoring Unix and Windows Servers* document, and will hence not be discussed again. However, for the *Active Directory* server alone, the **Operating System** layer is mapped to an additional **Net Logon** test. The section that follows will discuss this test in detail. All other sections in this chapter will focus only on the top 3 layers of Figure 3.1.

3.1 The Operating System Layer

The **Operating System** layer of a monitored *Active Directory* server typically runs all the tests that are mapped to the same layer for a *Windows* server or a *Windows Generic* server. The only difference however is that for the Active Directory server, an additional **Net Logon** test is mapped to this layer. This section provides details of the **Net Logon** test.

3.1.1 Net Logon Test

The Netlogon service is responsible for communication between systems in response to a logon request, a domain synchronization request, and a request to promote a Backup Domain Controller (BDC) to a Primary Domain Controller (PDC). The Netlogon service performs several tasks when servicing network logon requests. They are as follows:

- Selects the target domain for logon authentication
- Identifies a domain controller in the target domain to perform authentication
- Creates a secure channel for communication between Netlogon services on the originating and target systems
- Passes an authentication request to the appropriate domain controller
- Returns authentication results to Netlogon on the originating system

Delays in the Netlogon authentication process can often scar a user's overall experience with not just the domain controller, but also with the application that requests for the authentication. In order to avoid undue authentication delays, you can use the **Net Logon** test. This test monitors the Netlogon authentication feature, proactively detects potential authentication bottlenecks, and promptly alerts administrators to what is causing the bottleneck, so that remedial actions can be initiated in good time.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the Netlogon authentication feature, proactively detects potential authentication bottlenecks, and promptly alerts administrators to what is causing the bottleneck, so that remedial actions can be initiated in good time | | |
| Target of the test | An Active Directory server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for every AD server being monitored | | |
| Measurements | Measurement | Measurement Unit | Interpretation |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|------------------|---|--------|--|
| made by the test | <p>Semaphore waiters:</p> <p>Indicates the number of threads currently waiting to acquire the semaphore.</p> | Number | <p>A consistent increase in the value of this measure is a cause for concern, as it indicates that the count of 'busy' semaphores is steadily increasing. This in turn could cause many threads/logon requests to be enqueued, due to the lack of adequate semaphores. Consequently, authentication will be delayed.</p> |
| | <p>Semaphore acquires:</p> <p>Indicates the number of times the semaphore has been acquired over this secure channel during the last measure period.</p> | Number | |
| | <p>Semaphore holders:</p> <p>Indicates the number of threads currently holding the semaphore.</p> | Number | <p>This is a good indicator of the current authentication workload over the secure channel.</p> <p>If the value of this measure is equal to the <i>MaxConcurrentApi</i> registry setting or is fast approaching that value, it indicates that the server is getting overloaded. Authentication delays and timeouts may occur as a result. The typical way to resolve the problem is to raise the maximum allowed worker threads that service that authentication. You can do this by altering the <i>MaxConcurrentApi</i> registry value and then restarting the Net Logon service on the servers.</p> |
| | <p>Semaphore timeouts:</p> <p>Indicates the number of times a thread has timed out waiting for the semaphore over the secure communication channel during the last measure period.</p> | Number | <p>Ideally, this measure has to be 0.</p> <p>A non-zero value for the measure indicates that one/more authentication threads have hit the time-out for the waiting and the logon was denied. This is a sign of a very bad user experience, and typically occurs when the secure channel is overloaded, hung or broken.</p> <p>The typical way to resolve the <i>overload</i> problem is to raise the maximum allowed worker threads that service that authentication. You can do this by altering the <i>MaxConcurrentApi</i> registry value and then restarting the Net Logon service on the servers.</p> |

3.2 The AD Server Layer

The **AD Server** layer verifies the availability and responsiveness of the Active Directory (AD) service from an external location. This layer also monitors the user accesses to the AD server and reports how well the server handles access requests. In the process, the layer also reports useful session-related metrics pertaining to the user sessions on the AD server. Besides, the layer also reports the overall health of the AD database (see Figure 3.2).

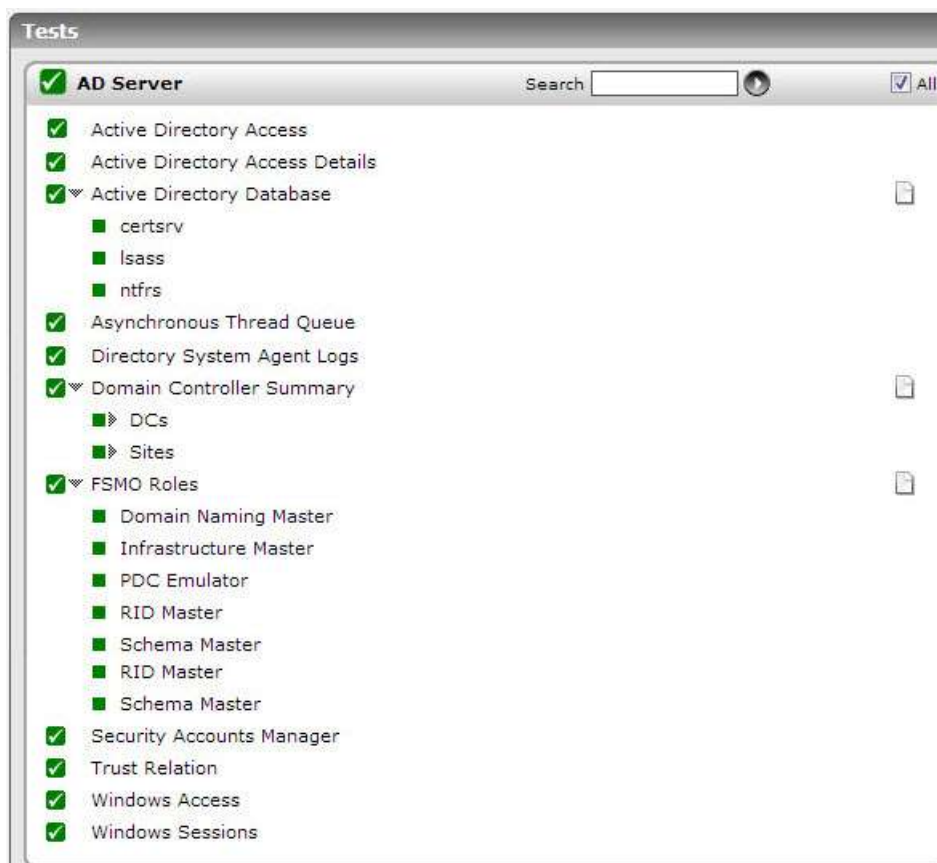


Figure 3.2: The tests associated with the AD Server layer

3.2.1 Asynchronous Thread Queue Test

Monitoring the asynchronous thread queue (ATQ) on an AD server will provide useful pointers to the request processing ability of the server. This test monitors the ATQ, reports the number and nature of requests queued in the ATQ, captures a steady growth (if any) in the length of the queue over time, and thus reveals potential processing bottlenecks on the AD server.

This test applies only to Active Directory Servers installed on Windows 2008.

| | |
|----------------|---|
| Purpose | Monitors the ATQ, reports the number and nature of requests queued in the ATQ, captures a |
|----------------|---|

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|--|
| | steady growth (if any) in the length of the queue over time, and thus reveals potential processing bottlenecks on the AD server | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | ATQ estimated queue delay: Indicates the estimated time the next request will spend in the queue prior to being serviced by the directory service. | Secs | |
| | ATQ outstanding queued requests: Indicate how many requests are queued at the domain controller. | Number | A high level of queuing indicates that requests are arriving at the domain controller faster than they can be processed. This can also lead to a high latency in responding to requests. Delay is the estimated time the next request will spend in the queue prior to being serviced by the directory service, 1.265 seconds. |
| | ATQ request latency: Indicates the average length of time to process a request, not including time spent on the queue. | Secs | A high value of this measure is a cause for concern, as it indicates a processing bottleneck on the AD server. |
| | ATQ threads ldap: Indicates the number of threads that ATQ has currently allocated to servicing LDAP requests. | Number | |

| | | | |
|--|--|--------|---|
| | <p>ATQ thread others:</p> <p>Indicates the number of threads that ATQ has currently allocate to DS services other than LDAP.</p> | Number | |
| | <p>ATQ threads total:</p> <p>Indicates the total number of threads that are either waiting to service an incoming request or are already servicing a request.</p> | Number | <p>If values for this counter and ATQ Threads ldap counter are equal, a queue is likely building on the LDAP port, which will result in long response times. If the two counters are always equal, use Server Performance Advisor to troubleshoot the problem.</p> |

3.2.2 ADAM Access Details Test

This test measures the load on the AD server in terms of the level of read-write activity on the server and the count of search operations performed by the server. In the process, the test reveals the following:

- Which AD services initiated the read-write operations? Which of these services generated the maximum I/O load on the server - is it the LSA? the NSPI? the NTDS? SAM? or the replication service? - this information is useful when administrators are faced with an AD overload, as it accurately points them to the probable sources of the load;
- Which AD service performed the maximum searches on the server? - in the event of an overload, this metric will help you identify that service which could be contributing to the overload;
- Is the server sized with adequate threads to handle the I/O load?

This test applies only to Active Directory Servers installed on Windows 2008.

| | |
|---------------------------------|--|
| Purpose | <p>Measures the load on the AD server in terms of the level of read-write activity on the server and the count of search operations performed by the server. In the process, the test reveals the following:</p> <ul style="list-style-type: none"> • Which AD services initiated the read-write operations? Which of these services generated the maximum I/O load on the server - is it the LSA? the NSPI? the NTDS? SAM? or the replication service? • Which AD service performed the maximum searches on the AD server? • Is the server sized with adequate threads to handle the I/O load? |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 |
| Agent deploying the test | An internal agent |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|---|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Schema cache hit ratio: Indicates the percentage of object name lookups serviced by the Schema Cache. | Percent | <p>All changes made to Active Directory are validated first against the schema. For performance reasons, this validation takes place against a version of the schema that is held in memory on the domain controllers. This "in-memory version," called the schema cache, is updated automatically after the on-disk version has been updated. The schema cache provides mapping between attribute identifiers such as a database column identifier or a MAPI identifier and the in-memory structures that describe those attributes. The schema cache also provides lookups for class identifiers to get in-memory structures describing those classes.</p> <p>A low value of this measure indicates that the Directory Service needs high disk read/write activity to perform its job. This results in poor response time of the components available in the Active Directory.</p> |
| | Notify queue size: Indicates the number of pending update notification requests that have been queued and not transmitted. | Number | <p>When any change in the Active Directory occurs, the originating domain controller sends an update notification requests to the other domain controllers.</p> <p>A high value of this measure indicates that the Active Directory is changing frequently but the update notification requests have not been transmitted to the other domain controllers. This results in a loss of data integrity in the directory store. This problem can be corrected by forcing the replication.</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|--------------|---|
| | <p>Current threads in use:</p> <p>Indicates the current number of threads in use by the directory service (which is different from the number of threads in the directory service process).</p> | Number | <p>This is the number of threads currently servicing client API calls; it can be used to indicate whether additional processors should be used.</p> <p>A fluctuating value for this measure indicates a change in the load.</p> <p>A low value could point to network problems that are preventing client requests from succeeding.</p> |
| | <p>Server binds:</p> <p>Indicates the number of domain controller-to-domain controller binds per second that are serviced by this domain controller.</p> | Binds/Sec | |
| | <p>Directory reads:</p> <p>Indicates the rate of directory reads.</p> | Reads/Sec | <p>These measures serve as effective indicators of the ability of the AD server to process read, write, and search requests.</p> |
| | <p>Directory writes:</p> <p>Indicates the rate of directory writes.</p> | Writes/Sec | |
| | <p>Directory searches:</p> <p>Indicates the number of directory searches per second.</p> | Searches/Sec | |
| | <p>DS reads from DRA:</p> <p>Indicates the percentage of reads on the directory by replication.</p> | Percent | <p>If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the <i>DS reads from KCC</i>, <i>DS reads from LSA</i>, <i>DS reads from NSPI</i>, <i>DS reads from NTDS</i>, and <i>DS reads from SAM</i> measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|----------------|--|
| | <p>DS reads from KCC:</p> <p>Indicates the percentage of reads performed by the Knowledge Consistency Checker on the directory.</p> | <p>Percent</p> | <p>The Knowledge Consistency Checker (KCC) generates the replication topology by specifying what domain controllers will replicate to which other domain controllers in the site. The KCC maintains a list of connections, called a replication topology, to other domain controllers in the site. The KCC ensures that changes to any object are replicated to all site domain controllers and updates go through no more than three connections.</p> <p>If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the <i>DS reads from DRA</i>, <i>DS reads from LSA</i>, <i>DS reads from NSPI</i>, <i>DS reads from NTDS</i>, and <i>DS reads from SAM</i> measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS reads from LSA:</p> <p>Indicates the percentage of reads performed by the Local Security Authority on the directory.</p> | <p>Percent</p> | <p>The Local Security Authority (LSA) is the security subsystem responsible for all interactive user authentication and authorization services on a local computer. The LSA is also used to process authentication requests made through the Kerberos V5 protocol or NTLM protocol in Active Directory.</p> <p>If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the <i>DS reads from DRA</i>, <i>DS reads from KCC</i>, <i>DS reads from NSPI</i>, <i>DS reads from NTDS</i>, and <i>DS reads from SAM</i> measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS?</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|----------------|---|
| | <p>DS reads from NSPI:</p> <p>Indicates the percentage of reads performed by the Name Service Provider Interface (NSPI) on the directory.</p> | <p>Percent</p> | <p>The Name Service Provider Interface (NSPI) is the protocol by which Messaging API (MAPI) clients access the AD DS.</p> <p>Exchange Address Book clients use the client MAPI provider Emsabp32.dll to look up e-mail addresses in the global catalog. The client-side MAPI provider communicates with the server through the proprietary Name Service Provider Interface (NSPI) RPC interface.</p> <p>If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the <i>DS reads from KCC</i>, <i>DS reads from LSA</i>, <i>DS reads from DRA</i>, <i>DS reads from NTDS</i>, and <i>DS reads from SAM</i> measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? or the NSPI?</p> |
| | <p>DS reads from NTDS:</p> <p>Indicates the percentage of reads performed by the name service directory APIs on the directory.</p> | <p>Percent</p> | <p>If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the <i>DS reads from KCC</i>, <i>DS reads from LSA</i>, and <i>DS reads from DRA</i>, <i>DS reads from NSPI</i>, and <i>DS reads from SAM</i> measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? or the SAM?</p> |
| | <p>DS reads from SAM:</p> <p>Indicates the percentage of reads performed by the Security Account Manager (SAM) on the directory.</p> | <p>Percent</p> | <p>The Security Accounts Manager (SAM) is used for verifying passwords and for checking passwords against any existing password policies that are in effect on a domain controller.</p> <p>If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the <i>DS reads from KCC</i>, <i>DS reads from LSA</i>, and <i>DS reads from DRA</i>, <i>DS reads from NSPI</i>, and <i>DS reads from NTDS</i> measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? or the NTDS?</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|---------|--|
| | <p>DS writes from DRA:</p> <p>Indicates the percentage of writes on the AD server by replication.</p> | Percent | <p>If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the <i>DS writes from KCC</i>, <i>DS writes from LSA</i>, <i>DS writes from NSPI</i>, <i>DS writes from NTDS</i>, and <i>DS writes from SAM</i> measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS writes from KCC:</p> <p>Indicates the percentage of writes performed by the Knowledge Consistency Checker on the directory.</p> | Percent | <p>If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the <i>DS writes from DRA</i>, <i>DS writes from LSA</i>, <i>DS writes from NSPI</i>, <i>DS writes from NTDS</i>, and <i>DS writes from SAM</i> measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the KCC? the LSA? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS writes from LSA:</p> <p>Indicates the percentage of writes performed by the Local Security Authority on the directory.</p> | Percent | <p>If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the <i>DS writes from DRA</i>, <i>DS writes from KCC</i>, <i>DS writes from NSPI</i>, <i>DS writes from NTDS</i>, and <i>DS writes from SAM</i> measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS writes from NSPI:</p> <p>Indicates the percentage of writes performed by the Name Service Provider Interface (NSPI) on the directory.</p> | Percent | <p>If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the <i>DS writes from DRA</i>, <i>DS writes from KCC</i>, <i>DS writes from LSA</i>, <i>DS writes from NTDS</i>, and <i>DS writes from SAM</i> measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS writes from NTDS:</p> <p>Indicates the percentage of writes performed by the name service directory APIs on the directory.</p> | Percent | <p>If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the <i>DS writes from DRA</i>, <i>DS writes from KCC</i>, <i>DS writes from LSA</i>, <i>DS writes from NSPI</i>, and <i>DS writes from SAM</i> measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|---------|---|
| | <p>DS writes from SAM:</p> <p>Indicates the percentage of writes performed by the Security Accounts Manager (SAM) on the directory.</p> | Percent | <p>If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the <i>DS writes from DRA</i>, <i>DS writes from KCC</i>, <i>DS writes from LSA</i>, <i>DS writes from NSPI</i>, and <i>DS writes from NTDS</i> measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS searches from DRA:</p> <p>Indicates the percentage of searches performed by the replication service on the AD server.</p> | Percent | <p>If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the <i>DS searches from KCC</i>, <i>DS searches from LSA</i>, <i>DS searches from NSPI</i>, <i>DS searches from NTDS</i>, and <i>DS searches from SAM</i> measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS searches from KCC:</p> <p>Indicates the percentage of searches performed by the Knowledge Consistency Checker on the directory.</p> | Percent | <p>If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the <i>DS searches from DRA</i>, <i>DS searches from LSA</i>, <i>DS searches from NSPI</i>, <i>DS searches from NTDS</i>, and <i>DS searches from SAM</i> measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS searches from LSA:</p> <p>Indicates the percentage of searches performed by the Local Security Authority on the directory.</p> | Percent | <p>If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the <i>DS searches from DRA</i>, <i>DS searches from KCC</i>, <i>DS searches from NSPI</i>, <i>DS searches from NTDS</i>, and <i>DS searches from SAM</i> measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |

| | | | |
|--|--|---------|---|
| | <p>DS searches from NSPI:</p> <p>Indicates the percentage of searches performed by the Name Service Provider Interface (NSPI) on the directory.</p> | Percent | <p>If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the <i>DS searches from DRA</i>, <i>DS searches from KCC</i>, <i>DS searches from LSA</i>, <i>DS searches from NTDS</i>, and <i>DS searches from SAM</i> measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS searches from NTDS:</p> <p>Indicates the percentage of searches performed by the name service directory APIs on the directory.</p> | Percent | <p>If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the <i>DS searches from DRA</i>, <i>DS searches from KCC</i>, <i>DS searches from LSA</i>, <i>DS searches from NSPI</i>, and <i>DS searches from SAM</i> measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |
| | <p>DS searches from SAM:</p> <p>Indicates the percentage of searches performed by the Security Accounts Manager (SAM) on the directory.</p> | Percent | <p>If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the <i>DS searches from DSA</i>, <i>DS searches from KCC</i>, <i>DS searches from LSA</i>, <i>DS searches from NSPI</i>, and <i>DS searches from NTDS</i> measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?</p> |

3.2.3 ADAM Database Test

This test reports critical statistics pertaining to the usage of the database caches, and the overall health of the AD database.

| | |
|---------------------------------|---|
| Purpose | Reports critical statistics pertaining to the usage of the database caches, and the overall health of the AD database |
| Target of the test | An Active Directory server |
| Agent deploying the test | An internal agent |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for every AD server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Database cache hits : Indicates the percentage of page requests of the database file that were occupied in a cache before responding to the request. | Percent | Ideally, the value of this measure should be moderate. A high value of this measure indicates the high utilization of physical memory. In such a case, you can add the required memory to the database. |
| | Database table cache hits: Indicates the percentage of database tables that were opened using cached schema information. | Percent | Ideally, the value of this measure should be high. |
| | Log records waiting: Indicates the rate of log record stalls, per second. | Records/Sec | |
| | Log threads waiting: Indicates the current number of threads waiting for data to be written to the log so that database updation will be executed. | Number | |

3.2.4 Active Directory Access Test

This test monitors the availability and response time from clients of an Active Directory server from an external perspective.

| | |
|---------------------------------|---|
| Purpose | Monitors the availability and response time from clients of an Active Directory server from an internal perspective |
| Target of the test | An Active Directory or Domain Controller |
| Agent deploying the test | An external agent |

| | | | |
|---|---|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. DOMAIN - The default value of the DOMAIN parameter will be <i>none</i>. In Windows 2003 environments however, the ADServerTest will function effectively only if a "fully qualified domain name" is provided in the DOMAIN text box. USER - Provide the name of a domain user in the USER text box. This can be <i>none</i> for Windows 2000 environments. PASSWORD - Provide the password for the domain user specified above, in the PASSWORD text box. This can be <i>none</i> for Windows 2000 environments. CONFIRM PASSWORD – Confirm the PASSWORD by retyping it here. CONNECTTIMEOUT - By default, this is set to 30 seconds. This implies that by default, the test will wait for 30 seconds to establish a connection with the target Active Directory server. If a connection is established within the default 30 second period, then the test will report that the server is available; if the test is unable to connect to the server within the default period, then it will report that the server is unavailable. If it generally takes a longer time for clients to connect to the AD server in your environment, then, you may want to change the CONNECTTIMEOUT period so that, the test does not time out before the connection is established, and consequently present an "untrue" picture of the availability of the server. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active directory availability: Indicates the availability of the server. | Percent | The availability is 100% when the server is responding to a request and 0% when it is not. Availability problems may be caused by a misconfiguration / malfunctioning of the server, or if the server has not been started. |
| | Active directory response time: Indicates the time taken by the server to respond to a user query | Secs | A sudden increase in response time is indicative of a bottleneck at the server. |

3.2.5 Windows Access Test

This test monitors the accesses to an AD server.

| | |
|---------------------------|---|
| Purpose | Monitors the accesses to the Windows server |
| Target of the test | An Active Directory server or a Domain Controller |
| Agent | An internal agent |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--------------------------------------|--|-------------------------|--|
| deploying the test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for every AD server or domain controller being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Blocking request rejects: The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items | Reqs/sec | If the number of blocking request rejects is high, you may need to adjust the <code>MaxWorkItem</code> or <code>MinFreeWorkItems</code> server parameters |
| | Permission errors: The number of times opens on behalf of clients have failed with <code>STATUS_ACCESS_DENIED</code> in the last measurement period | Number | Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected. |
| | File access denied errors: The number of times accesses to files opened successfully were denied in the last measurement period | Number | This number indicates attempts to access files without proper access authorization. |
| | Internal server errors: This value indicates the number of times an internal server error was detected in the last measurement period. | Number | Unexpected errors usually indicate a problem with the server. |
| | Data received: The rate at which the server has received data from the network | Kbytes/sec | This metric indicates how busy the server is. |
| | Data transmitted: The rate at which the server has sent data over the network | Kbytes/sec | This metric indicates how busy the server is. |
| | Resource shortage errors: The number of times <code>STATUS_DATA_NOT_ACCEPTED</code> was returned to clients in the last measurement period | Number | A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the <code>InitWorkItems</code> or <code>MaxWorkItems</code> server parameters might need to be adjusted. |

| | | | |
|--|---|------|--|
| | <p>Avg response time: Average time taken by the server to respond to client requests</p> | Secs | This is a critical measure of server health. |
|--|---|------|--|

3.2.6 Windows Sessions Test

This test reports various session-related statistics for an AD server.

| | | | |
|---|--|-------------------------|--|
| Purpose | Reports various session-related statistics for a Windows server | | |
| Target of the test | An AD server or a Windows Domain Controller | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for every AD server or domain controller being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Logons: Rate of logons to the server</p> | Reqs/sec | This measure reports the rate of all interactive, network, and service logons to a windows server. The measure includes both successful and failed logons. |
| | <p>Logon errors: Number of logons in the last measurement period that had errors</p> | Number | This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server. |
| | <p>Current sessions: The number of sessions currently active in a server</p> | Number | This measure is one of the indicators of current server activity. |
| | <p>Sessions with errors: The number of sessions in the last measurement period that were closed to unexpected error conditions</p> | Number | Sessions can be closed with errors if the session duration reaches the autodisconnect timeout. |

| | | | |
|--|---|--------|---|
| | <p>Sessions forced off:</p> <p>The number of sessions in the last measurement period that have been forced to logoff</p> | Number | This value indicates how many sessions were forced to logoff due to logon time constraints. |
| | <p>Sessions logged off:</p> <p>The number of sessions in the last measurement period that were terminated normally</p> | Number | Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low. |
| | <p>Sessions timed out:</p> <p>The number of sessions that have been closed in the last measurement period due to their idle time exceeding the AutoDisconnect parameter for the server</p> | Number | The number of session timed out gives an indication of whether the AutoDisconnect setting is helping to conserve server resources |

3.2.7 FSMO Roles Test

FSMO stands for Flexible Single Master Operations, and FSMO roles (also known as operations master roles) help you prevent conflicts in your Active Directory.

For most Active Directory objects, the task of updating can be performed by any Domain Controller except those Domain Controllers that are read-only. Updates such as computer object properties, renamed organizational units, and user account password resets can be handled by any writable domain controller.

After an object is changed on one domain controller, those changes are propagated to the other domain controllers through replication. During replication all of the Domain Controllers share their updates. So a user that has their password reset in one part of the domain may have to wait until those changes are replicated to the Domain Controller that they are signing in from.

This model works very well for most objects. In the case of any conflicts, such as a user's password being reset by both the central helpdesk as well as an administrator working at the user's site, then conflicts are resolved by whichever made the last change. However, there are some changes that are too important, and are not well suited to this model.

There are 5 specific types of updates to Active Directory that are very specific, and conflicts should be avoided. To help alleviate any potential conflicts, those updates are all performed on a single Domain Controller. And though each type of update must be performed on a single Domain Controller, they do not all have to be handled by the same Domain Controller.

These types of updates are handled by Domain Controllers Flexible Single Master Operations roles, or FSMO roles. Each of the five roles is assigned to only one domain controller.

There are five FSMO roles in every Active Directory forest. They are:

- Schema Master
- Domain Naming Master

MONITORING ACTIVE DIRECTORY SERVERS

- Infrastructure Master
- Relative ID (RID) Master
- Primary Domain Controller (PDC) Emulator

Among these five FSMO roles, the following three FSMO roles are needed only once in every domain in the forest:

- Infrastructure Master
- Relative ID (RID) Master
- Primary Domain Controller (PDC) Emulator

If a domain controller configured with a specific FSMO role is suddenly rendered unavailable or is unreachable, then that particular function cannot be performed. This in turn implies that the types of updates that will otherwise be handled by that domain controller can no longer be processed, thus creating a climate of conflict in the AD environment. With the help of the **FSMO Roles** test however, you can rapidly detect the unavailability of an FSMO domain controller over the network, isolate potential network connectivity issues and latencies, and spot real/probable delays in LDAP binding, so that such issues can be promptly remedied and conflicts prevented.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Helps rapidly detect the unavailability of an FSMO domain controller over the network, isolate potential network connectivity issues and latencies, and spot real/probable delays in LDAP binding, so that such issues can be promptly remedied and conflicts prevented. | | |
| Target of the test | An AD server or a Windows Domain Controller | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 4. HOST – The host for which the test is to be configured 5. PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for each FSMO role | | |
| Measurements | Measurement | Measurement Unit | Interpretation |

| | | | |
|------------------|---|------|---|
| made by the test | <p>LDAP bind time:</p> <p>Indicates the time taken for the last successful LDAP bind.</p> | Secs | <p>In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as binding. When a programmatic object, such as an IADs (Interface Adapter Device) or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be bound to the directory object.</p> <p>The method for programmatically binding to an Active Directory object will depend on the programming technology that is used.</p> <p>All bind functions and methods require a binding string. The form of the binding string depends on the provider. Active Directory Domain Services are supported by two providers, <i>LDAP</i> and <i>WinNT</i>.</p> <p>Beginning with Windows 2000, the LDAP provider is used to access Active Directory Domain Services. The LDAP binding string can take one of the following forms:</p> <p>"LDAP://<host name>/<object name>"</p> <p>"GC://<host name>/<object name>"</p> <p>Ideally, the value of this measure should be low. A high value for this measure could be a possible indication of network-related problems or of the hardware that needs to be upgraded immediately.</p> <p>This measure will not be reported if the value of the <i>Availability</i> measure is 0.</p> |
| | <p>Avg network delay:</p> <p>Indicates the average delay between transmission of packet to a target and receipt of the response to the packet at the source.</p> | Secs | <p>An increase in network latency could result from misconfiguration of the router(s) along the path, network congestion, retransmissions at the network, etc. The detailed diagnosis capability, if enabled, lists the hop-by-hop connectivity and delay.</p> <p>This measure will not be reported if the value of the <i>Availability</i> measure is 0.</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|---------|---|
| | <p>Minimum network delay:</p> <p>Indicates the minimum time between transmission of a packet and receipt of the response back.</p> | Secs | <p>A significant increase in the minimum round-trip time is often a sure sign of network congestion.</p> <p>This measure will not be reported if the value of the <i>Availability</i> measure is 0.</p> |
| | <p>Packet loss:</p> <p>Indicates the percentage of packets lost during transmission from source to target and back.</p> | Percent | <p>Packet loss is often caused by network buffer overflows at a network router or by packet corruptions over the network. The detailed diagnosis for this measure provides a listing of routers that are on the path from the external agent to target server, and the delays on each hop. This information can be used to diagnose the hop(s) that could be causing excessive packet loss/delays.</p> <p>This measure will not be reported if the value of the <i>Availability</i> measure is 0.</p> |
| | <p>Availability:</p> <p>Indicates whether/not this FSMO role is available over the network.</p> | Percent | <p>A value of 100 indicates that the FSMO role is available. The value 0 indicates that the FSMO role is not available.</p> <p>Typically, the value 100 corresponds to a <i>Pkt_loss_pct</i> of 0.</p> <p>If the FSMO role is not available over the network i.e., if this measure reports a value 0, all other measures applicable for this test will not be reported.</p> |

3.2.8 Directory System Agent Logs Test

This test monitors the Active Directory database files and log files for file size, and also monitors free disk space on the hosting volumes.

| | |
|---------------------------------|--|
| Purpose | Monitors the Active Directory database files and log files for file size, and also monitors free disk space on the hosting volumes |
| Target of the test | An AD server |
| Agent deploying the test | An internal agent |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for every AD server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Directory system agent DB size: Indicates the size of the database files on the AD server. | MB | |
| | System volume size: Indicates the size of the SYSVOL folder - SYSVOL is the shared directory on domain controllers that contains Group Policy and logon script information. | MB | |
| | Directory system agent log file size: Indicates the size of the log files on the AD server. | MB | |
| | Directory system agent free log space: Indicates the amount of free space on the volume hosting log files. | MB | Ideally, this value should be high. |
| | Directory system agent free DB space: Indicates the amount of free space on the volume hosting database files. | MB | Ideally, this value should be high. If the free space for database files is very low, then the AD server might be rendered unable to update objects. |
| | System volume share availability: Indicates whether the SYSVOL folder is available or not. | Percent | If the value of this measure is 100, it indicates the SYSVOL folder is available. The value 0 on the other hand, indicates that the folder is not available. |

3.2.9 Domain Controller Summary

Use this test to know the number and names of all domain controllers that manage the servers and users in the domains of interest to you.

This test runs only on Active Directory servers operating on Windows 2008.

| | | | |
|---|--|-------------------------|--|
| Purpose | Reports the number and names of all domain controllers that manage the servers and users in the domains of interest to you | | |
| Target of the test | An AD server on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the AD server 4. DNS NAME - Provide a comma-separated list of the fully qualified domain names of all the domains that you want the test to scan for domain controllers. For instance, your specification can be, <i>chn.eginnovations.com,maz.eginnovations.com</i>. | | |
| Outputs of the test | One set of results for every domain name configured against DNS NAME | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Domain Controllers: Indicates the number of domain controllers in this domain. | Number | The detailed diagnosis of this measure lists the names of all domain controllers in a chosen domain. |

3.2.10 Security Accounts Manager Test

Every Windows computer has a local Security Accounts Manager (SAM). The SAM is responsible for a few functions. First, it is responsible for storing the local users and groups for that computer. Second, the local SAM is responsible for authenticating logons. When a computer is not joined to a domain, the only option is to use the local SAM to perform the authentication.

If too many computer/user creations in SAM fail or if SAM takes too long to enumerate, evaluate, and authenticate users/user groups, the user experience with the computer is bound to be impacted adversely. By periodically monitoring the operations of SAM, administrators can proactively detect potential problem conditions and plug the holes, so that the user experience remains unaffected. The **Security Accounts Manager** test does just that. At configured intervals, this test checks how well SAM performs its core functions, and promptly reports real/probable failures and latencies to the administrator.

This test applies only to Active Directory Servers installed on Windows 2008 and above.

| | |
|---------------------------------|---|
| Purpose | At configured intervals, this test checks how well SAM performs its core functions, and promptly reports real/probable failures and latencies to the administrator. |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 and above |
| Agent deploying the test | An internal agent |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Machine creation attempts: Indicates the number of attempts per second to create computer accounts. | Number | |
| | User creation attempts: Indicates the number of attempts per second to create user accounts. | Number | |
| | Successful user creations: Indicates the number of user accounts successfully created per second. | Number | Ideally, the value of this measure should be equal to the value of the <i>User creation attempts</i> measure. A low value is a cause for concern, as it indicates that many user creation attempts are failing; the reasons for the same have to be ascertained and addressed soon. |
| | Successful computer creations: Indicates the number of computers successfully created per second. | Number | Ideally, the value of this measure should be equal to the value of the <i>Machine creation attempts</i> measure. A low value is a cause for concern, as it indicates that many machine creation attempts are failing; the reasons for the same have to be ascertained and addressed soon. |
| | GC evaluations: Indicates the number of SAM global catalog evaluations per second. | Number | |
| | Enumerations: Indicates the number of net user, net group, and net local function enumerations per second. | Connections/Sec | |

| | | | |
|--|---|-----------------|---|
| | <p>Display information queries:</p> <p>Indicates the number of queries per second to obtain display information.</p> | Connections/Sec | |
| | <p>Account group evaluation latency:</p> <p>Indicates the time taken by SAM to evaluate an account group.</p> | Secs | <p>This indicates the mean latency of the last 100 account and universal group evaluations performed for authentication.</p> <p>A high value could indicate a bottleneck.</p> |
| | <p>Resource group evaluation latency:</p> <p>Indicates the time taken by SAM to evaluate a resource group.</p> | Secs | <p>This indicates the mean latency of the last 100 resource group evaluations performed for authentication.</p> <p>A high value could indicate a bottleneck.</p> |

3.2.11 Trust Relation Test

Trusts are relationships that are established between domains or forests that enable users in one domain or forest to be authenticated by a domain controller in another domain or forest. Trusts allow users in one domain or forest to access resources in a different domain or forest.

This test automatically discovers the trust relationship that the configured domain shares with other domains, and brings to light problems (if any).

Note:
This test will not work on an Active Directory server running on Windows 2000.

| | |
|---------------------------------|--|
| Purpose | Automatically discovers the trust relationship that the configured domain shares with other domains, and brings to light problems (if any) |
| Target of the test | An AD server |
| Agent deploying the test | An internal agent |

| | | | |
|---|---|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows server 4. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every AD server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Trust errors: Indicates the number of errors in the trust relationship between the configured domain and other domains. | Number | Ideally, this value should be 0. In the event of the occurrence of one/more errors, you can use the detailed diagnosis capability of this measure to view elaborate error descriptions, and accordingly investigate the problem further. |

3.3 The DNS/DHCP Layer

The tests mapped to this layer perform periodic health checks on the DNS and DHCP services that AD relies on.

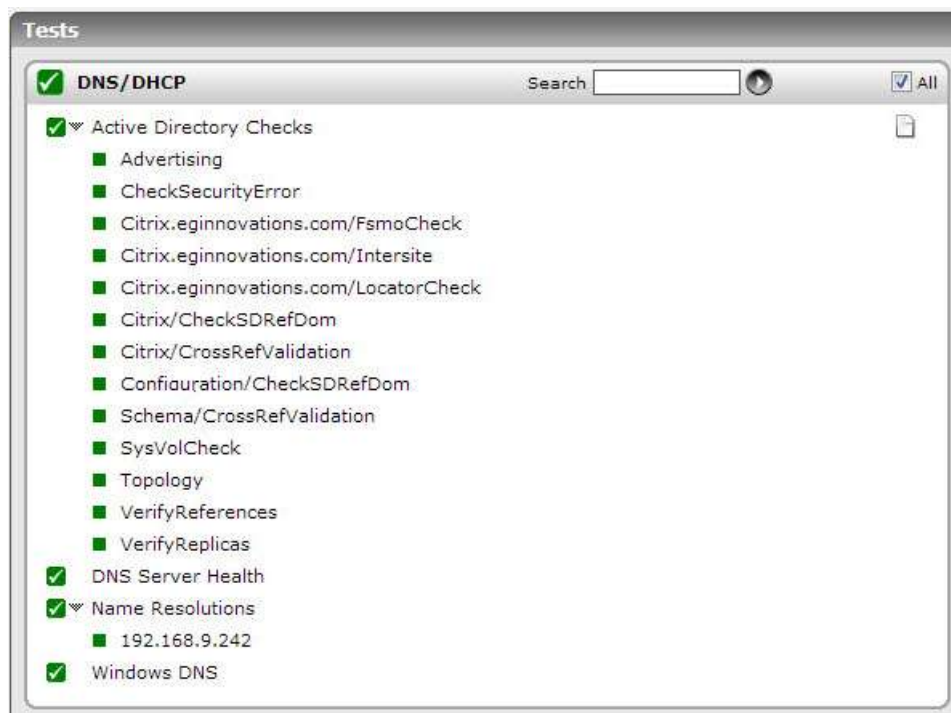


Figure 3.3: The tests mapped to the DNS/DHCP layer

3.3.1 Active Directory Checks Test

Domain controllers are the backbone of a Windows network. If your domain controllers are not working then the Active Directory does not work either. If the Active Directory does not work, then users cannot log on, group policies cannot be enforced, and a whole slew of other features become unavailable. To enable administrators to quickly detect and troubleshoot issues with the domain controller before they affect the operations of the AD server, Windows ships with a specialized tool called the Domain Controller Diagnostic (DCDIAG) Utility. DCDIAG is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in the system. The tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in troubleshooting. It consists of a framework for executing tests and a series of tests to verify different functional areas of the system - eg., replication errors, domain controller connectivity, permissions, proper roles, etc.

Using the **Active Directory Checks** test, the eG Enterprise Suite leverages the DCDIAG utility's ability to report on a wide variety of health parameters related to the domain controller. This ensures that even the smallest of aberrations in the performance of the domain controller is captured and promptly brought to the attention of the administrators. The **Active Directory Checks** test executes the DCDIAG command at configured intervals, and based on the output of the command, discovers the DCDIAG health checks that were performed, and the current status of each check - whether it reported a success or an error. In case the check resulted in an error/failure, you can use the detailed diagnosis of the test to understand the reason for the same, so that troubleshooting is easier!

Note:

For this test to run, the **DCDIAG.exe** should be available in the `<WINDOWS_INSTALL_DIR>\windows\system32` directory of the AD 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 system. 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 on any Windows 2003 (or higher) host in the environment to the same directory on the target host.

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Executes the DCDIAG command at configured intervals, and based on the output of the command, discovers the DCDIAG health checks that were performed, and the current status of each check - whether it reported a success or an error | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2003 or above | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. DOMAIN, USERNAME, PASSWORD, and CONFIRM PASSWORD - In order to execute the DCDIAG command, the eG agent has to be configured with a <i>domain administrator's</i> privileges. Therefore, specify the domain name and login credentials of the <i>domain administrator</i> in the DOMAIN, USERNAME and PASSWORD text boxes. Confirm the PASSWORD you provide by retyping it in the CONFIRM PASSWORD text box. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every DCDIAG health check that was performed | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| <p>test</p> | <p>Status: Indicates the status of this DCDIAG health check.</p> | <p>If the health check returns a positive result, the value of this measure will be <i>Pass</i>. If not, the value of this measure will be <i>Fail</i>. The numeric values that correspond to these measure values have been discussed in the table below:</p> <table border="1" data-bbox="933 401 1442 548"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> </tbody> </table> <p>Note: By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i>, you can use the detailed diagnosis of this measure to know the reason for the failure and the domain controller where the failure occurred. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 |
|--------------------|---|--|---------------|---------------|------|---|------|---|
| Measure Value | Numeric Value | | | | | | | |
| Pass | 1 | | | | | | | |
| Fail | 0 | | | | | | | |

3.3.2 AD Checks Test

Domain controllers are the backbone of a Windows network. If your domain controllers are not working then the Active Directory does not work either. If the Active Directory does not work, then users cannot log on, group policies cannot be enforced, and a whole slew of other features become unavailable. To enable administrators to quickly detect and troubleshoot issues with the domain controller before they affect the operations of the AD server, Windows ships with a specialized tool called the Domain Controller Diagnostic (**DCDIAG**) Utility. **DCDIAG** is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in the system. The tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in troubleshooting. It consists of a framework for executing tests and a series of tests to verify different functional areas of the system - eg., replication errors, domain controller connectivity, permissions, proper roles, etc.

Using the **AD Checks** test, the eG Enterprise Suite leverages the **DCDIAG** utility's ability to report on a wide variety of health parameters related to the domain controller. This ensures that even the smallest of aberrations in the performance of the domain controller is captured and promptly brought to the attention of the administrators. The **AD Checks** test executes the **DCDIAG** command at configured intervals, and based on the output of the command, reports the count of **DCDIAG** health checks (i.e., tests) that succeeded and failed in the last measurement period. The detailed diagnosis of the **AD Checks** test will provide detailed information pertaining to tests that failed, and thus assists in troubleshooting.

Note:

For this test to run, the **DCDIAG.exe** should be available in the <WINDOWS_INSTALL_DIR>\windows\system32 directory of the AD 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 system. 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 on any Windows 2003 (or higher) host in the environment to the same directory on the target host.

This test is disabled by default. To enable the test, follow the *Agents -> Tests -> Enable/Disable* menu sequence, pick **Active Directory** as the **Component type**, select **Performance** as the **Test type**, select this test from the **DISABLED TESTS** list and click the << button.

| | |
|---------------------------------|---|
| Purpose | Executes the DCDIAG command at configured intervals, and based on the output of the command, reports the count of DCDIAG health checks (i.e., tests) that succeeded and failed in the last measurement period |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 |
| Agent deploying the test | An internal agent |

| | | | |
|--|--|--------------------------------|---|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. DOMAIN, USERNAME, PASSWORD, and CONFIRM PASSWORD - In order to execute the DCDIAG command, the eG agent has to be configured with a <i>domain administrator's</i> privileges. Therefore, specify the domain name and login credentials of the <i>domain administrator</i> in the DOMAIN, USERNAME and PASSWORD text boxes. Confirm the PASSWORD you provide by retyping it in the CONFIRM PASSWORD text box. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| <p>Outputs of the test</p> | <p>One set of results for every Active Directory being monitored</p> | | |
| <p>Measurements made by the test</p> | <p>Measurement</p> | <p>Measurement Unit</p> | <p>Interpretation</p> |
| | <p>Passed tests: Indicates the number of DCDIAG tests that succeeded during the last measurement period.</p> | <p>Number</p> | |
| | <p>Failed tests: Indicates the number of DCDIAG tests that failed during the last measurement period.</p> | <p>Number</p> | <p>A non-zero value for this measure indicates the existence of one/more errors in the functioning of the domain controller. To know what these errors are, use the detailed diagnosis of this measure.</p> |

3.3.3 DNS Server Health Test

If the DNS component of the AD server is unable to provide domain name resolution services, then users may be denied access to their mission-critical servers managed by the AD server. Under such circumstances, you may want to quickly check what is stalling the operations of DNS, so that the source of the issue can be isolated and eliminated.

DCDIAG is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in the system. The tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in

MONITORING ACTIVE DIRECTORY SERVERS

troubleshooting. It consists of a framework for executing tests and a series of tests to verify different functional areas of the system.

DCDIAG also performs seven DNS-centric health checks to report on the overall DNS health of the domain controllers. To know the current status of each of these seven health checks, use the **DNS Server Health** test. The periodic health reports provided by the **DNS Server Health** test will enable administrators to proactively isolate potential DNS-related issues with their domain controllers, determine the reason for these issues, and work towards preventing them.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Reports the current status of the seven DNS-related health checks that DCDIAG performs on the domain controllers | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. DOMAIN, USERNAME, PASSWORD, and CONFIRM PASSWORD - In order to execute the DCDIAG command, the eG agent has to be configured with a <i>domain administrator's</i> privileges. Therefore, specify the domain name and login credentials of the <i>domain administrator</i> in the DOMAIN, USERNAME and PASSWORD text boxes. Confirm the PASSWORD you provide by retyping it in the CONFIRM PASSWORD text box. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every Active Directory server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| <p>test</p> | <p>Authentication:</p> <p>This test is run by default and checks the following:</p> <ul style="list-style-type: none"> • Are domain controllers registered in DNS? • Can they be pinged? • Do they have Lightweight Directory Access Protocol/Remote Procedure Call (LDAP/RPC)? <p>This measure reports the current status of the Authentication or Connectivity test.</p> | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="932 331 1438 533"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
|--------------------|--|---|---------------|---------------|------|---|------|---|---------|---|
| Measure Value | Numeric Value | | | | | | | | | |
| Pass | 1 | | | | | | | | | |
| Fail | 0 | | | | | | | | | |
| Warning | 2 | | | | | | | | | |

| | <p>Basic:</p> <p>The basic DNS test confirms the following:</p> <ul style="list-style-type: none"> a. Whether the DNS client, Netlogon, KDC, and DNS Server services are running and available on domain controllers tested by dcdiag b. Whether the DNS servers on all adapters are reachable. c. Whether A record of each domain controller is registered on at least one of the DNS servers configured on the client. d. If a domain controller is running the DNS Server service, whether the Active Directory domain zone and SOA record for the Active Directory domain zone are present. e. Whether the root (.) zone is present. <p>This measure reports the current status of the Basic test.</p> | <p>Number</p> | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="935 338 1438 533"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
|---------------|---|---------------|---|---------------|---------------|------|---|------|---|---------|---|
| Measure Value | Numeric Value | | | | | | | | | | |
| Pass | 1 | | | | | | | | | | |
| Fail | 0 | | | | | | | | | | |
| Warning | 2 | | | | | | | | | | |

| | <p>Forwarders:</p> <p>The forwarder test determines whether recursion is enabled. If forwarders or root hints are configured, the forwarder test confirms that all forwarders or root hints on the DNS server are functioning, and also confirms that the <code>_ldap._tcp.<Forest root domain></code> DC Locator record is resolved.</p> | | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="935 338 1438 533"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
|---------------|---|--|---|---------------|---------------|------|---|------|---|---------|---|
| Measure Value | Numeric Value | | | | | | | | | | |
| Pass | 1 | | | | | | | | | | |
| Fail | 0 | | | | | | | | | | |
| Warning | 2 | | | | | | | | | | |
| | <p>This measure reports the current status of the Forwarder test.</p> | | <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | | | | | | | | |
| | <p>Delegations:</p> <p>The delegation test confirms that the delegated name server is a functioning DNS Server. The delegation test checks for broken delegations by ensuring that all NS records in the Active Directory domain zone in which the target domain controller resides have corresponding glue A records.</p> <p>This measure reports the current status of the Delegation test.</p> | | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="935 1092 1438 1287"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
| Measure Value | Numeric Value | | | | | | | | | | |
| Pass | 1 | | | | | | | | | | |
| Fail | 0 | | | | | | | | | | |
| Warning | 2 | | | | | | | | | | |

| | <p>Dynamic update:</p> <p>The dynamic update test confirms that the Active Directory domain zone is configured for secure dynamic update and performs registration of a test record (_dcdiag_test_record).</p> <p>This measure reports the current status of the Dynamic Update test.</p> | | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="935 338 1438 533"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
|---------------|---|--|---|---------------|---------------|------|---|------|---|---------|---|
| Measure Value | Numeric Value | | | | | | | | | | |
| Pass | 1 | | | | | | | | | | |
| Fail | 0 | | | | | | | | | | |
| Warning | 2 | | | | | | | | | | |
| | <p>Record registration:</p> <p>The record registration test verifies the registration of all essential DC Locator records on all DNS Servers configured on each adapter of the domain controllers.</p> <p>This measure reports the current status of the Record Registration test.</p> | | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="935 1092 1438 1287"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
| Measure Value | Numeric Value | | | | | | | | | | |
| Pass | 1 | | | | | | | | | | |
| Fail | 0 | | | | | | | | | | |
| Warning | 2 | | | | | | | | | | |

| | <p>Resolve external name:</p> <p>The external name resolution test verifies basic resolution of external DNS from a given client, using a sample Internet name (www.microsoft.com), or user-provided Internet name.</p> <p>This measure reports the current status of the External name resolution test.</p> | <p>The values that this measure reports and their corresponding numeric values have been discussed in the table below:</p> <table border="1" data-bbox="933 333 1440 533"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Pass</td> <td>1</td> </tr> <tr> <td>Fail</td> <td>0</td> </tr> <tr> <td>Warning</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> <p>If the measure reports the value <i>Fail</i> or <i>Warning</i>, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</p> | Measure Value | Numeric Value | Pass | 1 | Fail | 0 | Warning | 2 |
|---------------|--|---|---------------|---------------|------|---|------|---|---------|---|
| Measure Value | Numeric Value | | | | | | | | | |
| Pass | 1 | | | | | | | | | |
| Fail | 0 | | | | | | | | | |
| Warning | 2 | | | | | | | | | |

3.3.4 Name Resolutions Test

Active Directory uses DNS as its domain controller location mechanism and leverages the namespace design of DNS in the design of Active Directory domain names. As a result, DNS is positioned within the discoverability and logical structure components of Active Directory technology components. If a user complains of being unable to access an AD domain, then administrators should first check whether the DNS component of AD is available and is able to resolve the IP address of the domain to its corresponding domain name and vice-versa. This is where, the **Name Resolutions** test will be useful!

This test emulates a client accessing DNS to issue a query. The query can either request DNS 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 component of the AD server.

| | |
|---------------------------------|---|
| Purpose | To measure the state of the DNS component of AD |
| Target of the test | An AD server |
| Agent deploying the test | An external agent |

| | | | |
|---|---|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT - The port on which the specified host is listening TARGETS - The IP address or host name to be resolved during the test. Multiple TARGETS can be specified as a comma-separated list. RECURSIVE - DNS supports two types of queries. For a non-recursive query, DNS 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 DNS. 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. | | |
| Outputs of the test | One set of results per TARGET configured | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | DNS availability: Whether a successful response is received from the DNS component of the target AD 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 DNS. |
| | DNS response time: Time taken (in seconds) by DNS 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 DNS, a network problem, etc. |

3.3.5 Windows DNS Test

This test measures the workload and processing ability of the DNS component of the AD server.

| | |
|---|--|
| Purpose | Measures the workload and processing ability of the DNS component of the AD server |
| Target of the test | An AD server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Windows server |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|--|--|
| Outputs of the test | One set of results for the AD server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total queries: The rate of queries received by DNS. | Reqs/sec | Indicates the workload of the DNS component of the AD server. |
| | Total responses: The rate of responses from DNS to clients. | Resp/sec | Ideally, the total responses should match the total queries. Significant differences between the two can indicate that DNS is not able to handle the current workload. |
| | Recursive queries: The rate of recursive queries successfully handled by DNS. | Reqs/sec | The ratio of recursive queries to total queries indicates the number of queries that required the DNS component on the AD 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 DNS. | 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. | |

| | | | |
|--|---|------|--|
| | <p>Zone transfers received: The number of zone transfer requests received by DNS.</p> | 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. |
| | <p>Zone transfers failed: The number of zone transfers that were not serviced by DNS in the last measurement period.</p> | 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. |

3.4 The AD Replication Service Layer

The tests mapped to this layer report on the health of the AD replication service.



Figure 3.4: The tests mapped to the AD Replication Service layer

3.4.1 File Replication Connections Test

This test reports metrics related to the file replication connections to Distributed File System roots (DFS) in an Active Directory.

| | |
|----------------------------|--|
| Purpose | Reports metrics related to the replica connections to Distributed File System roots (DFS) in an Active Directory |
| Target of the test | An Active Directory |
| Agent deploying the | An internal agent |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|---|
| test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Authentications: Indicates the number of successful authentications that were performed. | Number | |
| | Bindings: Indicates the number of successful RPC bindings that were completed. | Number | |
| | Joins: Indicates the number of joins. | Number | After FRS discovers a connection from Active Directory, FRS establishes a connection session with the remote connection partner based on the information provided by the connection object. The connection is called "joined" when a connection session is successfully established. |
| | Unjoins: Indicates the number of unjoins. | Number | |
| | Local change orders sent: Indicates the number of local change orders that were sent. | Number | A change order is a message that contains information about a file or folder that has changed on a replica. A local change order is a change order that is created because of a change to a file or folder on the local server. The local server becomes the originator of the change order and constructs a staging file – this file is nothing but a backup of the changed file or folder. |
| | Packets: Indicates the packets that were sent. | Number | |
| | Remote change orders sent: Indicates the number of remote change orders that were sent. | Number | A remote change order refers to a change order received from an inbound (or upstream) partner that originated elsewhere in the replica set. |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|--------|--|
| | Remote change orders received: Indicates the number of remote change orders that were received. | Number | |
|--|---|--------|--|

3.4.2 File Replication Events Test

This test reports statistical information about the File Replication Service events recorded in the File Replication Service event log. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Active Directory* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | |
|---------------------------------|--|
| Purpose | Reports statistical information about the File Replication Service events recorded in the File Replication Service event log |
| Target of the test | An Active Directory server |
| Agent deploying the test | An internal agent |

| | |
|--|--|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – 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 <i>application</i>. 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: <ul style="list-style-type: none"> • 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 <p>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.</p> 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: <i>{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}</i>. For example, assume that the FILTER text area takes the value, <i>OS_events:all:Browse,Print:all:none:all:none</i>. Here: <ul style="list-style-type: none"> • <i>OS_events</i> is the display name that will appear as a descriptor of the test in the monitor UI; • <i>all</i> 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 <i>none</i>. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, <i>Browse</i> and <i>Print</i> have been excluded from monitoring. Alternatively, you can use <i>all</i> to indicate that all the event sources have to be excluded from monitoring, or <i>none</i> 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 <i>all</i> 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.

| | | | |
|--------------------------------------|--|-------------------------|--|
| | <p>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.</p> <p>8. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY.</p> <p>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.</p> <p>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the FILTER configured | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>File replication errors:</p> <p>This refers to the number of File Replication Service events that were generated.</p> | Number | <p>A very low value (zero) indicates that the File Replication Service is in a healthy state without any potential problems.</p> <p>An increasing trend or high value indicates the existence of problems like loss of functionality or data.</p> <p>The detailed diagnosis capability, if enabled, lists the description of specific events.</p> <p>Please check the Application Logs in the Event Log Viewer for more details.</p> |
| | <p>File replication information count:</p> <p>This refers to the number of File Replication Service information events generated when the test was last executed.</p> | Number | <p>A change in the value of this measure may indicate infrequent but successful operations performed by the File Replication Service.</p> <p>The detailed diagnosis capability, if enabled, lists the description of specific events.</p> |

| | | | |
|--|---|--------|---|
| | <p>File replication warnings:</p> <p>This refers to the number of warnings that were generated when the test was last executed.</p> | Number | <p>A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems in the File Replication Service.</p> <p>The detailed diagnosis capability, if enabled, lists the description of specific events.</p> |
| | <p>File replication critical errors:</p> <p>Indicates the number of critical events that were generated when the test was last executed.</p> | Number | <p>This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems.</p> <p>A high value of this measure indicates that too many events have occurred, which the File Replication Service cannot automatically recover from.</p> <p>The detailed diagnosis capability, if enabled, provides the description of specific events.</p> |
| | <p>File replication verbose count:</p> <p>Indicates the number of verbose events that were generated when the test was last executed.</p> | Number | <p>This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems.</p> <p>Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better.</p> <p>The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.</p> |

3.4.3 File Replication Set Test

In the FRS, the replication of files and directories is according to a predefined topology and schedule on a specific folder. The topology and schedule are collectively called a replica set. A replica set contains a set of replicas, one for each machine that participates in replication.

This test reports statistics related to the health of the replication service provided by every replication set on an AD server.

| | |
|---|--|
| Purpose | Reports statistics related to the health of the replication service provided by every replication set on an AD server |
| Target of the test | An Active Directory |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|---|---|---|
| Outputs of the test | One set of results for every replication set on the Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Change orders received: Indicates the number of change orders that were currently received by this replica set. | Number | A change order is a message that contains information about a file or folder that has changed on a replica. These measures therefore serve as good indicators of the workload on the replica set. |
| | Change orders sent: Indicates the number of change orders that were currently sent by this replica set by this replica set. | Number | |
| Files installed: Indicates the number of file installations. | Number | Installation is the process by which FRS applies a change order to the local file system to restore the file or folder as it is in the upstream partner. If the change order is for a deletion, the file or folder in the local file system is deleted (staging file is not needed). If the change order is for a renaming, the file or folder in the local file system is renamed (staging file is needed). If the change order is for a copying or creation, the file or folder is copied or created (staging file is needed). Installing a file or folder may fail if the file or folder is already opened by another process. If the installation failed, FRS retries installing the file or folder at a later time. | |
| Packets received: Indicates the number of packets received currently. | Number | In an idle state, there should be no packets received unless a computer is having trouble joining with other computers in the replica set. | |
| Packets sent: Indicates the number of of packets sent currently. | Number | | |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|---------------|---|
| | <p>USN records accepted: Indicates the number of USN records that were currently accepted.</p> | <p>Number</p> | <p>Active Directory replication does not primarily depend on time to determine what changes need to be propagated. Instead it uses update sequence numbers (USNs) that are assigned by a counter that is local to each domain controller. Because these USN counters are local, it is easy to ensure that they are reliable and never run backward (that is, they cannot decrease in value).</p> <p>Domain controllers use USNs to simplify recovery after a failure. When a domain controller is restored following a failure, it queries its replication partners for changes with USNs greater than the USN of the last change it received from each partner.</p> |
| | <p>Staging space free: Indicates the staging space that is currently free.</p> | <p>KB</p> | <p>The Staging Directory is an area where modified files are stored temporarily either before being propagated to other replication partners or after being received from other replication partners. FRS encapsulates the data and attributes associated with a replicated file or directory object in a staging file. FRS needs adequate disk space for the staging area on both upstream and downstream machines in order to replicate files.</p> <p>Typically, if the Staging space free measure reports the value 0, or is found to be dangerously close to 0, it indicates that the staging directory is full. If the staging area is full, the FRS will stop functioning, and will resume only if disk space for the staging area becomes available or if the disk space limit for the staging area is increased.</p> <p>The staging area could get filled up owing to the following reasons:</p> <p>One or more downstream partners are not accepting changes. This could be a temporary condition due to the schedule being turned off and FRS waiting for it to open, or a permanent state because the service is turned off, or the downstream partner is in an error state.</p> <p>The rate of change in files exceeds the rate at which FRS can process them.</p> <p>A parent directory for files that have a large number of changes is failing to replicate, and so, all changes to subdirectories are blocked.</p> |
| | <p>Staging space in use: Indicates the staging space that is currently in use.</p> | <p>KB</p> | |

3.4.4 Replication Performance Test

Replication is the process by which the changes that are made on one domain controller are synchronized with all other domain controllers in the domain that store copies of the same information or replica.

Monitoring the replication operations on an AD server will shed light on the load generated by such operations and helps measure the ability of the AD server to process this load. The **Replication Performance** test does just that. In the process, the test points you to replication-related activities that could be contributing to processing delays (if any) and why. In addition, the test also promptly reports replication errors such as synchronization failures, and compels administrators to do what is necessary to ensure that no non-sync exists in the data that is replicated across the domain controllers in a forest.

This test applies only to Active Directory Servers installed on Windows 2008 or above.

| | | | |
|---|--|-------------------------|---|
| Purpose | Monitors the replication operations on an AD server and sheds light on the load generated by such operations and helps measure the ability of the AD server to process this load | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 or above | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | DRA inbound full sync objects remaining: Indicates the number of object updates received in the current directory replication update packet that have not yet been applied to the local server.. | Number | |
| | DRA inbound object updates remaining: Indicates the number of object updates received in the current directory replication update packet that have not yet been applied to the local server. | Number | The value of this measure should be low, with a higher value indicating that the hardware is incapable of adequately servicing replication (warranting a server upgrade). |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|-----------|---|
| | <p>Pending replication operations:</p> <p>Indicates the total number of replication operations on the directory that are queued for this server but not yet performed.</p> | Number | A steady increase in the value of this measure could indicate a processing bottleneck. |
| | <p>Pending replication synchronizations:</p> <p>Indicates the number of directory synchronizations that are queued for this server but not yet processed.</p> | Number | An unusually high value for a long duration may signify that the replication process is not being carried out at the desired rate. Forcing the replication activity may solve this problem. |
| | <p>Sync failures on schema mismatch:</p> <p>Indicates the number of synchronization requests made to neighbours that failed because their schema are not synchronized.</p> | Number | Ideally, the value of this measure should be 0. |
| | <p>Sync requests made:</p> <p>Indicates the number of synchronization requests made to neighbors.</p> | Number | |
| | <p>Sync requests successful:</p> <p>Indicates the number of synchronization requests made to neighbors that were successfully returned.</p> | Number | Ideally, the value of the <i>Sync requests made</i> measure should be equal to the value of the <i>Sync requests successful</i> measure - meaning, all sync request made should be successful, as one/more sync failures are a cause for concern. |
| | <p>DRA inbound objects applied rate:</p> <p>Indicates the rate at which replication updates received from replication partners are applied by the local directory service. This counter excludes changes that are received but not applied (because, for example, the change has already been made). This indicates how much replication update activity is occurring on the server as a result of changes generated on other servers.</p> | Appld/Sec | <p>A low value may indicate one of the following</p> <ul style="list-style-type: none"> • less changes to the objects in the other domains • this domain controller is not applying the changes to the objects at the desired rate. <p>If the object changes are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory. Forcing the replication activity may solve this problem.</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|----------------|---|
| | <p>DRA inbound properties applied rate:</p> <p>Indicates the number of properties that are updated due to the incoming property's winning the reconciliation logic that determines the final value to be replicated.</p> | Appld/Sec | <p>A low value may indicate one of the following less changes to the object properties in the other domains</p> <p>this domain controller is not applying the change to the object properties at the desired rate.</p> <p>If the object properties are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory. Forcing the replication activity may solve this problem.</p> |
| | <p>DRA inbound objects filtered rate:</p> <p>Indicates the number of objects received from inbound replication partners that contained no updates that needed to be applied.</p> | Filtrd/Sec | <p>A high value for this measure indicates that the objects are all static.</p> <p>This problem can be solved by increasing the replication frequency.</p> |
| | <p>DRA inbound properties filtered rate:</p> <p>Indicates the number of property changes (per second) already seen that were received during the replication.</p> | Filtrd/Sec | <p>A high value for this measure indicates that the properties are all static.</p> <p>This problem can be solved by increasing the replication frequency in the replicated domain.</p> |
| | <p>DRA inbound bytes total:</p> <p>Indicates the rate at which bytes were replicated in.</p> | Total/Sec | <p>This counter is the sum of the number of uncompressed bytes (never compressed) and the number of compressed bytes (after compression) per second.</p> |
| | <p>DRA outbound properties:</p> <p>Indicates the number of properties sent per second.</p> | Properties/Sec | <p>This counter tells you whether a source server is returning objects or not. Sometimes, the server might stop working correctly and not return objects quickly or at all.</p> |
| | <p>DRA outbound objects filtered rate:</p> <p>Indicates the number of objects per second that were determined by outbound replication to have no updates that the outbound partner did not already have.</p> | Filtrd/Sec | <p>A high value for this measure indicates that the objects are all static.</p> <p>This problem can be solved by increasing the replication frequency in the target domain.</p> |
| | <p>DRA outbound bytes total:</p> <p>Indicates the rate at which bytes were replicated out.</p> | Total/Sec | <p>This counter is the sum of the number of uncompressed bytes (never compressed) per second and the number of compressed bytes (after compression) per second.</p> |

3.4.5 Replication Traffic from Other Sites Test

Used in the Active Directory to express proximity of network connection, a **site** is defined as an IP subnetwork. A site consists of one or more subnets (unique network segments). Client machines use site information to find nearby DCs for logon operations. The Active Directory uses site information to help users find the closest machine that offers a needed network or a third-party service.

The Active Directory provides two methods of replication within the Active Directory environment: *intrasite replication* and *intersite replication*. *Intrasite replication* is replication within an Active Directory site. It is based assumption that the IP subnets within a site are well connected and that bandwidth is considered freely available and inexpensive. Because of this assumption, data is sent without compression.

Inter-site replication is replication between Active Directory sites. It is based on the assumption that the WAN is connected by slower links, so it is designed to minimize traffic rather than CPU cycles. Before being sent out, data is compressed to about 10% to 15% of original volume.

By monitoring the replication data flowing into each site, the **Replication Traffic from Other Sites** test helps determine the nature of the inbound traffic handled by every site - whether *inter-site* or *intrasite*, and reveals what type of inbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data compression is required at the source.

This test applies only to Active Directory Servers installed on Windows 2008 or above.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | By monitoring the replication data flowing into each site, the Replication Traffic from Other Sites test helps determine the nature of the inbound traffic handled by every site - whether <i>inter-site</i> or <i>intrasite</i> , and reveals what type of inbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data compression is required at the source. | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 or above | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory site being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| | | | |
|------|--|--------|--|
| test | <p>DRA inbound before bytes compression:</p> <p>Indicates the original size of inbound compressed replication data (kilobytes per second before compression, from DSAs in other sites).</p> | KB/Sec | |
| | <p>DRA inbound after bytes compression:</p> <p>Indicates the compressed size of inbound replication data (kilobytes per second received after compression, before DSAs in other sites).</p> | KB/Sec | <p>To save bandwidth on the network connection, the bridgehead servers in each site compress the traffic at the expense of additional CPU usage.</p> <p>A high value for this measure indicates that the bridgehead server is receiving high <i>inter-site</i> inbound replication traffic.</p> <p>Replication traffic is compressed down to about 40 percent when replication traffic is more than 32 KB in size.</p> |
| | <p>DRA inbound bytes not compression:</p> <p>Indicates the number of incoming bytes replicated per second that were not compressed at the source (that is, from DSAs in the same site).</p> | KB/Sec | <p>A high value for this measure indicates that the <i>intra-site</i> replication traffic is high.</p> <p>Compressing the replication data adds an additional load on the domain controller server. Uncompressed replication traffic preserves server performance at the expense of network utilization.</p> |

3.4.6 Replication Traffic to Other Sites Test

Used in the Active Directory to express proximity of network connection, a **site** is defined as an IP subnetwork. A site consists of one or more subnets (unique network segments). Client machines use site information to find nearby DCs for logon operations. The Active Directory uses site information to help users find the closest machine that offers a needed network or a third-party service.

The Active Directory provides two methods of replication within the Active Directory environment: *intrasite replication* and *intersite replication*. *Intrasite replication* is replication within an Active Directory site. It is based assumption that the IP subnets within a site are well connected and that bandwidth is considered freely available and inexpensive. Because of this assumption, data is sent without compression.

Inter-site replication is replication between Active Directory sites. It is based on the assumption that the WAN is connected by slower links, so it is designed to minimize traffic rather than CPU cycles. Before being sent out, data is compressed to about 10% to 15% of original volume.

By monitoring the replication data flowing from each site, the **Replication Traffic to Other Sites** test helps determine the nature of the outbound traffic handled by every site - whether *inter-site* or *intrasite*, and reveals what type of outbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data is to be compressed by the bridgehead server on each site.

This test applies only to Active Directory Servers installed on Windows 2008 or above.

| | | | |
|---|---|-------------------------|---|
| Purpose | By monitoring the replication data flowing from each site, the Replication Traffic to Other Sites test helps determine the nature of the outbound traffic handled by every site - whether <i>inter-site</i> or <i>intrasite</i> , and reveals what type of outbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data is to be compressed by the bridgehead server on each site. | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 or above | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory site being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | DRA outbound before bytes compression: Indicates the original size of outbound compressed replication data (kilobytes per second before compression, to DSAs in other sites). | KB/Sec | |
| | DRA outbound after bytes compression: Indicates the compressed size of outbound replication data (kilobytes per second sent after compression to DSAs in other sites). | KB/Sec | To save bandwidth on the network connection, the bridgehead servers in each site compress the traffic at the expense of additional CPU usage. A high value for this measure indicates that the bridgehead server is sending large high <i>inter-site</i> inbound replication traffic. Replication traffic is compressed down to about 40 percent when replication traffic is more than 32 KB in size. |

| | | | |
|--|---|--------|--|
| | <p>DRA outbound bytes not compression:</p> <p>Indicates the number of outgoing bytes replicated per second that were not compressed at the source (that is, to DSAs in the same site).</p> | KB/Sec | <p>A high value for this measure indicates that the <i>intra-site</i> replication traffic is high.</p> <p>Compressing the replication data adds an additional load on the domain controller server. Uncompressed replication traffic preserves server performance at the expense of network utilization.</p> |
|--|---|--------|--|

3.4.7 Replication Queue Test

As the domain controller formulates change requests, either by a schedule being reached or from a notification, it adds a work item for each request to the end of the queue of pending synchronization requests. Each pending synchronization request represents one <source domain controller, directory partition> pair, such as "synchronize the schema directory partition from DC1," or "delete the ApplicationX directory partition."

When a work item has been received into the queue, the domain controller processes the item (begins synchronizing from that source) as soon as the item reaches the front of the queue, and continues until either the destination is fully synchronized with the source domain controller, an error occurs, or the synchronization is pre-empted by a higher-priority operation.

A long replication queue is often an indication that synchronization requests are not swiftly processed by the AD server. If the reasons for the abnormal queue length are not determined quickly and addressed promptly, replication of some changes may be stalled indefinitely causing the source and destination domain controllers to remain 'out-of-sync' for long durations; this in turn may result in users having to work with obsolete data! To prevent such an eventuality, you can use this test to continuously track the replication queue length, so that you can be alerted as soon as the number of work items in the queue crosses an acceptable limit. You can also use the detailed diagnostics of this test to know what type of synchronization requests are in queue, so that you can figure out why the requests are taking too long to be processed.

| | |
|---------------------------------|---|
| Purpose | Continuously tracks the replication queue length, so that you can be alerted as soon as the number of work items in the queue crosses an acceptable limit |
| Target of the test | An Active Directory or Domain Controller |
| Agent deploying the test | An internal agent |

| | | | |
|---|--|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. 4. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every Active Directory site being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Replication queue size: Indicates the number of synchronization requests that are currently in the replication queue, awaiting processing. | Number | A high value for this measure is a cause for concern, as it indicates that too many synchronization requests are pending processing. This could be due to a severe processing bottleneck on the AD server. Very short replication schedules and large synchronization requests that require a lot of processing time are also factors that can increase the replication queue length. You can use the detailed diagnosis of this measure to know which requests are yet to be processed, so that you can figure out why there is a delay (if any) in processing. |

3.4.8 Lingering Objects Test

When restoring a backup file, Active Directory generally requires that the backup file be no more than 180 days old. If you attempt to restore a backup that has expired, you may encounter problems due to “lingering objects”.

A lingering object is a deleted AD object that re-appears (“lingers”) on the restored domain controller (DC) in its local copy of Active Directory. This can happen if, after the backup was made, the object was deleted on another DC more than 180 days ago.

When a DC deletes an object it replaces the object with a **tombstone** object. The tombstone object is a placeholder that represents the deleted object. When replication occurs, the tombstone object is transmitted to the other DCs, which causes them to delete the AD object as well.

Tombstone objects are kept for 180 days, after which they are garbage-collected and removed.

If a DC is restored from a backup that contains an object deleted elsewhere, the object will re-appear on the restored

MONITORING ACTIVE DIRECTORY SERVERS

DC. Because the tombstone object on the other DCs has been removed, the restored DC will not receive the tombstone object (via replication), and so it will never be notified of the deletion. The deleted object will “linger” in the restored local copy of Active Directory.

Such lingering objects tend to create problems during replication. For instance, if the source domain controller has outdated objects that have been out of replication for more than one tombstone lifetime a failure event will be logged in the Windows event log at the time of replicating from the source. You will have to promptly capture such events, identify the lingering objects, and delete them to ensure that replication resumes. In order to achieve this, you can use the **Lingering Objects** test. This test scans the event logs for replication events related to lingering objects, and promptly alerts you upon the occurrence of such events. Using the detailed diagnosis of the test, you can easily determine the location of the lingering objects, so that you can immediately proceed to remove them. This way, the test ensures that the replication engine operates without a glitch.

This test works only on Active Directory servers that operate on Windows 2008 or above.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Scans the event logs for replication events related to lingering objects, and promptly alerts you upon the occurrence of such events | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 or above | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every Active Directory server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| | | | |
|--------------------|---|---------------|---|
| <p>test</p> | <p>Lingering messages: Indicates the number of messages that are currently logged in the event log, which contains references to <i>lingering objects</i>.</p> | <p>Number</p> | <p>This measure typically captures and reports the number of events with event IDs 1388 and 1988 in the event log.</p> <p>Event ID 1388 indicates that a destination domain controller that does not have strict replication consistency enabled received a request to update an object that does not reside in the local copy of the Active Directory database. In response, the destination domain controller requested the full object from the source replication partner. In this way, a lingering object was replicated to the destination domain controller. Therefore, the lingering object was reintroduced into the directory.</p> <p>Event ID 1988 indicates that a destination domain controller that has strict replication consistency enabled has received a request to update an object that does not exist in its local copy of the Active Directory database. In response, the destination domain controller blocked replication of the directory partition containing that object from that source domain controller.</p> <p>The detailed diagnosis of this test provides the complete description of the events with IDs 1388 and/or 1988 that are logged in the event log. The source domain controller and the lingering objects can be inferred from the event description. Using this information, you can run the repadmin command on the source domain controller to delete the lingering objects.</p> |
|--------------------|---|---------------|---|

3.4.9 Replication Status Test

This test summarizes the replication state and relative health of an Active Directory forest by inventorying and contacting every domain controller in the forest, and collecting and reporting information such as replication deltas and replication failures. You can thus accurately identify the domain controllers that are prone to frequent failures.

| | |
|----------------------------------|--|
| <p>Purpose</p> | <p>Summarizes the replication state and relative health of an Active Directory forest by inventorying and contacting every domain controller in the forest, and collecting and reporting information such as replication deltas and replication failures</p> |
| <p>Target of the test</p> | <p>An Active Directory or Domain Controller</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|---|
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. 4. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every domain controller in an Active Directory forest being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total replication links: Indicates the number of replica links for this domain controller. | Number | A replica link exists for each naming context on a domain controller. This measure is the sum total of such replica links per domain controller. Please note that this is not the connection objects or replication partners per domain controller. You can use the detailed diagnosis of this measure to view the complete details of the replica links - this includes the source and destination sites, the source and destination domain controllers, the transport type, the number of link failures (if any), and details of the failures such as when the failure occurred and the failure status. |
| | Replication links failure: Indicates the total number of replica links on this domain controller that are failing to replicate for one reason or the other. This will never be greater than the <i>Total</i> field. | Number | Ideally, the value of this measure should be 0. |

| | | | |
|--|---|---------|---|
| | <p>Percent of replication links failure:</p> <p>Indicates the percentage of failures in relation to the total replica links on this domain controller.</p> | Percent | A low value is desired for this measure. A value close to 100% is a cause for concern, as it indicates that almost all replica links are failing. |
| | <p>Longest replication gap:</p> <p>Denotes the longest replication gap amongst all replication links on this domain controller.</p> | Secs | Ideally, this value should be less than 1 hour. |

3.4.10 Inter-Site Replication Test

Inter-site replication is based on the assumption that the WAN is connected by slower links or site links. It is designed to minimize traffic rather than CPU cycles. In inter-site replication, data is compressed and then sent out.

Bridgehead servers perform directory replication between sites. Only two designated domain controllers talk to each other. These domain controllers are called "Bridgehead servers".

After updates are replicated from one site to the bridgehead server in the other site, the updates are then replicated to other domain controllers within the site through intra-site replication process.

This test applies only to Active Directory Servers installed on Windows 2003.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | This test monitors the performance of the Active Directory Inter-site replication process in the target environment. | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2003 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| | | | |
|-------------|--|--------|--|
| test | In rate: This measure indicates the number of inbound kilobytes replicated between sites per second. | KB/Sec | A high value for this measure indicates that the bridgehead server is receiving high inter-site inbound replication traffic. |
| | Out rate: This measure indicates the number of outbound kilobytes replicated between sites per second. | KB/Sec | A high value indicates that bridgehead server is sending high inter-site outbound replication traffic. |

3.4.11 Intra-Site Replication Test

Intra-site replication means replication happening between domain controllers in the same site. Intra-site replication attempts to complete in the fewest CPU cycles possible. Intra-site replication avoids unnecessary network traffic by introducing a change notification mechanism that replaces the usual polling of replication partners for updates. When a change is performed in its database, a domain controller waits for a configurable interval (default 5 minutes) and accepts more changes during this time. Then it sends a notification to its replication partners, which will pull the changes from the source. If no changes are performed for a configurable period (default 6 hours) the domain controller initiates a replication sequence anyway, just to make sure that it did not miss anything.

This test applies only to Active Directory Servers installed on Windows 2003.

| | | | |
|---|--|-------------------------|--|
| Purpose | This test monitors the performance of the Active Directory Intra-site replication process in the target environment. | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2003 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | In rate: This measure indicates the number of inbound kilobytes replicated within the site per second. | KB/Sec | A high value for this measure indicates that the intra-site replication traffic is high. |

| | | | |
|--|---|--------|---|
| | <p>Out rate:</p> <p>This measure indicates the number of outbound kilobytes replicated within the site per second.</p> | KB/Sec | A high value for this measure indicates that the intra-site outbound replication traffic is high. |
|--|---|--------|---|

3.4.12 Replication Test

As the number of domain controllers increase, the replication process consumes more network bandwidth. So, replication process should be monitored within the target environment.

This test applies only to Active Directory Servers installed on Windows 2003.

| | | | |
|---|--|-------------------------|--|
| Purpose | This test monitors the performance of the Active Directory replication process in the target environment. | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2003 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>DRA inbound objects applied rate:</p> <p>This measure shows the number of replication updates applied per second that are occurring on this domain controller as a result of changes generated on other domain controllers.</p> | Appld/Sec | <p>A low value may indicate one of the following</p> <ol style="list-style-type: none"> less changes to the objects in the other domains this domain controller is not applying the changes to the objects at the desired rate. <p>If the object changes are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory.</p> <p>Forcing the replication activity may solve this problem.</p> |

| | | | |
|--|---|---------------|--|
| | <p>DRA inbound properties applied rate:</p> <p>This measure indicates the number of changes applied to object properties per second through inbound replication as a result of reconciliation logic. This logic is used to determine the final value to be replicated to the property.</p> | Appld/Sec | <p>A low value may indicate one of the following</p> <ol style="list-style-type: none"> 1. less changes to the object properties in the other domains 2. this domain controller is not applying the change to the object properties at the desired rate. <p>If the object properties are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory.</p> <p>Forcing the replication activity may solve this problem.</p> |
| | <p>DRA inbound objects filtered rate:</p> <p>This measure indicates the number of inbound replication objects received per second from the replication partners that contained no updates that needed to be applied.</p> | Filtrd/Sec | <p>A high value for this measure indicates that the objects are all static.</p> <p>Increasing the replication frequency may solve this problem.</p> |
| | <p>DRA inbound properties filtered rate:</p> <p>This measure indicates the number of inbound replication properties received per second from the replication partners that did not contain any updates to be applied.</p> | Filtrd/Sec | <p>A high value for this measure indicates that the properties are all static.</p> <p>Increasing the replication frequency in the replicated domain may solve this problem.</p> |
| | <p>DRA outbound objects filtered rate:</p> <p>This measure indicates the number of outbound replication objects that have not yet been received by the outbound replication partner per second.</p> | kerFiltrd/Sec | <p>A high value for this measure indicates that the objects are all static.</p> <p>Increasing the replication frequency in the target domain may solve this problem.</p> |
| | <p>Pending replication synchronizations:</p> <p>This measure indicates the number of directory synchronizations that are queued per second for this domain controller but not yet processed.</p> | Number | <p>An unusually high value for a long duration may signify that the replication process is not being carried out at the desired rate.</p> <p>Forcing the replication activity may solve this problem.</p> |

3.4.13 AD Replications Test

Replication is the process by which the changes that are made on one domain controller are synchronized with all other domain controllers in the domain that store copies of the same information or replica. Given the various types of information that Active Directory can store, changes to Active Directory can swiftly accumulate across multiple domain controllers in a large organization. It is therefore necessary for Windows to frequently synchronize the domain controllers through the replication process. If replication fails, it causes Active Directory objects that represent the replication topology, replication schedule, domain controllers, users, computers, passwords, security groups, group memberships, and Group Policy to be inconsistent between domain controllers. Directory inconsistency causes either operational failures or inconsistent results, depending on the domain controller that is contacted for the operation at hand.

To avoid such inconsistencies, it's best to capture failures promptly, isolate the source of failures, and fix them. The **AD Replications** test aids in this regard. This test closely monitors the replication activities on the domain controller and promptly reports replication failures, so that administrators can investigate such failures, discover the reasons for the same, fix them, and restore normalcy.

| | | | |
|---|--|-------------------------|---|
| Purpose | Closely monitors the replication activities on the domain controller and promptly reports replication failures, so that administrators can investigate such failures, discover the reasons for the same, fix them, and restore normalcy | | |
| Target of the test | An Active Directory or Domain Controller on Windows | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT - The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Replication failures: Indicates the number of replication failures in the target domain controller. | Number | Ideally, the value of this measure should be low. |
| | Total replications: Indicates the number of replication successes in the target domain controller. | Number | |

| | | | |
|--|--|----------------|--|
| | <p>Percent replication failures:</p> <p>Indicates the percentage of replication failures in the target domain controller.</p> | <p>Percent</p> | <p>Ideally, the value of this measure should be low. A high value is indicative of too many replication failures.</p> <p>Active Directory replication problems can have several different sources. For example, Domain Name System (DNS) problems, networking issues, or security problems can all cause Active Directory replication to fail.</p> <ul style="list-style-type: none"> • Network connectivity: The network connection might be unavailable or network settings are not configured properly. • Name resolution: DNS misconfigurations are a common cause for replication failures. • Authentication and authorization: Authentication and authorization problems cause "Access denied" errors when a domain controller tries to connect to its replication partner. • Directory database (store): The directory database might not be able to process transactions fast enough to keep up with replication timeouts. • Replication engine: If intersite replication schedules are too short, replication queues might be too large to process in the time that is required by the outbound replication schedule. In this case, replication of some changes can be stalled indefinitely — potentially, long enough to exceed the tombstone lifetime. • Replication topology: Domain controllers must have intersite links in Active Directory that map to real wide area network (WAN) or virtual private network (VPN) connections. If you create objects in Active Directory for the replication topology that are not supported by the actual site topology of your network, replication that requires the misconfigured topology fails. |
|--|--|----------------|--|

3.5 The AD Service Layer

This layer tracks the health of the Active Directory in a Windows environment using the ActiveDirectory test shown in Figure 3.5.

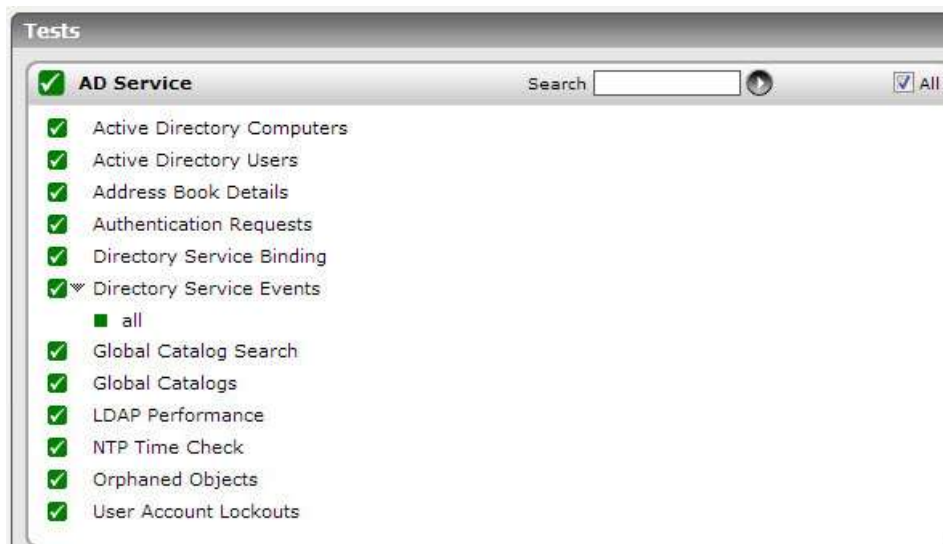


Figure 3.5: Tests mapping to the DC Service layer

3.5.1 Orphaned Objects Test

On a domain controller, the Lost and Found container contains Active Directory objects that have been orphaned. An object is orphaned when the object is created on one domain controller and the container in which the object is placed is deleted from the directory on another domain controller before the object has a chance to replicate. An orphaned object is automatically placed in the Lost and Found container where it can be found by an administrator, who must determine whether to move or delete the object.

The Orphaned Objects test periodically reports the number of orphaned objects on a domain controller.

| | |
|---------------------------------|--|
| Purpose | Periodically reports the number of orphaned objects on a domain controller |
| Target of the test | An AD server |
| Agent deploying the test | An internal agent |

| | | | |
|---|--|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows server 4. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every AD server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Orphaned objects: Indicates the number of objects in the Lost and Found container. | Number | If the value of this measure is greater than 0, it indicates the existence of orphaned objects. In such a case, you can use the detailed diagnosis capability of this measure to view the complete details of the objects, and accordingly decide whether to move the object or delete it. |

The detailed diagnosis of the *Orphaned objects* measure, if enabled, provides the complete details of the orphaned objects, which includes the named of the **Object class** and **Distinguished name**.

| Detailed Diagnosis measure of Lost and Found Container in Active Directory | | |
|--|--------------|--|
| Time | Object class | Distinguished name |
| 5:7:09 11/10/23 | lostAndFound | CN=LostAndFound,DC=TESTSUB2,DC=TESTMAIN,DC=COM |

Figure 3.6: The details of orphaned objects

3.5.2 Active Directory Status Test

This test tracks the performance of Active Directory existing in a Windows 2000 environment. Before getting into the details of this test, it is essential for the users to know that there are two choices for network authentication in a Windows 2000 environment. They are

1. **Kerberos Version 5.0:** This protocol is the default network authentication protocol for Windows 2000 servers.
2. **Windows NT LAN Manager (NTLM):** The NTLM protocol was the default network authentication protocol for Windows NT 4.0 operating system. NTLM is also used to authenticate logons to standalone computers with Windows 2000.

MONITORING ACTIVE DIRECTORY SERVERS

When a user first authenticates to Kerberos, he/she talks to the Authentication Service (AS) on the Kerberos Key Distribution Center (KDC) to get a Ticket Granting Ticket (TGT). This ticket is encrypted with the user's password. When the user wants to talk to a Kerberized service, he/she uses the Ticket Granting Ticket (TGT) to talk to the Ticket Granting Service (TGS), which also runs on the KDC. The Ticket Granting Service then verifies the user's identity using the TGT and issues a ticket for the desired service. The reason the Ticket Granting Ticket exists is that a user doesn't have to enter their password every time they wish to connect to a Kerberized service.

The outputs of the ActiveDirectoryStatus Test are given below:

| | | | |
|---|--|-------------------------|--|
| Purpose | This test monitors the performance of Active Directory in a Windows 2000 environment. | | |
| Target of the test | An Active Directory or Domain Controller | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Schema cache hit ratio:</p> <p>This measure shows the percentage of object name lookups available in the Schema Cache. This cache is present in the Domain Controller. All changes made to the Active Directory are first validated against this schema cache.</p> | Percent | A low value of this measure indicates that the Directory Service needs high disk read/write activity to perform its job. This results in poor response time of the components available in the Active Directory. |
| | <p>Notify queue size:</p> <p>When any change in the Active Directory occurs, the originating domain controller sends an update notification requests to the other domain controllers. This measure shows the number of pending update notification requests that have been queued and not transmitted.</p> | Number | A high value of this measure indicates that the Active Directory is changing frequently but the update notification requests have not been transmitted to the other domain controllers. This results in a loss of data integrity in the directory store. This problem can be corrected by forcing the replication process. |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|------------|---|
| | <p>Current threads:</p> <p>This measure shows the number of threads that are currently servicing the API calls by the users.</p> | Number | A fluctuating value for this measure indicates a change in the load. |
| | <p>Directory writes:</p> <p>This measure shows the number of successful write operations made by the directory service per second.</p> | Writes/Sec | A high value for this measure indicates that the directory service has made write operations in the Active Directory. This results in the fragmentation of the Active Directory. This problem can be corrected by forcing the replication process. |
| | <p>Kerberos requests:</p> <p>This measure shows the number of times per second that the user uses the user credentials to authenticate himself or herself with the domain controller that is being monitored.</p> | Reqs/Sec | <p>A high value for this measure indicates that the user requested some network resource, which requires authentication.</p> <p>Installing one or more Active Directory in the target environment can solve this problem</p> |
| | <p>NTLM requests:</p> <p>This measure shows the number of times per second that the user uses the user credentials to authenticate himself or herself with the domain controller, which is having the PDC emulator operation role.</p> | Reqs/Sec | <p>A high value for this measure indicates that the user requested some network resource, which basically belongs to the Windows NT network. Accessing this kind of resource needs authentication, which is serviced by the domain controller, who is having the PDC emulator operation role.</p> <p>Installing one or more domain controllers with PDC emulator operation role in the target environment can solve this problem.</p> |
| | <p>Ticket requests:</p> <p>This measure indicates the number of requests made by the Ticket Granting Service per second.</p> | Reqs/Sec | <p>A high value for this measure indicates that the user requested some network resources, which needs authentication.</p> <p>Installing one or more domain controllers in the target environment can solve this problem.</p> |
| | <p>Authentication requests:</p> <p>This measure indicates the number of requests made by the Authentication Server (to obtain the TGT) per second.</p> | Reqs/Sec | <p>A high value for this measure indicates that the user requested some network resources, which needs authentication.</p> <p>Installing one or more domain controllers in the target environment can solve this problem.</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|--------|--|
| | Ldap sessions: This measure indicates the number of Ldap clients currently connected to the Active Directory. | Number | This measure is just an indicator of the number of Ldap clients connected to the Active Directory. A high or low value for this measure does not always denote an error situation. |
|--|---|--------|--|

3.5.3 Directory Service Events Test

This test reports statistical information about the Directory Service events recorded in the event log. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Active Directory* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | |
|---------------------------------|--|
| Purpose | Reports statistical information about the Directory Service events recorded in the event log |
| Target of the test | An Active Directory server |
| Agent deploying the test | An internal agent |

| | |
|---|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – 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 <i>application</i>. 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: <ul style="list-style-type: none"> ➤ 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 <p>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.</p> 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: <i>{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}</i>. For example, assume that the FILTER text area takes the value, <i>OS_events:all:Browse,Print:all:none:all:none</i>. Here: <ul style="list-style-type: none"> • <i>OS_events</i> is the display name that will appear as a descriptor of the test in the monitor UI; • <i>all</i> 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 <i>none</i>. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, <i>Browse</i> and <i>Print</i> have been excluded from monitoring. Alternatively, you can use <i>all</i> to indicate that all the event sources have to be excluded from monitoring, or <i>none</i> 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 <i>all</i> 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.

| | | | |
|--------------------------------------|--|-------------------------|---|
| | <p>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.</p> <p>8. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY.</p> <p>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.</p> <p>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the FILTER configured | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Directory service errors:</p> <p>This refers to the number of Directory Service events that were generated.</p> | Number | <p>A very low value (zero) indicates that the Directory Service is in a healthy state without any potential problems.</p> <p>An increasing trend or high value indicates the existence of problems like loss of functionality or data.</p> <p>The detailed diagnosis capability, if enabled, lists the description of specific events.</p> <p>Please check the Application Logs in the Event Log Viewer for more details.</p> |
| | <p>Directory service information count:</p> <p>This refers to the number of Directory Service Service information events generated when the test was last executed.</p> | Number | <p>A change in the value of this measure may indicate infrequent but successful operations performed by the Directory Service.</p> <p>The detailed diagnosis capability, if enabled, lists the description of specific events.</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|--|--------|---|
| | <p>Directory service warnings:</p> <p>This refers to the number of warnings that were generated when the test was last executed.</p> | Number | <p>A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems in the Directory Service.</p> <p>The detailed diagnosis capability, if enabled, lists the description of specific events.</p> |
| | <p>Directory service critical errors:</p> <p>Indicates the number of critical events that were generated when the test was last executed.</p> | Number | <p>This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems.</p> <p>A high value of this measure indicates that too many errors have occurred, which the Directory Service cannot automatically recover from.</p> <p>The detailed diagnosis capability, if enabled, provides the description of specific events.</p> |
| | <p>Directory service verbose count:</p> <p>Indicates the number of verbose events that were generated when the test was last executed.</p> | Number | <p>This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems.</p> <p>Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better.</p> <p>The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.</p> |

3.5.4 User Account Lockouts Test

Account lockout is a feature of password security that disables a user account when a certain number of failed logons occur due to wrong passwords within a certain interval of time. The purpose behind account lockout is to prevent attackers from brute-force attempts to guess a user's password.

Other ways accounts can get locked out include:

- Applications using cached credentials that are stale.
- Stale service account passwords cached by the Service Control Manager (SCM).
- Stale logon credentials cached by Stored User Names and Passwords in Control Panel.
- Scheduled tasks and persistent drive mappings that have stale credentials.
- Disconnected Terminal Service sessions that use stale credentials.
- Failure of Active Directory replication between domain controllers.
- Users logging into two or more computers at once and changing their password on one of them.

Any one of the above situations can trigger an account lockout condition, and the results can include applications behaving unpredictably and services inexplicably failing.

This is why, whenever a user complains of inability to login to his/her desktop, help desk should be able to instantly figure out whether that user's account has been locked out, and if so, why. The **User Account Lockouts** test provides answers to these questions. This test, at configured intervals, reports the count of locked user accounts and names the users who have been affected by this anomaly.

| | |
|---------------------------------|--|
| Purpose | Reports the count of locked user accounts and names the users who have been affected by this anomaly |
| Target of the test | An Active Directory |
| Agent deploying the test | An internal agent; this test cannot be run in an 'agentless' manner |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|---|--|-------------------------|-----------------------|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. 4. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| test | <p>Account lockout events:</p> <p>Indicates the number of account lockouts that occurred during the last measurement period.</p> | Number | <p>A very high value for this measure could indicate a malicious attack, and may require further investigation.</p> <p>If the high lockout rate is not due to any such attacks, then it is recommended that you alter the lockout policy in your environment to minimize the count and consequently, the impact of account lockouts. Microsoft recommends the following policies for high, medium, and low security environments:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Security Level</th> <th style="text-align: left;">Lockout Policy</th> </tr> </thead> <tbody> <tr> <td style="vertical-align: top;">Low</td> <td> Account Lockout Duration = Not Defined Account Lockout Threshold = 0 (No lockout) Reset account lockout counter after = Not Defined </td> </tr> <tr> <td style="vertical-align: top;">Medium</td> <td> Account Lockout Duration = 30 minutes Account Lockout Threshold = 10 invalid logon attempts Reset account lockout counter after = 30 minutes </td> </tr> <tr> <td style="vertical-align: top;">High</td> <td> Account lockout duration = 0 (an administrator must unlock the account) Account lockout threshold = 10 invalid logon attempts Reset account lockout counter after = 30 minutes </td> </tr> </tbody> </table> | Security Level | Lockout Policy | Low | Account Lockout Duration = Not Defined Account Lockout Threshold = 0 (No lockout) Reset account lockout counter after = Not Defined | Medium | Account Lockout Duration = 30 minutes Account Lockout Threshold = 10 invalid logon attempts Reset account lockout counter after = 30 minutes | High | Account lockout duration = 0 (an administrator must unlock the account) Account lockout threshold = 10 invalid logon attempts Reset account lockout counter after = 30 minutes |
|--|--|---|---|----------------|----------------|------------|---|---------------|--|-------------|--|
| | Security Level | Lockout Policy | | | | | | | | | |
| | Low | Account Lockout Duration = Not Defined Account Lockout Threshold = 0 (No lockout) Reset account lockout counter after = Not Defined | | | | | | | | | |
| Medium | Account Lockout Duration = 30 minutes Account Lockout Threshold = 10 invalid logon attempts Reset account lockout counter after = 30 minutes | | | | | | | | | | |
| High | Account lockout duration = 0 (an administrator must unlock the account) Account lockout threshold = 10 invalid logon attempts Reset account lockout counter after = 30 minutes | | | | | | | | | | |
| <p>Unique users locked out:</p> <p>Indicates the number of distinct users who were locked out during the last measurement period.</p> | Number | Use the detailed diagnosis of this measure to view the names of these users. | | | | | | | | | |
| <p>Users currently locked out:</p> <p>Indicates the number of users who are currently locked out.</p> | Number | Use the detailed diagnosis of this measure to know which users are currently locked out. | | | | | | | | | |

3.5.5 Active Directory Lost and Found Test

On a domain controller, the Lost and Found container contains Active Directory objects that have been orphaned. An object is orphaned when the object is created on one domain controller and the container in which the object is placed is deleted from the directory on another domain controller before the object has a chance to replicate. An orphaned object is automatically placed in the Lost and Found container where it can be found by an administrator.

This test reports the number of orphaned objects currently in the Lost and Found container, provides the details of these objects, so that administrators can determine which objects to move and which ones to delete.

This test applies only to Active Directory Servers installed on Windows 2008.

This test is disabled by default. To enable the test, follow the *Agents -> Tests -> Enable/Disable* menu sequence, pick **Active Directory** as the **Component type**, select **Performance** as the **Test type**, select this test from the **DISABLED TESTS** list and click the << button.

| | | | |
|---|---|-------------------------|--|
| Purpose | Reports the number of orphaned objects currently in the Lost and Found container, provides the details of these objects, so that administrators can determine which objects to move and which ones to delete | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. 4. 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: <ul style="list-style-type: none"> ➤ The eG manager license should allow the detailed diagnosis capability <p>f. Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</p> | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Lost and Found objects: Indicates the number of objects currently available in the Lost and Found container. | Number | A non-zero value indicates the existence of orphaned objects. Use the detailed diagnosis of this measure to know which objects to move and which ones to delete. |

3.5.6 Global Catalog Search Test

The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain Active Directory Domain Services (AD DS) forest. The global catalog is stored on domain controllers that have been designated as global catalog servers and is distributed through multimaster replication.

The Global Catalog enables searching for Active Directory objects in any domain in the forest without the need for subordinate referrals, and users can find objects of interest quickly without having to know what domain holds the object. The global catalog makes the directory structure within a forest transparent to users who perform a search. For example, if you search for all printers in a forest, a global catalog server processes the query in the global catalog and then returns the results. Without a global catalog server, this query would require a search of every domain in the forest.

This test reveals whether the server being monitored is a global catalog server or not. If it is, then the test attempts to search the global catalog server for a configured user and reports whether that user was found or not. The test also reports the time taken to search for that user. This information helps administrators assess how efficient the global catalog is in minimizing the time taken to locate a user across domains.

This test applies only to Active Directory Servers installed on Windows 2008.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Reveals whether the server being monitored is a global catalog server or not. If it is, then the test attempts to search the global catalog server for a configured user and reports whether that user was found or not. The test also reports the time taken to search for that user. This information helps administrators assess how efficient the global catalog is in minimizing the time taken to locate a user across domains | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. 4. USERNAME - Specify the name of the user who has to be searched in the global catalog. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

| | | | |
|-------------|--|---------|---|
| test | Is it a global catalog server?: Indicates whether the monitored server is a global catalog server or not. | Boolean | This measure reports the value <i>True</i> if the AD server being monitored is a global catalog server, and the value <i>False</i> if it is not. If this measure reports the value <i>False</i> , the remaining measures of the test will not report any values. |
| | Was user found? Indicates whether the configured USERNAME was found or not in the global catalog server. | Boolean | This measure reports the value <i>True</i> if the configured USERNAME was found in the global catalog server and the value <i>False</i> if the user name was not found. |
| | Catalog search time: Indicates the time taken by the global catalog server to search and find the configured USERNAME . | Secs | A high value for this measure would warrant an investigation. |

3.5.7 Address Book Details Test

The Address Book is a client for the Active Directory database. It performs lookups and search operations on the Active Directory database to look for details such as account email ID, and so forth. Using the **Address Book Details** test, you can determine the number of Address Book clients currently connected to the AD database and the rate at which search operations are performed by each AD server. In the event that the AD database gets inundated with search queries, you can use this test to figure out whether or not the Address Book clients are contributing to the query load.

This test applies only to Active Directory Servers installed on Windows 2008.

| | |
|---|---|
| Purpose | You can determine the number of Address Book clients currently connected to the AD database and the rate at which search operations are performed by each AD server. In the event that the AD database gets inundated with search queries, you can use this test to figure out whether or not the Address Book clients are contributing to the query load |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. |

| | | | |
|--------------------------------------|---|-------------------------|--|
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Client sessions: Indicates the number of client sessions that are currently connected to the AD database. | Number | A high value is indicative of heavy load. A consistent increase in the value of this measure could indicate a potential overload condition. |
| | Search operations: Indicate the rate at which the key search operations are performed on the AD database. | Searches/Sec | If the value of this measure decreases while the number of Client sessions keeps increasing, it indicates that search queries are not being processed as quickly; this in turn is indicative of a processing bottleneck, which can consequently choke the AD server database. |

3.5.8 ADAM LDAP Performance Test

The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that runs on a layer above the TCP/IP stack. It provides a mechanism used to connect to, search, and modify Internet directories. The LDAP directory service is based on a client-server model. The function of LDAP is to enable access to an existing directory. LDAP is one of the protocols used to query and modify items on the Active Directory server.

To monitor the interactions between clients and the AD server over LDAP, and to promptly capture slowdowns in LDAP searches and binds, use the **ADAM LDAP Performance** test.

This test applies only to Active Directory Servers installed on Windows 2008.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | To monitor the interactions between clients and the AD server over LDAP, and to promptly capture slowdowns in LDAP searches and binds | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT - The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |
| | | | |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|------|--|-----------------|--|
| test | <p>Ldap searches:</p> <p>Indicates the rate at which LDAP clients perform search operations.</p> | Searches/Sec | This counter should show activity over time. If it does not, network problems are probably hindering the processing of client requests. |
| | <p>Ldap writes:</p> <p>Indicates the rate at which clients perform write operations on the AD server.</p> | Writes/Sec | |
| | <p>Ldap active threads:</p> <p>Indicates the current number of threads in use by the LDAP subsystem of the local directory service.</p> | Number | A high number indicates a high level of LDAP activity on the directory service. |
| | <p>Ldap bind time:</p> <p>Indicates the time, in milliseconds, taken for the last successful LDAP bind.</p> | Secs | <p>In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as <i>binding</i>. When a programmatic object, such as an IADs or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be <i>bound to</i> the directory object.</p> <p>This measure should be as low as possible. If it is not, hardware or network-related problems are indicated.</p> |
| | <p>Ldap sessions:</p> <p>Indicates the number of currently connected LDAP client sessions.</p> | Number | This measure is just an indicator of the number of Ldap clients connected to the Active Directory. A high or low value for this measure does not always denote an error situation. |
| | <p>Ldap closed connections:</p> <p>Indicates the LDAP connections that have been closed in the last second.</p> | Connections/Sec | |
| | <p>Ldap new connections:</p> <p>Indicates the number of new LDAP connections that have arrived in the last second.</p> | Connections/Sec | |
| | <p>Ldap new ssl connections:</p> <p>Indicates the number of new SSL or TLS connections that arrived in the last second.</p> | Connections/Sec | |

| | | | |
|--|---|-----------|---|
| | <p>Ldap successful binds:</p> <p>Indicates the number of successful LDAP binds per second.</p> | Binds/Sec | <p>In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as <i>binding</i>. When a programmatic object, such as an IADs or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be <i>bound</i> to the directory object.</p> <p>A high value is desired for this measure. A very low value could indicate network problems.</p> |
|--|---|-----------|---|

3.5.9 Authentication Performance Test

Authentication of domain user logins is a core function of an Active Directory server. The default authentication protocol used by the AD server is **Kerberos**. Kerberos authentication is based on specially formatted data packets known as tickets. In Kerberos, these tickets pass through the network instead of passwords. Transmitting tickets instead of passwords makes the authentication process more resistant to attackers who can intercept the network traffic.

In a Kerberos environment, the authentication process begins at logon. The following steps describe the Kerberos authentication process:

1. When a user enters a user name and password, the computer sends the user name to the KDC (Key Distribution Center). The Key Distribution Center (KDC) maintains a database of account information for all security principals in the domain. The KDC stores a cryptographic key known only to the security principal and the KDC. This key is used in exchanges between the security principal and the KDC and is known as a long term key. The long term key is derived from a user's logon password.
2. Upon the receipt of a user name, the KDC looks up the user's master key (KA), which is based on the user's password. The KDC then creates two items: a session key (SA) to share with the user and a Ticket-Granting Ticket (TGT). The TGT includes a second copy of the SA, the user name, and an expiration time. The KDC encrypts this ticket by using its own master key (KKDC), which only the KDC knows.
3. The client computer receives the information from the KDC and runs the user's password through a one-way hashing function, which converts the password into the user's KA (i.e., master key). The client computer now has a session key and a TGT so that it can securely communicate with the KDC. The client is now authenticated to the domain and is ready to access other resources in the domain by using the Kerberos protocol.
3. When a Kerberos client needs to access resources on a server that is a member of the same domain, it contacts the KDC. The client will present its TGT and a timestamp encrypted with the session key that is already shared with the KDC. The KDC decrypts the TGT using its KKDC. The TGT contains the user name and a copy of the SA. The KDC uses the SA to decrypt the timestamp. The KDC can confirm that this request actually comes from the user because only the user can use the SA.
4. Next, the KDC creates a pair of tickets, one for the client and one for the server on which the client needs to access resources. Each ticket contains the name of the user requesting the service, the recipient of the request, a timestamp that declares when the ticket was created, and a time duration that says how long the tickets are valid. Both tickets also contain a new key (KAB) that will be shared between the client and the server so they can securely communicate.
5. The KDC takes the server's ticket and encrypts it using the server master key (KB). Then the KDC nests the

MONITORING ACTIVE DIRECTORY SERVERS

server's ticket inside the client's ticket, which also contains the KAB. The KDC encrypts the whole thing using the session key that it shares with the user from the logon process. The KDC then sends all the information to the user.

- When the user receives the ticket, the user decrypts it using the SA. This exposes the KAB to the client and also exposes the server's ticket. The user cannot read the server's ticket. The user will encrypt the timestamp by using the KAB and send the timestamp and the server's ticket to the server on which the client wants to access resources. When it receives these two items, the server first decrypts its own ticket by using its KB. This permits access to the KAB, which can then decrypt the timestamp from the client.

In situations where a domain controller is not available or is unreachable, **NTLM (the NT LAN Manager)** is used as the authentication protocol. For example, NTLM would be used if a client is not Kerberos capable, the server is not joined to a domain, or the user is remotely authenticating over the web.

In some other environments **Digest** authentication is supported. Digest authentication offers the same functionality as Basic authentication; however, Digest authentication provides a security improvement because a user's credentials are not sent across the network in plaintext. Digest authentication sends credentials across the network as a Message Digest 5 (MD5) hash, which is also known as the MD5 message digest, in which the credentials cannot be deciphered from the hash.

Regardless of the protocol/authentication mode used, the quality of a user's experience with the AD server largely relies on how fast his/her login is authenticated by the AD server. The slightest of delays will hence not be tolerated! Administrators therefore need to keep their eyes open at all times for authentication-related latencies, isolate their source, and fix the problems, so that users are able to login to their systems quickly. The **Authentication Performance** test helps administrators in this regard.

This test reports the rate at which Kerberos, NTLM, and Digest authentication requests are serviced by the AD server and thus promptly reveals delays in authentication (if any). Where latencies are noticed in Kerberos requests, the test goes one step further and indicates the probable source of the latencies - could it be because the KDC took too long to grant TGTs to the clients? or is it because the KDC took too long to process the TGTs and grant the clients access to authorized resources?

This test applies only to Active Directory Servers installed on Windows 2008.

| | |
|---|---|
| Purpose | Reports the rate at which Kerberos and NTLM authentication requests are serviced by the AD server and thus promptly reveals delays in authentication (if any). Where latencies are noticed in Kerberos requests, the test goes one step further and indicates the probable source of the latencies - could it be because the KDC took too long to grant TGTs to the clients? or is it because the KDC took too long to process the TGTs and grant the clients access to authorized resources? |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none">TEST PERIOD - How often should the test be executedHOST - The IP address of the machine where the Active Directory is installed.PORT - The port number through which the Active Directory communicates. The default port number is 389. |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--------------------------------------|---|-------------------------|---|
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Kerberos requests: Indicates the number of times per second that clients use a ticket to authenticate to the domain controller. | Reqs/Sec | A low value indicates a bottleneck when processing Kerberos requests. |
| | Digest requests: Indicates the rate at which requests from a potential user were received by a network server and then sent to a domain controller. | Reqs/Sec | A low value indicates a bottleneck when processing Digest requests. |
| | Ntlm requests: Indicates the rate at which NTLM authentication requests were serviced by the domain controller. | Reqs/Sec | A low value indicates a bottleneck when processing NTLM requests. A high value for this measure indicates that the user requested some network resource, which basically belongs to the Windows NT network. Accessing this kind of resource needs authentication, which is serviced by the domain controller, who is having the PDC emulator operation role. Installing one or more domain controllers with PDC emulator operation role in the target environment can solve this problem. |
| | Authentication requests: Indicates the number of Authentication Server (AS) requests serviced by the Kerberos Key Distribution Center (KDC) per second. | Reqs/Sec | AS requests are used by the client to obtain a ticket-granting ticket. If the AD server appears to be taking too long to process Kerberos requests - i.e., if the value of the <i>Kerberos requests</i> measure is too high - then you can compare the value of this measure with that of the <i>Ticket requests</i> measure to know where the request spent too much time - when granting TGTs to clients? or when processing the TGTs to allow users access to a resource? |

| | | | |
|--|---|----------|---|
| | <p>Ticket requests:</p> <p>Indicates the number of Ticket Granting Server (TGS) requests serviced by the KDC per second.</p> | Reqs/Sec | <p>TGS requests are used by the client to obtain a ticket to a resource.</p> <p>If the AD server appears to be taking too long to process Kerberos requests - i.e., if the value of the <i>Kerberos requests</i> measure is too high - then you can compare the value of this measure with that of the <i>Authentication requests</i> measure to know where the request spent too much time - when granting TGTs to clients? or when processing the TGTs to allow users access to a resource?</p> |
|--|---|----------|---|

3.5.10 ADAM Binding Test

In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as *binding*. When a programmatic object, such as an IADs or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be *bound to* the directory object.

This test reports the type of binds that exist in an AD environment, and for each bind type, reports how fast the AD server bound the programmatic objects to the directory object.

This test applies only to Active Directory Servers installed on Windows 2008.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the type of binds that exist in an AD environment, and for each bind type, reports how fast the AD server bound the programmatic object to the directory object | | |
| Target of the test | An Active Directory or Domain Controller on Windows 2008 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT - The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Ntlm binds:</p> <p>Indicates the rate at which programmatic and directory objects were bound to one another using <i>NTLM binds</i>.</p> | Binds/Sec | |

| | | | |
|--|--|-----------|--|
| | <p>Simple binds:</p> <p>Indicates the rate at which programmatic and directory objects were bound to one another using <i>Simple binds</i>.</p> | Binds/Sec | In a simple bind, the client either binds anonymously, that is, with an empty bind Distinguished Name, or by providing a Distinguished Name and a password. |
| | <p>External binds:</p> <p>Indicates the rate at which programmatic and directory objects were bound to one another using <i>External binds</i>.</p> | Binds/Sec | |
| | <p>Fast binds:</p> <p>Indicates the rate at which programmatic and directory objects were bound to one another using <i>Fast binds</i>.</p> | Binds/Sec | Fast bind mode allows a client to use the LDAP bind request to simply validate credentials and authenticate the client without the overhead of establishing the authorization information. |
| | <p>Negotiated binds:</p> <p>Indicates the rate at which programmatic and directory objects were bound to one another using <i>Negotiated binds</i>.</p> | Binds/Sec | |

3.5.11 Global Catalogs Test

The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain active directory domain services (AD DS) forest. The global catalog is stored on domain controllers that have been designated as global catalog servers and is distributed through multimaster replication.

The global catalog enables searching for active directory objects in any domain in the forest without the need for subordinate referrals, and users can find objects of interest quickly without having to know what domain holds the object. The global catalog makes the directory structure within a forest transparent to users who perform a search. For example, if you search for all printers in a forest, a global catalog server processes the query in the global catalog and then returns the results. Without a global catalog server, this query would require a search of every domain in the forest.

This test monitors the global catalogs in the target domain controller and reports the number of catalogs that are currently available and unavailable. This way, the test enables administrators to determine whether/not adequate global catalogs are available in the domain controller to handle the request load.

| | |
|---------------------------------|--|
| Purpose | Monitors the global catalogs in the target domain controller and reports the number of catalogs that are currently available and unavailable |
| Target of the test | An Active Directory or Domain Controller on Windows |
| Agent deploying the test | An internal agent |

| | | | |
|---|--|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total global catalogs: Indicates the total number of global catalogs on the domain controller being monitored. | Number | |
| | Available global catalogs: Indicates the number of global catalogs that are currently available on the domain controller. | Number | |
| | Unavailable global catalogs: Indicates the number of global catalogs that are currently unavailable on the domain controller. | Number | If the value of this measure is equal to the value of the <i>Total global gatalogs</i> measure or is higher than that of the <i>Available global catalogs</i> measure, it indicates that enough global catalogs may not be available on the domain controller to process user logon requests and search requests. As a result, requests may fail. |
| | Percent unavailable global catalogs: Indicates percentage of global catalogs that are currently unavailable. | Percent | A high value indicates that too many global catalogs are unavailable for request processing. This in turn can cause many user logon and search requests to the domain controller to fail. Ideally therefore, the value of this measure should be very low. |

3.5.12 Active Directory Users

This test reports the status of user accounts configured in the Active Directory server and thus, quickly points you to 'unused' accounts that can be deleted to make room for those that are actively used.

| | |
|---------------------------------|---|
| Purpose | Reports the status of user accounts configured in the Active Directory server and thus, quickly points you to 'unused' accounts that can be deleted to make room for those that are actively used |
| Target of the test | An Active Directory or Domain Controller on Windows |
| Agent deploying the test | An internal agent |

| | | | |
|---|--|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The IP address of the machine where the Active Directory is installed. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Never logged on users: Indicates the number of AD users who have never logged on to the network. | Number | A healthy AD server is one that has no or very few 'unused' user accounts. A high value is therefore not desired for this measure. To know who these users are, use the detailed diagnosis of this measure. |
| | Inactive users: Indicates the number of users who are currently inactive in the AD server. | Number | To identify the inactive users, use the detailed diagnosis of this measure. The detailed diagnosis displays the Distinguished Name of the user, the date/time he/she logged in last, and the date/time at which the user account was created. This will help you in figuring out how long that user has been inactive. If you think that the user will never again become active, you can proceed to delete that user account. |
| | Disabled users: Indicates the number of user accounts that are currently disabled on the AD server. | Number | To identify the disabled users, use the detailed diagnosis of this measure. The detailed diagnosis displays the Distinguished Name of the user and the date/time at which the user account was created. This will help you in figuring out how long each user account has remained disabled. If you think that the user will never again become active, you can proceed to delete that user account. |

3.5.13 Account Management Events Test

The addition of new users/computers/groups to an Active Directory domain, changes to existing user/computer/group accounts, and deletion of accounts are important to verify that they were performed only by authorized personnel and with no malicious intent. To track such operations, "Audit account management events" provides specific event IDs. Using the **Account Management Events** test, you can continuously track events with the event IDs grouped under *Audit account management events*, and be proactively alerted to the sudden addition/modification/deletion of users/groups/computers in the Active Directory. You can also use the detailed diagnosis of the test to know which user performed the addition/modification/deletion and when.

| | |
|----------------|---|
| Purpose | Continuously tracks events with the event IDs grouped under <i>Audit account management events</i> , and be proactively alerted to the sudden addition/modification/deletion of |
|----------------|---|

MONITORING ACTIVE DIRECTORY SERVERS

| | |
|---------------------------------|--|
| | users/groups/computers in the Active Directory |
| Target of the test | An Active Directory server |
| Agent deploying the test | An internal agent |

| | |
|--|---|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the EventLog Service. Here it is null. 4. SUCSESSEVENTSINDD - By default, this parameter displays <i>none</i>, indicating that by default none of the successful log audits will be reflected in the detailed diagnosis. If you set this parameter to, say 10, then the test will display only the 10 most recent successful log audits in the detailed diagnosis page. Setting this parameter to <i>all</i>, on the other hand will make sure that all successful log audits are listed in the detailed diagnosis. 5. FAILUREEVENTSINDD - By default, this parameter displays <i>all</i>, indicating that by default all the failed log audits will be reflected in the detailed diagnosis. If you set this parameter to, say 10, then the test will display only the 10 most recent log audits that failed, in the detailed diagnosis page. Setting this parameter to <i>none</i>, on the other hand will make sure that none of the failed log audits are listed in the detailed diagnosis. 6. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 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. EVENTS DURING RESTART - By default, the EVENTS DURING RESTART flag is set to Yes. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to No ensures that the agent, when restarted, ignores the events that occurred during the time it was not available. 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: <ul style="list-style-type: none"> ➤ The eG manager license should allow the detailed diagnosis capability ➤ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
|--|---|

10. **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 users 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. 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. This is the default selection.
11. **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 users 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}:{users_to_be_included}:{users_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.

- In the same way, you can also ensure that events generated by specific users on the target host are alone tracked by providing a comma-separated list of users to be monitored – for example, *john,elvis*. In our example however, *all* is specified, indicating that *all* users need be monitored.
- You can similarly indicate if specific users need to be excluded from monitoring. In our example however, *none* is provided to ensure that no users are 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 users 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}:{users_to_be_included}:{users_to_be_excluded}
```

To monitor a specific combination of event sources, event IDs, and users, 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.

12. **STATELESS ALERTS** - Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **CRITICAL** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **CRITICAL**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the **STATELESS ALERTS** flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.

| | | | |
|--------------------------------------|--|-------------------------|--|
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>User password reset by administrator:</p> <p>Indicates the number of times the user password was changed by the administrator since the last measurement period.</p> | Number | <p>Typically, such an event occurs when the administrator attempts to change some other user's password in response to a 'forgot password' call.</p> <p>You can use the detailed diagnosis of this measure to know which admin user attempted the password change on which computer.</p> |
| | <p>User password reset by users:</p> <p>Indicates the number of times the user password was changed by the users themselves since the last measurement period.</p> | Number | <p>You can use the detailed diagnosis of this measure to know which user attempted the password change on which computer.</p> |
| | <p>User accounts created:</p> <p>Indicates the number of user accounts that have been created since the last measurement period.</p> | Number | <p>New user accounts are important to audit to verify that they correspond to a legitimate employee, contractor or application. Outside intruders often create new user accounts to facilitate continued access to the penetrated system. Therefore, you need to eye any sudden increase in the value of this measure with suspicion. You can use the detailed diagnosis of this measure to know which user created new users on which computer.</p> |
| | <p>User accounts deleted:</p> <p>Indicates the number of user accounts that have been deleted since the last measurement period.</p> | Number | <p>You can use the detailed diagnosis of this measure to know which user deleted user accounts on which computer.</p> |
| | <p>User account changed:</p> <p>Indicates the number of times the user account has been changed since the last measurement period.</p> | Number | <p>Certain changes to user accounts are important to audit since they can be a tip-off to compromised accounts. For instance, both insider and outsider computer criminals often gain access to a system by socially engineering the help desk to a user's password. Or a previously disabled account being re-enabled may be suspicious depending on the history and type of the account.</p> <p>You can use the detailed diagnosis of this measure to know which user made changes to user accounts on which computer.</p> |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|--------|---|
| | <p>Computer accounts created:</p> <p>Indicates the number of times computer accounts have been created since the last measurement period.</p> | Number | You can use the detailed diagnosis of this measure to know which user created computer accounts on which computer. |
| | <p>Computer accounts deleted:</p> <p>Indicates the number of computer accounts that have been deleted since the last measurement period.</p> | Number | You can use the detailed diagnosis of this measure to know which user deleted computer accounts on which computer. |
| | <p>Computer accounts changed:</p> <p>Indicates the number of times the computer accounts that have been changed since the last measurement period.</p> | Number | You can use the detailed diagnosis of this measure to know which user changed computer accounts on which computer. |
| | <p>User/Computer object disabled:</p> <p>Indicates the number of times the user/computer object was disabled during the last measurement period.</p> | Number | You can use the detailed diagnosis of this measure to know which user disabled user/computer objects on which computer. |
| | <p>User/Computer object enabled:</p> <p>Indicates the number of times the user/computer object was enabled during the last measurement period.</p> | Number | You can use the detailed diagnosis of this measure to know which user enabled user/computer objects on which computer. |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|--|---|---------------|---|
| | <p>User added to security group:</p> <p>Indicates the number of users who were added to the security group during the last measurement period.</p> | <p>Number</p> | <p>Group changes, especially changes to the group's membership, are very useful to track since groups are used to control access to resources, link security policies and control wireless and remote access all over a Windows network.</p> <p>Security groups are the only group type that you can assign permissions and rights. Security groups are referred to as "security enabled" groups in the security log.</p> <p>You can use the detailed diagnosis of this measure to know which user added users to the security group on which computer.</p> |
| | <p>Security groups deleted:</p> <p>Indicates the number of security groups that were deleted during the last measurement period.</p> | <p>Number</p> | <p>You can use the detailed diagnosis of this measure to know which user deleted security groups on which computer.</p> |
| | <p>Security groups created:</p> <p>Indicates the number of security groups that were created during the last measurement period.</p> | <p>Number</p> | <p>You can use the detailed diagnosis of this measure to know which user created security groups on which computer.</p> |
| | <p>Security groups changed:</p> <p>Indicates the number of security groups that were changed during the last measurement period.</p> | <p>Number</p> | <p>You can use the detailed diagnosis of this measure to know which user changed security groups on which computer.</p> |

Note:

The **STATELESS ALERTING** capability is currently available for the following tests alone, by default:

- EventLog test
- ApplicationEventLog test
- SystemEventLog test
- ApplicationEvents test
- SystemEvents test
- SecurityLog test
- Account Management Events test

If need be, you can enable the **stateless alerting** capability for other tests. To achieve this, follow the steps given below:

- Login to the eG manager host.
- Edit the **eg_specs.ini** file in the `<EG_INSTALL_DIR>\manager\config` directory.
- Locate the test for which the **Stateless Alarms** flag has to be enabled.
- Insert the entry, **-statelessAlerts yes**, into the test specification as depicted below:

```
EventLogTest::$hostName:$portNo=$hostName, -auto, -host $hostName -port
$portNo -eventhost $hostIp -eventsrc all -excludedSrc none -useWmi yes -
statelessAlerts yes -ddFreq 1:1 -rptName $hostName, 300
```

- Finally, save the file.
- If need be, you can change the status of the **statelessAlerts** flag by reconfiguring the test in the eG administrative interface.

Once the **stateless alerting capability** is enabled for a test (as discussed above), you will find that everytime the test reports a problem, the eG manager does the following:

- Closes the alarm that pre-exists for that problem;
- Sends out a normal alert indicating the closure of the old problem;
- Opens a new alarm and assigns a new alarm ID to it;
- Sends out a fresh email alert to the configured users, intimating them of the new issue.

In a redundant manager setup, the secondary manager automatically downloads the updated **eg_specs.ini** file from the primary manager, and determines whether the stateless alerting capability has been enabled for any of the tests reporting metrics to it. If so, everytime a threshold violation is detected by such a test, the secondary manager will perform the tasks discussed above for the problem reported by that test. Similarly, the primary manager will check whether the stateless alert flag has been switched on for any of the tests reporting to it, and if so, will automatically perform the above-mentioned tasks whenever those tests report a deviation from the norm.

Note:

- Since alerts will be closed after every measurement period, alarm escalation will no longer be relevant for tests that have **statelessAlerts** set to **yes**.
- For tests with **statelessAlerts** set to **yes**, **statelessAlerts** will apply for all measurements of that test (i.e., it will not be possible to only have one of the measurements with stateless alerts and others without).
- If **statelessAlerts** is set to **yes** for a test, an alarm will be opened during one measurement period (if a threshold violation happens) and will be closed prior to the next measurement period. This way, if a threshold violation happens in successive measurement periods, there will be one alarm per measurement period. This will reflect in all the corresponding places in the eG Enterprise system. For example, multiple alerts in successive measurement periods will result in multiple trouble tickets being opened (one for each measurement period). Likewise, the alarm history will also show alarms being opened during a measurement period and closed during the next measurement period.

3.5.14 Active Directory Computers Test

This test takes stock of the total number of computers managed by the AD server and the status of these computers, so that administrators can determine from a single glance which computers are inactive/unused.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Takes stock of the total number of computers managed by the AD server and the status of these computers, so that administrators can determine from a single glance which computers are inactive/unused | | |
| Target of the test | An Active Directory or Domain Controller on Windows | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The IP address of the machine where the Active Directory is installed. 3. PORT – The port number through which the Active Directory communicates. The default port number is 389. | | |
| Outputs of the test | One set of results for every Active Directory being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING ACTIVE DIRECTORY SERVERS

| | | | |
|------|--|--------|---|
| test | <p>Never logged on computers:</p> <p>Indicates the number of computers to which no user has ever logged in.</p> | Number | To know which computers are unused, use the detailed diagnosis of this measure. You can consider removing such computers to reduce the workload of the AD server. |
| | <p>Inactive computers:</p> <p>Indicates the number of computers that are currently inactive.</p> | Number | To identify the inactive computers, use the detailed diagnosis of this measure. The detailed diagnosis displays the Distinguished Name of the computer, the age of the computer, and the date/time at which the computer was created. This will help you in figuring out how long that computer has been inactive. If the computer has been inactive for too long, you may think about deleting it from the AD server. |
| | <p>Disabled computers:</p> <p>Indicates the number of computers that are currently disabled on the AD server.</p> | Number | To identify the disabled computers, use the detailed diagnosis of this measure. The detailed diagnosis displays the Distinguished Name of the computer and the date/time at which the computer was created. |
| | <p>Total computers:</p> <p>Indicates the total number of computers managed by the AD server.</p> | Number | Use the detailed diagnosis of this measure to know the Distinguished Names of the computers. |

Monitoring the BizTalk Server

Microsoft BizTalk server provides a powerful web-based development and execution environment that integrates loosely coupled, long-running business processes, both within and between businesses. The server provides a standard gateway for sending and receiving documents across the Internet, as well as providing a range of services that ensure data integrity, delivery, security, and support for the BizTalk Framework and other key document formats.

As mission-critical business processes are integrated via the BizTalk server, it is imperative that the BizTalk server itself stays in good health at all times. To ensure the continuous availability and smooth functioning of the BizTalk server, you need to constantly monitor the server, and promptly detect performance issues, so that the issues can be fixed before they prove fatal to the critical business processes that ride on the server.

eG Enterprise offers two dedicated models for monitoring the BizTalk server - one each for BizTalk Server 2000 and BizTalk Server 2010. Both these models are capable of monitoring the entire pipeline of the processes happening within the BizTalk server. This chapter takes a closer look at both the models.

4.1 Monitoring the BizTalk Server 2000

BizTalk server 2000 includes a document interchange engine, a business process execution engine, a business document editor, a business document mapper, and a set of business document and server management tools. Initially, an agreement should be made between the organizations, to determine the following:

- the source and destination locations of the business documents
- the transportation medium to be used,
- the source and destination formats of the business documents

After the agreement, the business process diagram should be drawn by using the VISIO style-drawing tool. The business process diagram is then compiled to a XLANG file using XLANG Scheduler tool given by the BizTalk Server environment. The XLANG engine loads the XLANG file at runtime environment.

The sender application (say Application 1 of Organization A) is responsible for generating business documents in well-defined XML format (for e.g., a purchase order). This business document is submitted to the BizTalk server. Then, the business document has to be transformed using Schema transformations. Here, a mapping is done to transform the business document from the source organization's native representation to the representation requested by the destination organization (for e.g., the source organization may submit an XML document, but the destination organization may require the document in EDI format). The source XML document is parsed to determine the well-defined XML standard. Encoding and encryption is done when specified. Until this stage, the documents are

MONITORING THE BIZTALK SERVER

available in the **work queue**. Then, the document is serialized to the standard that is ready for transmission. The document in the interchange form will be available in the **scheduled queue**. By using the specified transportation medium, the document interchanges are transmitted to the destination location that has been specified in the agreement. Decryption and decoding of the business document is done at the receiving end (Application 2 of Organization B) if necessary. At this stage, the business document is in the target representation form. It is received by the target application that is running in Organization B. The business documents and interchanges will be in the **retry queue** when the BizTalk server is overloaded. In this case, the documents and interchanges are re-submitted to the BizTalk server automatically. When any error happens during the above stages, the documents and interchanges are moved to the **suspended queue** and cannot be re-submitted to the BizTalk server.

Since a BizTalk server acts as a bridge between systems having heterogeneous inputs, it is critical for the BizTalk server to perform optimally so as not to choke the performance of the system being integrated. The eG Enterprise suite of products is capable of monitoring the BizTalk server 2000 inside out. The *BizTalk* monitoring model that is used by the eG Enterprise suite for monitoring the BizTalk server is shown in Figure 4.11.

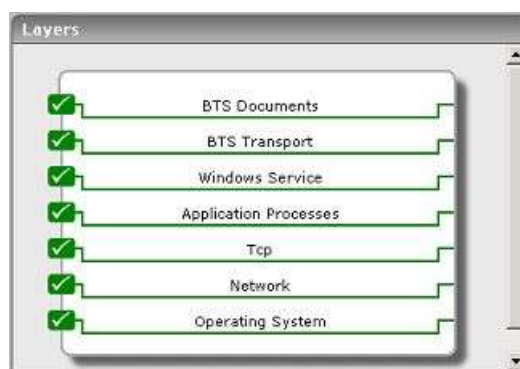


Figure 4.1: Layer model of a BizTalk server

Each layer of Figure 4.1 is mapped to tests that report a wide variety of metrics revealing the internal health of the BizTalk Server 2000. Using the metrics so reported, administrators can find quick and easy answers for many persistent performance queries, such as the following:

- Is the rate of interchange decodes and interchange decrypts unusually low?
- How is the transport mechanism functioning? Could problems in this mechanism be causing a slowdown in the reception and transmission of the interchange?
- Can the BizTalk server encode, encrypt, and serialize interchanges?
- Are applications able to receive and submit documents quickly to the BizTalk server?
- Is the BizTalk server experiencing any delays in document processing?
- Is the BizTalk server able to map documents?

The details about the 5 layers at the bottom of Figure 4.1 are available in the *Monitoring Unix and Windows Servers* document. The sections to come will therefore discuss the top 2 layers only.

4.1.1 The BTS Transport Layer

This layer monitors the transportation of the BizTalk documents and interchanges using the InterChangeRcvd test and InterChangeXmit test shown in Figure 4.2. A business document is an XML document containing the business transaction data. This transaction data may represent a purchase order, invoice, sales forecast, or any other

MONITORING THE BIZTALK SERVER

business information. A BizTalk document is a combination of one or more business documents, and zero or more binary file(s). BizTalk interchanges refer to a collection of one or more document instances that comprises a single transmission. This is exchanged from application to application within an organization or from one trading partner to another.



Figure 4.2: Tests mapping to the BTS Transport layer

4.1.1.1 Inter Changes Received Test

BizTalk messaging service enables the administrator to send, receive, parse, and verify the integrity of the documents, track interchanges and documents, and provide secure methods for exchanging documents with trading partners and applications. This test tracks the performance of the messaging service while receiving interchanges from the BizTalk server.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | This test measures the performance of the messaging service while receiving the interchanges from the BizTalk server. | | |
| Target of the test | A BizTalk server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – This indicates how often should the test be executed 2. HOST – The IP address of the machine where BizTalk has been installed. 3. PORT - Not applicable to this test. Set to NULL. | | |
| Outputs of the test | One set of results for every BizTalk server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING THE BIZTALK SERVER

| | | | |
|------|--|----------------|--|
| test | <p>Decode rate:</p> <p>This measure indicates the number of interchanges being decoded per second by the runtime process.</p> | Intchanges/Sec | <p>A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>In case of an unusually low value, verify the status of the interchange in the suspended queue using the BizTalk server administration.</p> |
| | <p>Decrypt rate:</p> <p>This measure indicates the number of interchanges being decrypted per second by the runtime process.</p> | Intchanges/Sec | <p>A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>If the value of this measure is unusually low, then it indicates that the certificate might have expired. Verify the validity of the certificate in the Certificate Microsoft Management Console Snap-in.</p> |
| | <p>Receive rate:</p> <p>This measure indicates the number of interchanges received by the BizTalk messaging service between trading partners.</p> | Intchanges/Sec | <p>A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>If the value of this measure is unusually low, then it indicates that the transport mechanism (HTTP/MSMQ/FTP) used may not be functioning.</p> |

a.

4.1.1.2 Inter Changes Transmitted Test

This test tracks the performance of the messaging service while receiving interchanges from the BizTalk server. The outputs of the test are given below:

| | |
|---|---|
| Purpose | This test measures the performance of the messaging service while transmitting the interchanges to the BizTalk server. |
| Target of the test | A BizTalk server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – This indicates how often should the test be executed HOST – The IP address of the machine where BizTalk has been installed. PORT - Not applicable to this test. Set to NULL. |
| Outputs of the test | One set of results for every BizTalk server being monitored |

MONITORING THE BIZTALK SERVER

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|--|------------------|---|
| | <p>Encode rate:</p> <p>This measure indicates the number of interchanges being encoded per second by the runtime process.</p> | Intchanges/Sec | <p>A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>In case of an unusually low value, verify the status of the interchange in the suspended queue using the BizTalk server administration. If the status corresponding to an interchange is Encoding, then it implies that the BizTalk server could not encode the interchange. Resubmitting the interchange to the BizTalk server may solve this problem.</p> |
| | <p>Encrypt rate:</p> <p>This measure indicates the number of interchanges being encrypted per second by the runtime process.</p> | Intchanges/Sec | <p>A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>In case of an unusually low value, verify the status of the interchange in the suspended queue using the BizTalk server administration. If the status corresponding to the interchange is Encrypting, then it signifies that the BizTalk server could not encrypt this interchange. Also, verify the expiration of the certificate in the Certificate Microsoft Management Console snap-in.</p> |
| | <p>Serialize rate:</p> <p>This measure indicates the number of interchanges being serialized per second by the BizTalk runtime process.</p> | Intchanges/Sec | <p>A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>In case of an unusually low value, verify the status of the interchange in the suspended queue. If the status corresponding to the interchange is Serializing, then it implies that the BizTalk server could not convert the interchange to its native format.</p> <p>Resubmitting the interchange can solve this problem.</p> |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|-----------------------|--|
| | <p>Transmit rate:</p> <p>This measure indicates the number of interchanges being transmitted per second by the BizTalk messaging service.</p> | <p>Intchanges/Sec</p> | <p>A high value over a period may indicate that transmission took a long time to attain completion.</p> <p>A value of -1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>If the value of this measure is unusually low, verify the transport address in the channel. Correct the problem in the channel and resubmit the interchange.</p> <p>Alternatively, the BizTalk server might have taken a long time to transmit the interchange. Verify the transport mechanism used.</p> <p>Another reason could be that the BizTalk administrator might have moved the interchange to the suspended queue, resubmitted the interchange from the suspended queue.</p> <p>Alternatively, the computer on which the BizTalk server could be running out of memory, restart the server and resubmit all the interchanges in the suspended queue.</p> |
|--|--|-----------------------|--|

4.1.2 The BTS Documents Layer

This layer reports the statistics about the various attributes of the documents being handled by the BizTalk server using the DocReceive test and DocSubmit test shown in Figure 4.3.



Figure 4.3: Tests mapping to the BTS Documents layer

4.1.2.1 Documents Received Test

This test tracks the performance of the messaging service while it is receiving documents from the BizTalk server.

| | |
|----------------|--|
| Purpose | This test measures the performance of the messaging service while receiving documents from the |
|----------------|--|

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|---|
| | BizTalk server. | | |
| Target of the test | A BizTalk server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – This indicates how often should the test be executed HOST – The IP address of the machine where BizTalk has been installed. PORT - Not applicable to this test. Set to NULL. | | |
| Outputs of the test | One set of results for every BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Receive rate: This measure indicates the number of documents being received per second by the application from the BizTalk server. | Docs/Sec | A value of –1 for this measure indicates that either the BizTalk messaging service or XLANG Schedule Restart Service or Distributed Transaction Coordinator (MSDTC) may not be running. If the value of this measure is unusually low, then verify the interface between the BizTalk server and the application. |

4.1.2.2 Documents Submitted Test

The DocSubmitTest tracks the performance of the messaging service while submitting documents to the BizTalk server.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | This test measures the performance of the messaging service while submitting documents to the BizTalk server. | | |
| Target of the test | A BizTalk server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – This indicates how often should the test be executed HOST – The IP address of the machine where BizTalk has been installed. PORT - Not applicable to this test. Set to NULL. | | |
| Outputs of the test | One set of results for every BizTalk server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |
| | | | |

MONITORING THE BIZTALK SERVER

| | | | |
|------|---|----------|---|
| test | <p>Submit rate:</p> <p>This measure shows the number of documents submitted asynchronously per second to the BizTalk server from the application. Once submitted, the BizTalk server holds the documents in the work queue for further processing.</p> | Docs/Sec | <p>A sudden increase in the value of this measure denotes a change in the workload.</p> <p>A value of -1 for this measure indicates that either the BizTalk messaging service or XLANG Schedule Restart Service or Distributed Transaction Coordinator (MSDTC) may not be running.</p> <p>If the value of this measure is unusually low, then verify the interface between the BizTalk server and the application or check the event log entry in the BizTalk server administration.</p> |
| | <p>Map rate:</p> <p>BizTalk runtime process maps the actual document content from one structural form to another.</p> <p>This measure shows the number of documents that have been mapped per second by the runtime process.</p> | Docs/Sec | <p>A sudden increase in the value of this measure indicates that the BizTalk runtime process is mapping larger number of documents. This scenario indicates an increased workload.</p> <p>A value of -1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>For other reasons, verify the status of the document available in the suspended queue using the BizTalk server administration. If the status corresponding to a document is Mapping then it indicates that the document has been failed to map. To rectify this problem, delete the document from the suspended queue, correct the map and resubmit the document.</p> |
| | <p>Parse rate:</p> <p>This measure shows the number of documents in the work queue that is being parsed per second by the appropriate parser.</p> | Docs/Sec | <p>A sudden increase in the value of this measure indicates that the parser is parsing larger number of documents. This scenario may be due to the deletion of large number of documents from the suspended queue, which affects the performance of the parser.</p> <p>A value of -1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>In case of an unusually low value, verify the status of the documents available in the suspended queue using the BizTalk server administration. If the status corresponding to the document is parsing then it indicates that the BizTalk server was unable to parse the data. The other reasons could be that the timestamp of the document is no longer valid, or the document does not contain enough information to locate the channel.</p> |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|----------|--|
| | <p>Process rate:</p> <p>This measure indicates the number of documents being processed successfully (necessary changes to the document) per second by the runtime process.</p> | Docs/Sec | <p>A high value for this measure over a period may indicate that the runtime system is processing larger number of documents. This scenario may indicate a change in the workload.</p> <p>A value of -1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running.</p> <p>Incase of an unusually low value, verify the status of Microsoft SQL server in the Service Manager tool available in the Microsoft SQL server environment. Also check the status of the document available in the suspended queue using the BizTalk server administration.</p> |
|--|---|----------|--|

b.

4.2 Monitoring the BizTalk Server 2010

BizTalk Server is Microsoft's Integration and connectivity server solution. A mature product on its seventh release, BizTalk Server 2010 provides a solution that allows organizations to more easily connect disparate systems. Including over 25 multi-platform adapters and a robust messaging infrastructure, BizTalk Server provides connectivity between core systems both inside and outside your organization. In addition to integration functionality, BizTalk also provides strong durable messaging, a rules engine, EDI connectivity, Business Activity Monitoring (BAM), RFID capabilities and IBM Host/Mainframe connectivity.

The BizTalk Server includes a range of technologies. The figure below illustrates the product's major components.

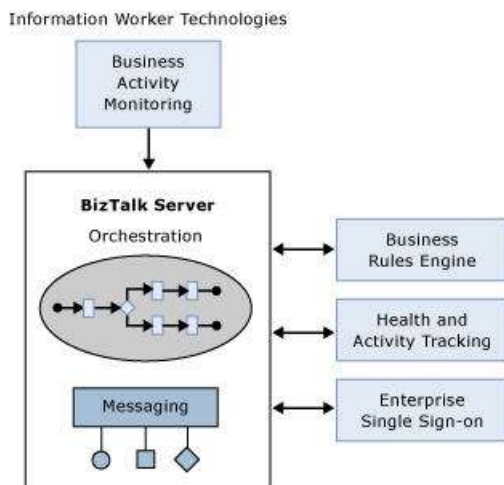


Figure 4.4: The major components of a BizTalk server

As the figure suggests, the heart of the product is the BizTalk Server Engine. The engine has two main parts:

- A messaging component that provides the ability to communicate with a range of other software. By relying on adapters for different kinds of communication, the engine can support a variety of protocols and data formats, including Web services and many others.
- Support for creating and running graphically-defined processes called orchestrations. Built on top of the engine's messaging components, orchestrations implement the logic that drives all or part of a business

MONITORING THE BIZTALK SERVER

process.

Several other BizTalk components can also be used in concert with the engine, including:

- A Business Rule Engine that evaluates complex sets of rules.
- A Group Hub that lets developers and administrators monitor and manage the engine and the orchestrations it runs.
- An Enterprise Single Sign-On (SSO) facility that provides the ability to map authentication information between Windows and non-Windows systems.

On top of this foundation, BizTalk Server includes Business Activity Monitoring, which information workers use to monitor a running business process. The information is displayed in business rather than technical terms, and business users determine what information is displayed.

As the present era is all about business process management, the BizTalk server plays a vital role in connecting and communicating with disparate business processes that may be operating within an organization or across organizations. If this 'connector' malfunctions, it could break the only link that exists between the processes, thereby significantly affecting the way the enterprise functions. All software-dependent activities of the enterprise - from the performance of simple, routine operations to the execution of critical business transactions - could either experience delays or could come to a virtual standstill. If such adversities are to be avoided, the BizTalk server has to be monitored 24x7.

eG Enterprise provides a *BizTalk 2010* monitoring model that provides in-depth monitoring of the BizTalk Server 2010. Each layer of this model is mapped to a series of tests that report issues in the overall health of the adapters and protocols supported by the BizTalk server, thus shedding light on applications with which the server is unable to communicate.



Figure 4.5: The layer model of the BizTalk Server 2010

The metrics extracted by these tests enable administrators to find answers to persistent performance queries such as the following:

- Which host instance is heavily loaded in terms of documents processed?
- Is any host instance experiencing processing bottlenecks?
- Have any documents been suspended by a host instance? If so, which host instance is it?
- Have any request messages timed out without response messages?

MONITORING THE BIZTALK SERVER

- How are the receive and send adapters on a host instance handling the load? Is any receive/send adapter experiencing a slowdown in processing? Which adapter is it - the file adapter, FTP adapter, HTTP adapter, Msmq adapter, POP3 adapter, SMTP adapter, SOAP adapter, or the SQL adapter?
- Is the messaging engine experiencing any latencies - if so, where did the delay originate? while delivering messages to the MessageBox, or while delivering messages to a target application?
- Are too many messages pending processing in the host queue?
- Are any SQL agent jobs taking too long to complete? If so, which ones?
- Is the depth of the spool table optimal, or is it growing continuously?
- Is the tracking data table growing uncontrollably in size?
- Have too many orchestrations been suspended or discarded?
- What is the rate at which dehydrations and rehydrations take place?
- Have the orchestrations acknowledged all the messages they received, or are there too many pending messages?
- Is there a contention for physical memory resources on any host instance?
- Have any BAM (Business Activity Monitoring) events failed?
- Has the tracking data decode service failed to process any batches?
- How is the host throttling mechanism functioning? Are message processing and/or message publishing throttled? Were any delays imposed on the message processing/publishing rates?
- Has process memory consumption exceeded its threshold?
- Has thread count exceeded its threshold?

The sections that follow will discuss each layer of Figure 4.5 in great detail.

4.2.1 The Messaging Engine Layer

The BizTalk Server Messaging engine enables users to create business processes that spans multiple applications by providing two primary things:

- A way to specify and implement the logic driving that business process
- A mechanism for communicating across the applications that the business process uses

The figure below illustrates the main components of the engine that address these two problems.

MONITORING THE BIZTALK SERVER

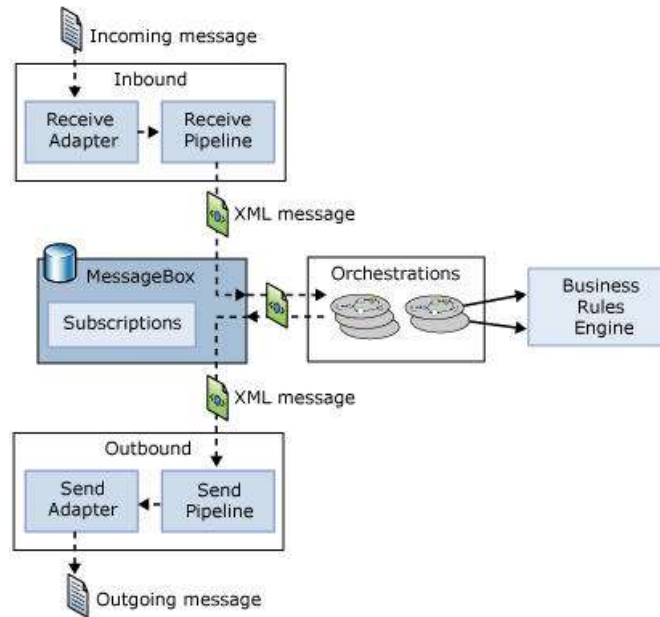


Figure 4.6: Messaging architecture

As the diagram shows, a message is received through a **receive adapter**. Different adapters provide different communication mechanisms, so a message might be acquired by accessing a Web service, reading from a file, or in some other way. The message is then processed through a **receive pipeline**. This pipeline can contain various components that do things such as converting the message from its native format into an XML document, validating a message's digital signature, and more. The message is then delivered into a database called the **MessageBox**, which is implemented using Microsoft SQL Server.

The logic that drives a business process is implemented as one or more **orchestrations**, each of which consists of executable code. These orchestrations are not created by writing code in a language such as C#, however. Instead, a business analyst or (more likely) a developer uses an appropriate tool to graphically organize a defined group of shapes to express conditions, loops, and other behavior. Orchestrations can optionally use the **Business Rule Engine**, which provides a simpler and more easily modified way to express complex sets of rules in a business process.

Each orchestration creates **subscriptions** to indicate the kinds of messages it wants to receive. When an appropriate message arrives in the MessageBox, that message is dispatched to its target orchestration, which takes whatever action the business process requires. The result of this processing is typically another message, produced by the orchestration and saved in the MessageBox. This message, in turn, is processed by a **send pipeline**, which may convert it from the internal XML format used by BizTalk Server to the format required by its destination, add a digital signature, and more. The message is then sent out using a **send adapter**, which uses an appropriate mechanism to communicate with the application for which this message is destined.

This layer monitors the messaging engine of the BizTalk server, measures the load on the engine, reports how quickly every send and receive adapter processes the message load, and sheds light on current / potential processing bottlenecks (if any) in the engine. All the tests mapped to this layer report metrics for each host instance on the BizTalk server. A *host* is a logical representation of a Microsoft Windows process that executes BizTalk Server artifacts such as send ports and orchestrations. A *host instance* is the physical representation of a host on a specific server.

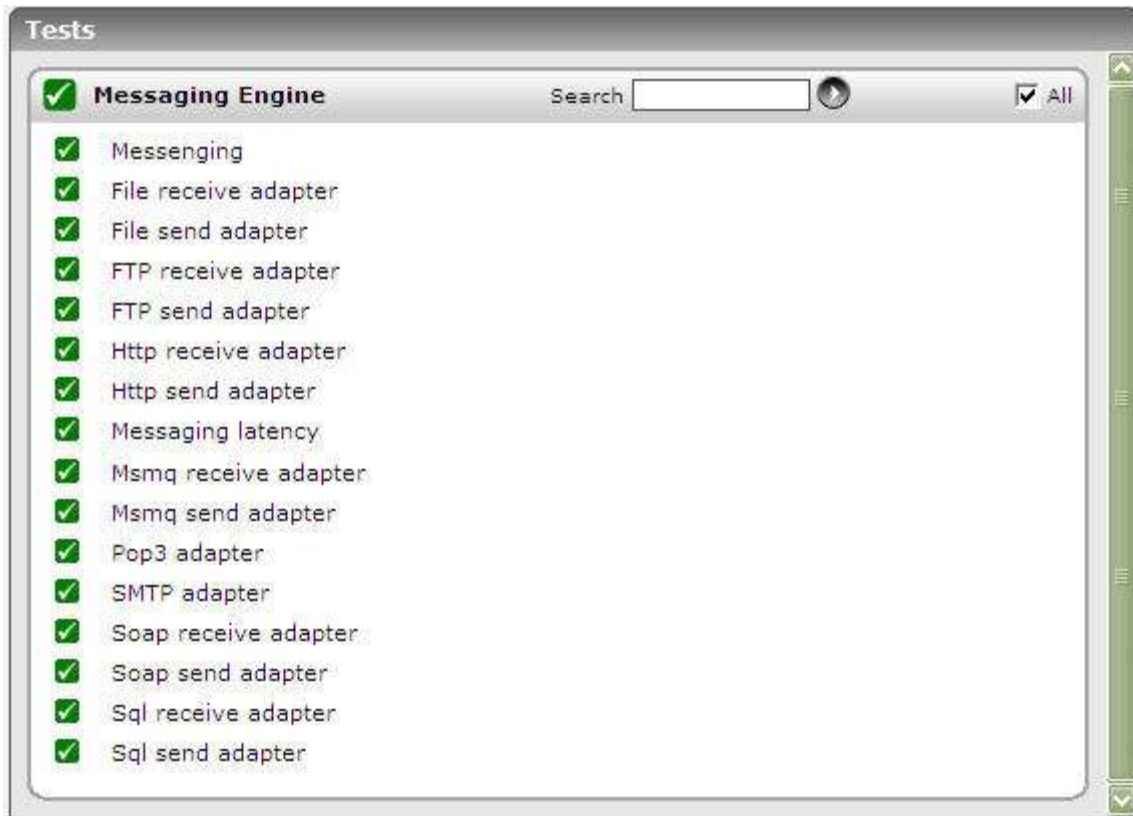


Figure 4.7: The tests mapped to the Messaging Engine layer

4.2.1.1 BT Messaging Test

This test monitors the documents received and sent by each host instance on the BizTalk server, and reports the load on that host instance and delays experienced by the host instance while processing the documents.

Using this test, administrators can easily isolate host instances that are overloaded or are experiencing bottlenecks in processing.

| | |
|---|---|
| Purpose | Monitors the documents received and sent by each host instance on the BizTalk server, and reports the load on that host instance and delays experienced by the host instance while processing the documents |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. |

MONITORING THE BIZTALK SERVER

| | | | |
|--------------------------------------|--|-------------------------|--|
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active receive locations: Indicates the number of receive locations currently enabled in this host instance. | Number | A <i>receive location</i> is the configuration of a single endpoint (URL) to receive messages. |
| | Documents processed: Indicates the number of documents processed by this host instance. | Number | This is a good indicator of the load handled by a host instance. Comparing the value of this measure across host instances will reveal which instance is currently overloaded. |
| | Documents processed: Indicates the rate at which this host instance processed documents. | Docs/Sec | A very low value or a consistent decrease in the value of this measure indicates a slowdown in the corresponding host instance. |
| | Documents received: Indicates the number of documents received by this host instance from a target source. | Number | This is a good indicator of the load handled by a host instance. Comparing the value of this measure across host instances will reveal which instance is currently overloaded. |
| | Documents received: Indicates the rate at which documents were received by this host instance. | Number | |

| | | | |
|--|---|---------------|---|
| | <p>Documents suspended:</p> <p>Indicates the number of documents that have been suspended by this host instance.</p> | <p>Number</p> | <p>By default, the BizTalk server places failed messages/documents in the Suspended queue. The value of this measure indicates the number of documents in the Suspended queue.</p> <p>A message failure can occur in one of the following instances:</p> <ul style="list-style-type: none"> • Failures in the disassembly phase: Processing might also fail during the disassembly phase; that is, failure in one of the pipeline components. For example, decryption failed due to absence of decryption cert on the processing server, or parsing failure due to problem either in the schema or in the message. • Failures in routing: After a message disassembles successfully, the next potential failure point is routing; for example, users enable a corresponding receive location of an orchestration and forget to enlist the orchestration. In this case, the message picked up from the receive location fails routing and the MessageBox database generates a Routing Failure report. <p>Routing Failure reports are listed in the BizTalk Server Administration Console as non-resumable suspended messages. Each Routing Failure report contains a message property snap shot taken when the routing failure occurred. You can use the information in each report to determine why routing failed for its associated message. If the associated message is resumable, you can correct the routing problem and resume the message so that processing continues.</p> |
|--|---|---------------|---|

| | | | |
|--|---|-----------------|---|
| | | | <ul style="list-style-type: none"> Failures during the transformation phase: When a message is received from Receive Location, the message is disassembled (for example, decrypted and parsed), the message might optionally be transformed to a different format via an Inbound Map specified on a receive Port, and published to the MessageBox for routing to an orchestration or a Send Port. In this case, processing may fail during transformation phase due to incorrect Inbound Map, or problems in the schema or in the message received. <p>When a message is to be sent to a Send Location, an Outbound Map configured on Send Port might optionally transform the message. Then the transformed message is assembled and handed to the adapter for final transmission to the Send Location. In this case, processing may fail during transformation phase due to incorrect Outbound Map or problem in schema or source message.</p> <ul style="list-style-type: none"> Failures in the message assembly phase: Processing can also fail during message assembly phase – in other words, failing in pipeline component. After a message successfully assembles, the next potential failure point becomes transmission to Send Location; for example, the Send Location (which belongs to the partner) might be down or not exist. |
| | <p>Documents suspended:</p> <p>Indicates the rate at which documents were suspended by this host instance.</p> | <p>Docs/Sec</p> | |

| | | | |
|--|---|--------|--|
| | <p>Request/Response timeouts:</p> <p>Indicates the number of request messages that have not received a response message within the time limit specified by the adapter associated with this host instance.</p> | Number | A high value of this measure could indicate that too many messages are getting timed out. You may want to consider reconfiguring the timeout period. |
|--|---|--------|--|

4.2.1.2 BT Messaging Agents Test

Most of the processing that takes place on a BizTalk server occurs within a logical entity known as a BizTalk Server host instance, which is a process running as a Windows service or an isolated host process on the BizTalk server. To manage the use of resources by a host instance process, BizTalk Server utilizes an adjustable throttling mechanism that governs the flow and processing of messages through a host instance.

The throttling mechanism moderates the workload of the host instance to ensure that the workload does not exceed the capacity of the host instance or any downstream host instances. The throttling mechanism also prevents a condition known as resource contention that can lower the overall performance of the host instance process or other system processes. Resource contention occurs when one or more processes consume a limited resource to the detriment of themselves and/or another process. For example, the consumption of excessive memory or threads can lead to memory allocation failure or high thread context-switches, which can impact the performance of the process. Resource contention like this can be detrimental to the overall performance of BizTalk Server.

The host throttling mechanism also detects when available resources are being underutilized. If available resources are underutilized then the throttling mechanism allows additional messages to be processed by a host instance. The host throttling mechanism continually monitors if available resources are being over or underutilized and adjusts message flow through the host instance accordingly.

The BizTalk Server host throttling mechanism helps to ensure that the system operates at an optimal and sustainable level.

This test measures the efficiency of the host throttling mechanism.

| | |
|---------------------------------|--|
| Purpose | Measures the efficiency of the host throttling mechanism |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|-----------------------|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

| | | | |
|--|--|-----------------|--|
| | <p>Publishing delay:</p> <p>Indicates the current delay imposed on each message publishing batch.</p> | <p>MilliSec</p> | <p>This measure is applicable only if the message publishing is throttled and if the message publishing batch is not exempted from throttling.</p> <p><i>Message publishing throttling</i> in BizTalk Server, is applied to host instances that contain receive adapters or orchestrations that publish messages to the MessageBox database. An inbound host throttling condition can be triggered under the following conditions:</p> <ul style="list-style-type: none"> • The amount of memory, the number of threads, or the number of database connections used by the host instance exceeds the throttling thresholds defined • Downstream hosts are unable to process the messages that are published. • The Message publishing incoming rate for the host instance exceeds the Message publishing outgoing rate * the specified Rate overdrive factor (percent) value. • The default throttling behavior has been modified by setting a registry value or values to control the throttling behavior of a host process. <p>Depending on the severity of the throttling condition, the following actions are taken:</p> <ul style="list-style-type: none"> • A progressive delay in the processing logic of the host instance is implemented. The delay may be implemented when an End Point Manager (EPM) thread receives a batch of messages from the transport adapter, and/or when the EPM submits a batch of messages to be published into the MessageBox database. Both the duration of the processing delay and the rate at which the duration is incremented scale with the severity of the throttling condition. |
|--|--|-----------------|--|

| | | | |
|--|---|-----------------|---|
| | | | <ul style="list-style-type: none"> • The number of threads that are available to the End Point Manager (EPM) is restricted. The EPM receives batches of messages from adapters and publishes the messages to the MessageBox database. By default, the EPM is configured to use 20 threads per CPU. If the host throttling mechanism detects a stress condition for inbound processing then it can temporarily reduce the number of threads available to the EPM until the stress condition is eliminated. The EPM cannot process messages from transport adapters or deliver message batches to the MessageBox database unless an EPM thread is available to service the inbound message batch. • The use of memory and other resources is reduced as applicable. BizTalk Server can send instructions to other service classes to limit memory use by dehydrating running schedules, shrinking memory cache size, and by limiting the usage of memory-intensive threads. |
| | <p>Publishing incoming rate:</p> <p>Indicates the rate at which the messages are being sent by the message agent to the database of this host instance for publishing.</p> | <p>Msgs/Sec</p> | <p>A <i>message publishing throttling condition</i> is also triggered when the Message publishing incoming rate for the host instance exceeds the Message publishing outgoing rate * the specified Rate overdrive factor (percent) value. The Rate overdrive factor (percent) value is defined on the Message Publishing Throttling Settings dialog box available from the Advanced page of the Host Properties dialog box.</p> |
| | <p>Publishing outgoing rate:</p> <p>Indicates the rate at which the messages are actually published by the message agent in the database of this host instance.</p> | <p>Msgs/Sec</p> | |

| | <p>Publishing throttling state:</p> <p>Indicates whether the system is throttling the message publishing i.e., indicates whether the XLANG message processing and inbound transports are affected.</p> | <p>Number</p> | <p>This measure indicates any one of the following values while indicating whether the system is throttling the message publishing or not.</p> <table border="1" data-bbox="980 338 1370 1236"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Not throttling</td> </tr> <tr> <td>2</td> <td>Throttling due to imbalanced message publishing rate (input rate exceeds output rate)</td> </tr> <tr> <td>4</td> <td>Throttling due to process memory pressure</td> </tr> <tr> <td>5</td> <td>Throttling due to system memory pressure</td> </tr> <tr> <td>6</td> <td>Throttling due to database growth</td> </tr> <tr> <td>8</td> <td>Throttling due to high session count</td> </tr> <tr> <td>9</td> <td>Throttling due to high thread count</td> </tr> <tr> <td>11</td> <td>Throttling due to user override on publishing</td> </tr> </tbody> </table> | Value | State | 0 | Not throttling | 2 | Throttling due to imbalanced message publishing rate (input rate exceeds output rate) | 4 | Throttling due to process memory pressure | 5 | Throttling due to system memory pressure | 6 | Throttling due to database growth | 8 | Throttling due to high session count | 9 | Throttling due to high thread count | 11 | Throttling due to user override on publishing |
|-------|---|---------------|---|-------|-------|---|----------------|---|---|---|---|---|--|---|-----------------------------------|---|--------------------------------------|---|-------------------------------------|----|---|
| Value | State | | | | | | | | | | | | | | | | | | | | |
| 0 | Not throttling | | | | | | | | | | | | | | | | | | | | |
| 2 | Throttling due to imbalanced message publishing rate (input rate exceeds output rate) | | | | | | | | | | | | | | | | | | | | |
| 4 | Throttling due to process memory pressure | | | | | | | | | | | | | | | | | | | | |
| 5 | Throttling due to system memory pressure | | | | | | | | | | | | | | | | | | | | |
| 6 | Throttling due to database growth | | | | | | | | | | | | | | | | | | | | |
| 8 | Throttling due to high session count | | | | | | | | | | | | | | | | | | | | |
| 9 | Throttling due to high thread count | | | | | | | | | | | | | | | | | | | | |
| 11 | Throttling due to user override on publishing | | | | | | | | | | | | | | | | | | | | |

| | | | |
|--|--|-----------------|---|
| | <p>Delivery delay:</p> <p>Indicates the current delay imposed on each message delivery batch.</p> | <p>MilliSec</p> | <p>This measure is applicable only if message delivery is throttled.</p> <p><i>Message processing throttling</i> in BizTalk Server, is applied to host instances that contain orchestrations or send adapters that receive and deliver or process messages that have been published to the MessageBox. An outbound host throttling condition can be triggered under the following conditions:</p> <ul style="list-style-type: none"> • The amount of memory, the number of threads, or the number of database connections used by the host instance exceeds the throttling thresholds defined • The Message delivery incoming rate for the host instance exceeds the Message delivery outgoing rate * the specified Rate overdrive factor (percent) value. • The number of messages being processed concurrently by the host instance exceeds the In-process messages per CPU * the number of CPUs available on the box. • The default throttling behavior has been modified by setting a registry value or values to control the throttling behavior of a host process. <p>Depending upon the severity of the throttling condition, the following actions are taken:</p> <ul style="list-style-type: none"> • A progressive delay in the processing logic of the host instance is implemented before delivering the messages to the outbound transport adapter or the orchestration engine for processing the messages. Both the duration of the processing logic delay and the rate at which the duration is incremented scale with the severity of the throttling condition. |
|--|--|-----------------|---|

| | | | |
|--|--|-----------------|---|
| | | | <ul style="list-style-type: none"> • The number of messages that can be held by the in-memory queue is limited. The in-memory queue serves as a temporary placeholder for delivering messages from the MessageBox to the Message Agent which in turn delivers messages to XLANG and send adapters. By default, the in-memory queue is set to hold 100 messages per CPU. When the queue is full, no more messages are de-queued from the MessageBox until the in-memory queue is freed up. • The size of the Message Agent thread pool is limited. By limiting the Message Agent thread pool size, the host throttling mechanism effectively reduces the amount of messages that are delivered to XLANG and adapters. • The use of memory and other resources is reduced as applicable. BizTalk Server can send instructions to other service classes to limit memory use by dehydrating running schedules, shrinking memory cache size, and by limiting the usage of memory intensive threads. |
| | <p>Delivery incoming rate: Indicates the rate at which the messages are delivered to the Orchestration engine or the Messaging engine of this host instance.</p> | <p>Msgs/Sec</p> | <p>A <i>message processing throttling</i> condition can also be triggered if the message Delivery incoming rate for the host instance exceeds the message Delivery outgoing rate * the specified Rate overdrive factor (percent) value. The Rate overdrive factor (percent) value is defined on the Message Processing Throttling Settings dialog box available from the Advanced page of the Host Properties dialog box.</p> |
| | <p>Delivery outgoing rate: Indicates the rate at which the messages are processed and sent to the recipients by the Orchestration engine or the Messaging engine of this host instance.</p> | <p>Msgs/Sec</p> | |

MONITORING THE BIZTALK SERVER

| | <p>Delivery throttling state:</p> <p>Indicates whether the system is throttling the message delivery i.e., indicates whether the XLANG message processing and outbound transports are affected or not.</p> | <p>Number</p> | <p>Indicates whether the system is throttling the message delivery i.e., indicates whether the XLANG message processing and outbound transports are affected or not.</p> <table border="1" data-bbox="935 338 1419 1066"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Not throttling</td> </tr> <tr> <td>1</td> <td>Throttling due to imbalanced message delivery rate (input rate exceeds output rate)</td> </tr> <tr> <td>3</td> <td>Throttling due to high in-process message count</td> </tr> <tr> <td>4</td> <td>Throttling due to process memory pressure</td> </tr> <tr> <td>5</td> <td>Throttling due to process memory pressure</td> </tr> <tr> <td>9</td> <td>Throttling due to high thread count</td> </tr> <tr> <td>10</td> <td>Throttling due to user override on delivery</td> </tr> </tbody> </table> | Value | State | 0 | Not throttling | 1 | Throttling due to imbalanced message delivery rate (input rate exceeds output rate) | 3 | Throttling due to high in-process message count | 4 | Throttling due to process memory pressure | 5 | Throttling due to process memory pressure | 9 | Throttling due to high thread count | 10 | Throttling due to user override on delivery |
|-------|---|---------------|--|-------|-------|---|----------------|---|---|---|---|---|---|---|---|---|-------------------------------------|----|---|
| Value | State | | | | | | | | | | | | | | | | | | |
| 0 | Not throttling | | | | | | | | | | | | | | | | | | |
| 1 | Throttling due to imbalanced message delivery rate (input rate exceeds output rate) | | | | | | | | | | | | | | | | | | |
| 3 | Throttling due to high in-process message count | | | | | | | | | | | | | | | | | | |
| 4 | Throttling due to process memory pressure | | | | | | | | | | | | | | | | | | |
| 5 | Throttling due to process memory pressure | | | | | | | | | | | | | | | | | | |
| 9 | Throttling due to high thread count | | | | | | | | | | | | | | | | | | |
| 10 | Throttling due to user override on delivery | | | | | | | | | | | | | | | | | | |

| | <p>High database session: Indicates whether the database session is within normal limits or not for this host instance.</p> | <p>Number</p> | <p>This measure reports any one of the following values to indicate whether the database session is within normal limits or not.</p> <table border="1" data-bbox="935 352 1421 583"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Database session count exceeds threshold</td> </tr> </tbody> </table> <p>The database session count is nothing but the number of concurrent MessageBox database connections being used. The threshold for database session count is initially set to the value specified for Database connections per CPU on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. This value is auto-tuned based on the database session usage of the process. If the number of concurrent database sessions exceeds this threshold at any time, host throttling is implemented.</p> | Value | State | 0 | Normal | 1 | Database session count exceeds threshold |
|-------|---|---------------|---|-------|-------|---|--------|---|--|
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | Database session count exceeds threshold | | | | | | | | |
| | <p>High database size: Indicates whether the size of the database is within normal limits or not for this host instance.</p> | <p>Number</p> | <p>This measure indicates any one of the following values while indicating whether the size of the database is within normal limits or not.</p> <table border="1" data-bbox="935 1209 1421 1440"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Database size has grown beyond threshold</td> </tr> </tbody> </table> <p>Database size is represented by the number of messages in the database queues that a host instance has published. This value is measured by the number of items in the queue tables for all hosts and the number of items in the spool and tracking tables. If a process is publishing to multiple queues, this counter reflects the weighted average of all the queues. If the threshold set for database size is violated, then throttling is implemented.</p> | Value | State | 0 | Normal | 1 | Database size has grown beyond threshold |
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | Database size has grown beyond threshold | | | | | | | | |

MONITORING THE BIZTALK SERVER

| | <p>High in-process message count:</p> <p>Indicates whether the In-process message count is within normal limits or not.</p> | Number | <p>This measure reports any one of the following values to indicate whether the In-process message count is within normal limits or not.</p> <table border="1" data-bbox="935 407 1419 638"><thead><tr><th>Value</th><th>State</th></tr></thead><tbody><tr><td>0</td><td>Normal</td></tr><tr><td>1</td><td>In-process message count exceeds limit</td></tr></tbody></table> <p>The in-process message count indicates the number of in-memory messages delivered to the XLANG engine or the outbound messaging engine that are not yet processed.</p> | Value | State | 0 | Normal | 1 | In-process message count exceeds limit |
|-------|--|--------|---|-------|-------|---|--------|---|--|
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | In-process message count exceeds limit | | | | | | | | |

MONITORING THE BIZTALK SERVER

| | <p>High message delivery rate:</p> <p>Indicates whether the message delivery rate is within normal limits or not.</p> | <p>Number</p> | <p>This measure reports any one of the following values to indicate whether the message delivery rate is within normal limits or not.</p> <table border="1" data-bbox="935 438 1419 667"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Message delivery rate exceeds the message processing rate</td> </tr> </tbody> </table> | Value | State | 0 | Normal | 1 | Message delivery rate exceeds the message processing rate |
|-------|--|---------------|---|-------|-------|---|--------|---|---|
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | Message delivery rate exceeds the message processing rate | | | | | | | | |
| | <p>High message publishing rate:</p> <p>Indicates whether the message publishing rate is within normal limits or not.</p> | <p>Number</p> | <p>This measure reports any one of the following values to indicate whether the message publishing rate is within normal limits or not.</p> <table border="1" data-bbox="935 963 1419 1192"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Message delivery rate exceeds the message processing rate</td> </tr> </tbody> </table> | Value | State | 0 | Normal | 1 | Message delivery rate exceeds the message processing rate |
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | Message delivery rate exceeds the message processing rate | | | | | | | | |

| | <p>High process memory: Indicates whether the process memory is within normal limits or not.</p> | <p>Number</p> | <p>This measure reports any one of the following values to indicate whether the process memory is within normal limits or not.</p> <table border="1" data-bbox="935 390 1409 621"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Process memory exceeds threshold</td> </tr> </tbody> </table> <p>Process memory consumption is the maximum of the process's working set size and the total space allocated for the page file for the process. The threshold for process memory consumption is initially set to the value specified for Process memory usage on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. If a percentage value is specified, it is computed based on the available memory to commit.</p> | Value | State | 0 | Normal | 1 | Process memory exceeds threshold |
|-------|---|---------------|--|-------|-------|---|--------|---|----------------------------------|
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | Process memory exceeds threshold | | | | | | | | |
| | <p>High system memory: Indicates whether the system memory is within normal limits or not.</p> | <p>Number</p> | <p>This measure reports any one of the following values to indicate whether the system memory is within normal limits or not.</p> <table border="1" data-bbox="935 1150 1409 1381"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>System memory exceeds threshold</td> </tr> </tbody> </table> | Value | State | 0 | Normal | 1 | System memory exceeds threshold |
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | System memory exceeds threshold | | | | | | | | |

MONITORING THE BIZTALK SERVER

| | <p>High thread count: Indicates whether the thread count is within normal limits or not for this host instance.</p> | Number | <p>This measure reports any one of the following values to indicate whether the thread count is within normal limits or not.</p> <table border="1" data-bbox="935 359 1425 552"> <thead> <tr> <th>Value</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Normal</td> </tr> <tr> <td>1</td> <td>Thread count exceeds threshold</td> </tr> </tbody> </table> <p>The thread count indicates the number of threads being used in the process. The threshold for this count is initially set to the value specified for Threads per CPU on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. This value is auto-tuned depending on the thread requirements of the current process. If the number of threads in the process exceeds this threshold value at any point in time, host throttling is implemented.</p> | Value | State | 0 | Normal | 1 | Thread count exceeds threshold |
|-------|--|--------|---|-------|-------|---|--------|---|--------------------------------|
| Value | State | | | | | | | | |
| 0 | Normal | | | | | | | | |
| 1 | Thread count exceeds threshold | | | | | | | | |
| | <p>Thread count: Indicates the number of thread being used in the process.</p> | Number | | | | | | | |
| | <p>Thread count threshold: Indicates the current threshold for the number of threads in the process.</p> | Number | <p>The threshold for the thread count is initially set to the value specified for Threads per CPU on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. This value is auto-tuned depending on the thread requirements of the current process. If the number of threads in the process exceeds this threshold value at any point in time, host throttling is implemented.</p> | | | | | | |
| | <p>Database size: Indicates the number of messages in the database queues that this process has published.</p> | Number | <p>This value is measured by the number of items in the queue tables for all hosts and the number of items in the spool and tracking tables. If a process is publishing to multiple queues, this counter reflects the weighted average of all the queues.</p> | | | | | | |

| | | | |
|--|---|--------|--|
| | <p>Database session:</p> <p>Indicates the number of concurrent message box database connections that is being used.</p> | Number | |
| | <p>Process memory usage:</p> <p>Indicates the memory consumption of the process.</p> | MB | |
| | <p>Process memory usage threshold:</p> <p>Indicates the current threshold for the memory consumption of the process.</p> | MB | <p>This threshold value is initially set to the value specified for the process memory consumption on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. If a percentage value is specified, the threshold value is computed based on the available memory to commit.</p> |

4.2.1.3 BT File Receive Adapter Test

The file receive adapter is used to read messages from files and submit them to the server. The receive adapter reads the file and creates a BizTalk message object, so that BizTalk server can process the message. While reading from the file, the adapter locks the file to ensure that no modifications can be made to the file content. The file receive adapter does **not** pick up read-only files or system files.

This test reports how efficient the file receive adapter on each host instance is. The test monitors the inflow of messages to the file receive adapter, measures the load on the adapter, and reveals how well the adapter handled the load; lock failures encountered by the adapter while attempting to read files are also revealed by this test, so that reasons for the same can be diagnosed.

| | |
|---------------------------------|--|
| Purpose | This test reports how efficient the file receive adapter on each host instance is. The test monitors the inflow of messages to the file receive adapter, measures the load on the adapter, and reveals how well the adapter handled the load; lock failures encountered by the adapter while attempting to read files are also revealed by this test, so that reasons for the same can be diagnosed. |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes received: Indicates the total number of bytes received by the file receive adapter on this host instance. | Bytes | The counter is incremented after a message is completely read by the adapter from the file system. |
| | Bytes received: Indicates the rate at which bytes were received by the file receive adapter on this host instance. | Bytes/Sec | |
| | Messages received: Indicates the number of messages received by the file receive adapter on this host instance. | Number | The counter is incremented after a message is completely read by the file receive adapter from the file system. |
| | Messages received: Indicates the rate at which messages were received by the file receive adapter on this host instance. | Msgs/Sec | The counter applies only to messages that have been completely read by the file receive adapter from the file system. Ideally, the value of this measure should be high. A low value indicates that the file receive adapter is not reading files quickly. Further investigation may be required to diagnose the root-cause of the slowdown. |
| | Lock failures: Indicates the number of times the file receive adapter on this host instance failed to lock the file. | Number | Ideally, the value of this measure should be 0. A non-zero value indicates a lock failure. This in turn implies that the adapter could not prevent changes from being made to one/more files that were being read. |

4.2.1.4 BT File Send Adapter Test

The File send adapter transmits messages from the message box database to a specified destination address (URL). You define the URL, which is a file path and file name, by using wildcard characters related to the message context properties. The File send adapter resolves the wildcard characters to the actual file name before writing the message to the file.

When writing a message to a file, the File send adapter gets the message content from the body part of the BizTalk message object. The File send adapter ignores other message parts in the BizTalk Message object. After the File adapter writes the message to a file, it deletes the message from the MessageBox database. The File adapter writes files to the file system either directly or by using the file system cache, which can improve performance, particularly for large files.

This test monitors the outflow of data and messages from the file send adapter on each host instance and reports the load on the adapter and the slowdowns (if any) suffered by the adapter while processing the load.

| | | | |
|---|---|-------------------------|--|
| Purpose | Monitors the outflow of data and messages from the file send adapter on each host instance and reports the load on the adapter and the slowdowns (if any) suffered by the adapter while processing the load | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes sent: Indicates the total number of bytes sent by the file send adapter on this host instance. | Bytes | The counter is incremented only for messages that have been completely written to file system. |
| | Bytes sent: Indicates the rate at which bytes were sent by the file send adapter on this host instance. | Bytes/Sec | The counter applies only to messages that have been completely written to file system. |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|----------|---|
| | <p>Messages sent:</p> <p>Indicates the number of messages sent by the file send adapter on this host instance.</p> | Number | The counter is incremented only for messages that have been completely written to file system. |
| | <p>Messages sent:</p> <p>Indicates the rate at which messages were sent by the file send adapter on this host instance.</p> | Msgs/Sec | <p>The counter applies only to messages that have been completely written to file system.</p> <p>Ideally, the value of this measure should be high. A low value indicates that the file send adapter is experiencing delays while writing files to the file system. Further investigation may be required to diagnose the root-cause of the slowdown.</p> |

4.2.1.5 BT FTP Receive Adapter Test

The FTP receive adapter enables you to move data from an FTP server to BizTalk Server.

Key features of the FTP receive adapter are:

- Pulling files from the FTP server on demand
- Running polls based on a configurable schedule
- Polling the FTP server and sending data directly to BizTalk Server
- Specifying the FTP server as an IP address, port, password, and host name
- Guaranteed file delivery

The FTP receive adapter also works with the BizTalk Administration console and BizTalk Explorer to configure and administer each receive function, which is composed of the following configuration items:

- Poll interval to run an FTP command (for example, 60 minutes).
- Information with which to route the document to a specific BizTalk send port or receive location.

The FTP receive adapter does **not** support receiving files from a **partitioned data set**.

With the help of this test, you can measure the current load on the FTP receive adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing.

| | |
|---------------------------------|--|
| Purpose | Measure the current load on the FTP receive adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes received: Indicates the total number of bytes received by the FTP receive adapter on this host instance. | Bytes | The counter is incremented after a message is completely read by the FTP receive adapter from the FTP server. |
| | Bytes received: Indicates the rate at which bytes were received by the FTP receive adapter on this host instance. | Bytes/Sec | The counter applies only to messages that have been completely read by the FTP receive adapter from the FTP server. |
| | Messages received: Indicates the number of messages received by the FTP receive adapter on this host instance. | Number | The counter applies only to messages that have been completely read by the FTP receive adapter from the FTP server. This measure is a good indicator of the load on the FTP receive adapter. |
| | Messages received: Indicates the rate at which messages were received by the FTP receive adapter on this host instance. | Msgs/Sec | The counter applies only to messages that have been completely read by the FTP receive adapter from the FTP server. Ideally, a value of this measure should be high. A low value indicates that the FTP receive adapter is experiencing delays while moving files and data from the FTP server to the BizTalk server. Further investigation may be required to diagnose the root-cause of the slowdown. |

4.2.1.6 BT FTP Send Adapter Test

The FTP send adapter enables you to move data from BizTalk Server to an FTP server.

Key features of the FTP send adapter are:

- Ability to run sends on demand

MONITORING THE BIZTALK SERVER

- Guaranteed delivery

With the help of this test, you can measure the current load on the FTP send adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing.

| | | | |
|---|---|-------------------------|--|
| Purpose | Measure the current load on the FTP send adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes sent: Indicates the total number of bytes sent by the FTP send adapter on this host instance. | Bytes | The counter is incremented only for messages that have been written to the destination FTP server. |
| | Bytes sent: Indicates the rate at which bytes were sent by the FTP send adapter on this host instance. | Bytes/Sec | The counter applies only to messages that have been written to the destination FTP server. |
| | Messages sent: Indicates the number of messages sent by the FTP send adapter on this host instance. | Number | The counter is incremented only for messages that have been written to the destination FTP server. |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|----------|---|
| | <p>Messages sent:</p> <p>Indicates the rate at which messages were sent by the FTP send adapter on this host instance.</p> | Msgs/Sec | <p>The counter applies only to messages that have been written to destination FTP server.</p> <p>Ideally, a value of this measure should be high. A low value indicates that the FTP send adapter is experiencing delays while writing files to the destination FTP server. Further investigation may be required to diagnose the root-cause of the slowdown.</p> |
|--|---|----------|---|

4.2.1.7 BT Http Receive Adapter Test

The HTTP adapter is used to exchange information between the BizTalk server and an application by means of the HTTP protocol. HTTP is the primary protocol for inter-business message exchange. Applications can send messages to a server by sending HTTP POST or HTTP GET requests to a specified HTTP URL. The HTTP receive adapter is an Internet Information Services (IIS) Internet Server Application Programming Interface (ISAPI) extension that the IIS process hosts, and controls the receive locations that use the HTTP adapter. The receive location for the HTTP receive adapter is a distinct URL configured through BizTalk Explorer.

Using this test, you can monitor the flow of messages to and from the HTTP receive adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP receive adapter of a host instance, and evaluate the load processing ability of that adapter.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitor the flow of messages to and from the HTTP receive adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP receive adapter of a host instance, and evaluate the load processing ability of that adapter | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|----------|--|
| | <p>Messages received:</p> <p>Indicates the total number of HTTP requests received by the HTTP receive adapter on this host instance.</p> | Number | The counter is incremented after a request message is completely read by the HTTP receive adapter from the HTTP client. |
| | <p>Messages received:</p> <p>Indicates the rate at which the HTTP requests are received by the HTTP receive adapter on this host instance.</p> | Msgs/Sec | <p>The counter applies only to request messages that have been completely read by the HTTP receive adapter from the HTTP client.</p> <p>Ideally, the value of this measure should be high. A low value indicates that the HTTP receive adapter is experiencing delays while accepting requests from the HTTP client. Further investigation may be required to diagnose the root-cause of the slowdown.</p> |
| | <p>Messages sent:</p> <p>Indicates the total number of HTTP responses sent by the HTTP receive adapter on this host instance.</p> | Number | The counter is incremented only for response messages that have been successfully sent to HTTP clients. |
| | <p>Messages sent:</p> <p>Indicates the rate at which messages were sent by the FTP send adapter on this host instance on this host instance.</p> | Number | <p>The counter applies only to response messages that have been successfully sent to HTTP clients.</p> <p>Ideally, the value of this measure should be high. A low value indicates that the HTTP receive adapter is experiencing delays while responding to requests from HTTP clients. Further investigation may be required to diagnose the root-cause of the slowdown.</p> |

4.2.1.8 BT Http Send Adapter Test

The HTTP send adapter gets messages from BizTalk Server and sends them to a destination URL on an HTTP POST request. The HTTP send adapter gets the message content from the body part of the BizTalk Message object. The HTTP send adapter ignores all other parts of the BizTalk Message object.

Using this test, you can monitor the flow of messages to and from the HTTP send adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP send adapter of a host instance, and evaluate the load processing ability of that adapter.

| | |
|---------------------------|---|
| Purpose | Monitor the flow of messages to and from the HTTP send adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP send adapter of a host instance, and evaluate the load processing ability of that adapter |
| Target of the test | A BizTalk Server 2010 |
| Agent | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|--|
| deploying the test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Messages received: Indicates the total number of HTTP response messages received by the HTTP send adapter on this host instance. | Number | The counter is incremented after a response message is completely read by the HTTP send adapter from HTTP servers. |
| | Messages received: Indicates the rate at which HTTP response messages are received by the HTTP send adapter on this host instance. | Msgs/Sec | The counter applies only to response messages that have been completely read by the HTTP send adapter. Ideally, the value of this measure should be high. A low value indicates that the HTTP send adapter is experiencing delays while receiving messages from the BizTalk server. Further investigation may be required to diagnose the root-cause of the slowdown. |
| | Messages sent: Indicates the total number of HTTP requests sent by the HTTP send adapter on this host instance to the destination URL. | Number | The counter is incremented only for request messages that have reached the destination URL. |
| | Messages sent: Indicates the rate at which HTTP requests were sent by the HTTP send adapter on this host instance to the destination URL. | Msgs/Sec | The counter applies only to request messages that have reached the destination URL. Ideally, the value of this measure should be high. A low value indicates that the HTTP send adapter is experiencing delays while sending messages to the destination URL on an HTTP POST request. Further investigation may be required to diagnose the root-cause of the slowdown. |

4.2.1.9 BT Messaging Latency Test

One of the key services provided by the BizTalk server messaging engine is the mechanism for communicating across the applications that a business process uses. As the first steps towards enabling this communication, the messaging engine receives messages from a source application through a **receive adapter**. The message is then processed through a **receive pipeline** and delivered into a database called the **MessageBox**. Depending upon the nature of the messages delivered to the MessageBox, the messaging engine dispatches the messages to their appropriate orchestrations; each orchestration then takes whatever action the business process requires. The result of this processing is typically another message, produced by the orchestration and saved in the MessageBox. This message, in turn, is processed by a **send pipeline**, and sent out to the application for which it is destined, using a **send adapter**.

The health of the messaging engine relies heavily on how quickly messages are processed at each step of the electronic data exchange that has been described above. Administrators should be promptly notified of even the slightest of latencies in this communication, so that they can take the measures to curb it before it causes a significant delay in the delivery of messages to the target. The **Messaging latency** test serves this purpose.

The test closely observes the time taken by the messaging engine to send a message to the MessageBox and to send a message in the MessageBox to the target destination, and proactively alerts administrators to delays; this way, administrators will not only be able to promptly detect latencies experienced by the messaging engine, but will also be able to pin-point where the delay originated - while delivering messages to the MessageBox? or while delivering messages to the destination application?

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Closely observes the time taken by the messaging engine to send a message to the MessageBox and to send a message in the MessageBox to the target destination, and proactively alerts administrators to delays; this way, administrators will not only be able to promptly detect latencies experienced by the messaging engine, but will also be able to pin-point where the delay originated - while delivering messages to the MessageBox? or while delivering messages to the destination application? | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|------|---|
| | <p>Inbound Latency:</p> <p>Indicates the time taken by the messaging engine to receive a document from the receive adapter and publish it to the MessageBox.</p> | Secs | Ideally, the value of this measure should be low. A very high value is indicative of a slowdown while publishing documents to the MessageBox. |
| | <p>Outbound Latency:</p> <p>Indicates the time taken by the messaging engine to receive a document from the MessageBox and send it to the adapter.</p> | Secs | Ideally, the value of this measure should be low. A very high value is indicative of a slowdown in publishing documents to the destination. |
| | <p>Request-Response Latency:</p> <p>Indicates the total time taken by the messaging engine to process a request document from the adapter and send back a response document to the adapter.</p> | Secs | Ideally, the value of this measure should be low. A very high value is indicative of delays in communication across applications. |

4.2.1.10 BT Msmq Receive Adapter Test

With the BizTalk Server Adapter for MSMQ (the MSMQ adapter), you can send and receive messages to Microsoft Message Queuing (also known as MSMQ) queues using Microsoft BizTalk Server. The MSMQ adapter works with transactional and non-transactional, public and private, and local and remote queues. Additionally, the MSMQ adapter provides large (greater than 4 MB) message support and gives you access to Message Queuing features such as messaging over HTTP and multi-cast messaging. The key features of the MSMQ adapter are:

- Can be configured to deliver messages in order.
- Provides large message support by breaking the message into parts, accumulating the parts in memory, and delivering the parts in order to the destination (more memory intensive than MSMQT).
- Provides better performance than MSMQT.
- Enables other non-BizTalk applications to use MSMQ services at the same time on the same computer.
- Requires intermediate storage of MSMQ queues. Inbound messages are written to the MSMQ queue and then picked up from the MSMQ queue by the MSMQ adapter.

By continuously tracking the messages and data received and processed by the MSMQ receive adapter for every host instance on the BizTalk server, administrators can receive an overview of the load on the adapter, and will be able to accurately judge its processing ability. This test does just that.

| | |
|----------------|---|
| Purpose | Continuously tracks the messages and data received and processed by the MSMQ receive adapter for every host instance on the BizTalk server, and provides administrators with an |
|----------------|---|

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|---|
| | overview of the load on the adapter, and will be able to accurately judge its processing ability | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes received: Indicates the total number of bytes received by the MSMQ receive adapter on this host instance. | Bytes | The counter is incremented after a message is completely read by the MSMQ receive adapter from the source queue. |
| | Bytes received: Indicates the rate at which bytes were received by the MSMQ receive adapter on this host instance. | Bytes/Sec | The counter applies only to messages that have been completely read by the MSMQ receive adapter from the source queue. |
| | Messages received: Indicates the number of messages received by the MSMQ receive adapter on this host instance. | Number | The counter is incremented after a message is completely read by the MSMQ receive adapter from the source queue. This measure is a good indicator of the load on the MSMQ receive adapter. |
| | Messages received: Indicates the rate at which messages were received by the MSMQ receive adapter on this host instance. | Msgs/Sec | The counter applies only to messages that have been completely read by the MSMQ receive adapter from the source queue. Ideally, the value of this measure should be high. A low value indicates that the MSMQ receive adapter is experiencing delays while reading messages from the source queue. Further investigation may be required to diagnose the root-cause of the slowdown. |

4.2.1.11 BT Msmq Send Adapter Test

With the BizTalk Server Adapter for MSMQ (the MSMQ adapter), you can send and receive messages to Microsoft Message Queuing (also known as MSMQ) queues using Microsoft BizTalk Server. The MSMQ adapter works with transactional and non-transactional, public and private, and local and remote queues. Additionally, the MSMQ adapter provides large (greater than 4 MB) message support and gives you access to Message Queuing features such as messaging over HTTP and multi-cast messaging. The key features of the MSMQ adapter are:

- Can be configured to deliver messages in order.
- Provides large message support by breaking the message into parts, accumulating the parts in memory, and delivering the parts in order to the destination (more memory intensive than MSMQT).
- Provides better performance than MSMQT.
- Enables other non-BizTalk applications to use MSMQ services at the same time on the same computer.
- Requires intermediate storage of MSMQ queues. Inbound messages are written to the MSMQ queue and then picked up from the MSMQ queue by the MSMQ adapter.

By continuously tracking the messages and data sent by the MSMQ send adapter for every host instance on the BizTalk server, administrators can receive an overview of the load on the adapter, and will be able to accurately judge its processing ability. This test does just that.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Measure the current load on the MSMQ send adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|-----------|--|
| | Bytes sent: Indicates the total number of bytes sent by the MSMQ send adapter on this host instance. | Bytes | The counter is incremented only for messages that have reached the destination queue. |
| | Bytes sent: Indicates the rate at which bytes were sent by the MSMQ send adapter on this host instance. | Bytes/Sec | The counter applies only to messages that have reached the destination queue. |
| | Messages sent: Indicates the number of messages sent by the MSMQ send adapter on this host instance. | Number | The counter is incremented only for messages that have reached the destination queue. |
| | Messages sent: Indicates the rate at which messages were sent by the MSMQ send adapter on this host instance. | Msgs/Sec | The counter applies only to messages that have reached the destination queue. Ideally, the value of this measure should be high. A low value indicates that the MSMQ send adapter is experiencing delays while sending messages to the destination queue. Further investigation may be required to diagnose the root-cause of the slowdown. |

4.2.1.12 BT Pop3 Adapter Test

The Post Office Protocol 3 (POP3) adapter is used to retrieve data from a server that houses POP3 mailboxes into a BizTalk Server by means of the POP3 protocol. The POP3 adapter consists of only one adapter, a receive adapter. This receive adapter controls the receive locations that use the POP3 adapter.

The POP3 receive adapter retrieves e-mail from a specified mailbox on a specified POP3 server. By default, the POP3 receive adapter applies MIME processing to the e-mail messages that it downloads and submits these messages to BizTalk Server as multipart BizTalk messages. The POP3 receive adapter can receive and process e-mail in the following formats:

- Plain text
- MIME encoded
- MIME encrypted
- MIME encoded and signed
- MIME encrypted and signed

To monitor the session and message load on the POP3 adapter so that, overload conditions and processing bottlenecks are accurately identified, use the **POP3 adapter** test.

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|---|
| Purpose | To monitor the session and message load on the POP3 adapter so that, overload conditions and processing bottlenecks are accurately identified | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active sessions: Indicates the number of open POP3 connections that the POP3 adapter on this host instance is currently managing. | Number | This is a good indicator of the session load on the adapter. |
| | Bytes received: Indicates the total number of bytes downloaded by the POP3 adapter on this host instance from a mail server. | Bytes | This is a good indicator of the data load on the adapter. |
| | Bytes received: Indicates the rate at which bytes that the POP3 adapter on this host instance downloaded from a mail server. | Bytes/Sec | A consistent decrease in this value could indicate a processing bottleneck. |
| | Messages received: Indicates the number of messages that the POP3 adapter on this host instance downloaded from the mail server. | Number | This is a good indicator of the load on the adapter. |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|----------|---|
| | <p>Messages received:</p> <p>Indicates the rate at which the POP3 adapter on this host instance downloaded messages from the mail server.</p> | Msgs/Sec | A consistent decrease in this value could indicate a processing bottleneck. |
|--|--|----------|---|

4.2.1.13 BT SMTP Adapter Test

The Simple Mail Transfer Protocol (SMTP) adapter is used to exchange information between a BizTalk Server and other applications by means of the SMTP protocol. BizTalk Server can send messages to other applications by creating an e-mail message and delivering it to a specified e-mail address. The SMTP adapter consists of only one adapter, a send adapter. The send adapter controls the send ports that use the SMTP adapter. Internally, the SMTP send adapter creates an SMTP-based e-mail message and sends it to a target e-mail address. The target e-mail address is a property of the SMTP adapter. The SMTP send adapter gets messages from the server and posts them to an SMTP server that sends them to e-mail recipients.

Using this test, you can figure out how quickly the SMTP send adapter sends out messages to other applications, and thus promptly detect slowdowns in message delivery.

| | | | |
|---|---|-------------------------|---|
| Purpose | Helps figure out how quickly the SMTP send adapter sends out messages to other applications, and thus promptly detects slowdowns in message delivery | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Messages sent:</p> <p>Indicates the total number of messages sent by the SMTP adapter on this host instance to the target e-mail address.</p> | Number | The counter is incremented only for messages that have been transmitted to the SMTP server. |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|----------|---|
| | <p>Messages sent:</p> <p>Indicates the rate at which messages were sent by the SMTP adapter on this host instance to the target e-mail address.</p> | Msgs/Sec | <p>The counter applies only to messages that have been transmitted to the SMTP server.</p> <p>Ideally, the value of this measure should be high. A low value indicates that the SMTP send adapter is experiencing delays while sending messages to the target e-mail address. Further investigation may be required to diagnose the root-cause of the slowdown.</p> |
|--|--|----------|---|

4.2.1.14 BT Soap Receive Adapter Test

The SOAP adapter enables you to publish orchestrations as Web services and consume external Web services. The SOAP adapter consists of two adapters — a send adapter and receive adapter.

The SOAP receive adapter is used to receive Web service requests. The SOAP receive adapter creates a BizTalk Message object, and promotes the associated properties to the message context.

This test enables you to determine the web service request load on the SOAP receive adapter at any given point in time, and helps you assess the processing capability of the adapter.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Enables you to determine the web service request load on the SOAP receive adapter at any given point in time, and helps you assess the processing capability of the adapter | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|----------|---|
| | <p>Messages received:</p> <p>Indicates the total number of messages that are received by the SOAP receive adapter on this host instance.</p> | Number | The counter is incremented after a request message is completely read by the adapter from the SOAP client. |
| | <p>Messages received:</p> <p>Indicates the rate at which the messages are received by the SOAP receive adapter on this host instance.</p> | Msgs/Sec | <p>The counter applies only to request messages that have been completely read by the adapter from the SOAP client.</p> <p>Ideally, the value of this measure should be high. A consistent decrease in this value indicates that the SOAP receive adapter is experiencing delays while reading messages from the SOAP client. Further investigation may be required to diagnose the root-cause of the slowdown.</p> |

4.2.1.15 BT Soap Send Adapter Test

The SOAP send adapter is used to call a web service. The SOAP send adapter reads the message context on the BizTalk Message object to get the proxy name and calls the associated external Web service proxy.

Monitor the load on the SOAP send adapter and be proactively alerted to processing bottlenecks in the adapter with the help of the **Soap send adapter** test.

| | |
|---|---|
| Purpose | Monitor the load on the SOAP send adapter and be proactively alerted to processing bottlenecks in the adapter |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored |

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|---|------------------|--|
| | Messages sent: Indicates the total number of messages that are sent by the SOAP send adapter on this host instance. | Number | The counter is incremented only for messages that have reached the destination URL. |
| | Messages sent: Indicates the rate at which the messages are sent by the SOAP send adapter on this host instance. | Msgs/Sec | The counter applies only to messages that have reached the destination URL. Ideally, the value of this measure should be high. A consistent decrease in this value indicates that the SOAP send adapter is experiencing delays while writing messages to the destination URL. Further investigation may be required to diagnose the root-cause of the slowdown. |

4.2.1.16 BT Sql Receive Adapter Test

The SQL adapter exchanges data between the BizTalk Server and a SQL Server database. You can use the SQL adapter to poll data from one or more data tables and transmit the data as one or more XML messages to BizTalk Server. You can also use the SQL adapter to move large amounts of data to or from the SQL Server database as part of a BizTalk Server messaging or orchestration solution. In addition, you can use the SQL adapter to insert, update, and delete data in SQL Server tables by using SQL updategrams or by invoking stored procedures. The SQL adapter consists of two adapters—a receive adapter and a send adapter.

The SQL receive adapter is a polling adapter that periodically polls for SQL result sets.

This test monitors the load on the **SQL Receive Adapter** and proactively alerts you to potential overload conditions.

| | |
|---|---|
| Purpose | Monitors the load on the SQL Receive Adapter and proactively alerts you to potential overload conditions |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored |

MONITORING THE BIZTALK SERVER

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|---|------------------|---|
| | Messages received: Indicates the total number of messages that are read by the SQL receive adapter from the SQL server. | Number | A high value could indicate an overload condition. |
| | Messages received: Indicates the rate at which the messages are read by the SQL receive adapter from the SQL server. | Msgs/Sec | A consistent decrease in the value of this measure points you to current/potential bottlenecks in the processing of messages. |

4.2.1.17 BT File Send Adapter Test

The SQL adapter exchanges data between the BizTalk Server and a SQL Server database. You can use the SQL adapter to poll data from one or more data tables and transmit the data as one or more XML messages to BizTalk Server. You can also use the SQL adapter to move large amounts of data to or from the SQL Server database as part of a BizTalk Server messaging or orchestration solution. In addition, you can use the SQL adapter to insert, update, and delete data in SQL Server tables by using SQL updategrams or by invoking stored procedures. The SQL adapter consists of two adapters—a receive adapter and a send adapter.

The SQL send adapter is used to send dynamically created updategrams or dynamically invoked stored procedures to SQL Server. An updategram is an XML fragment that inserts, updates, or deletes data in a SQL Server database by mapping XML nodes against database tables and columns. SQL Server returns an optional response document after the updategram completes, which contains the success status of the update. If a failure occurs during the update, the SQL adapter throws an exception that the BizTalk Messaging Engine handles. When the SQL send adapter is configured to invoke a stored procedure, it returns any results in the form of a single XML-formatted record set.

This test monitors the load on the **SQL Send Adapter** and proactively alerts you to potential overload conditions and processing bottlenecks.

| | |
|---------------------------------|---|
| Purpose | Monitors the load on the SQL Send Adapter and proactively alerts you to potential overload conditions and processing bottlenecks |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Messages sent: Indicates the total number of messages that are sent by the SQL send adapter to the destination SQL table in the SQL server database. | Number | |
| | Messages sent: Indicates the rate at which the messages are sent by the SQL send adapter to the destination SQL table in the SQL server database. | Msgs/Sec | A consistent decrease in the value of this measure points you to current/potential bottlenecks in the processing of messages. |

4.2.2 The Message Box Layer

The heart of the publish/subscribe engine in Microsoft BizTalk Server is the MessageBox database. The MessageBox is made up of two components: one or more Microsoft SQL Server databases and the Messaging Agent. The SQL Server database provides the persistence store for many things including messages, message parts, message properties, subscriptions, orchestration state, tracking data, host queues for routing, and others. The BizTalk Server group may have one or more MessageBox databases into which it publishes messages and from which subscribers to those messages extract messages.

The database provides some of the logic related to routing messages and fulfilling subscriptions. The Message Agent, however, is the component that encapsulates and abstracts the database component and is the interface used by BizTalk Server to interact with the MessageBox. The Message Agent is a Component Object Model (COM) component that provides interfaces for publishing messages, subscribing to messages, retrieving messages, and so on. This interface is the only mechanism used by other BizTalk Server components, including the adapter framework and orchestrations, to interact with the MessageBox.

Using the tests mapped to this layer you can monitor the health of the BizTalk server MessageBox and the efficiency of the SQL Server agent jobs.

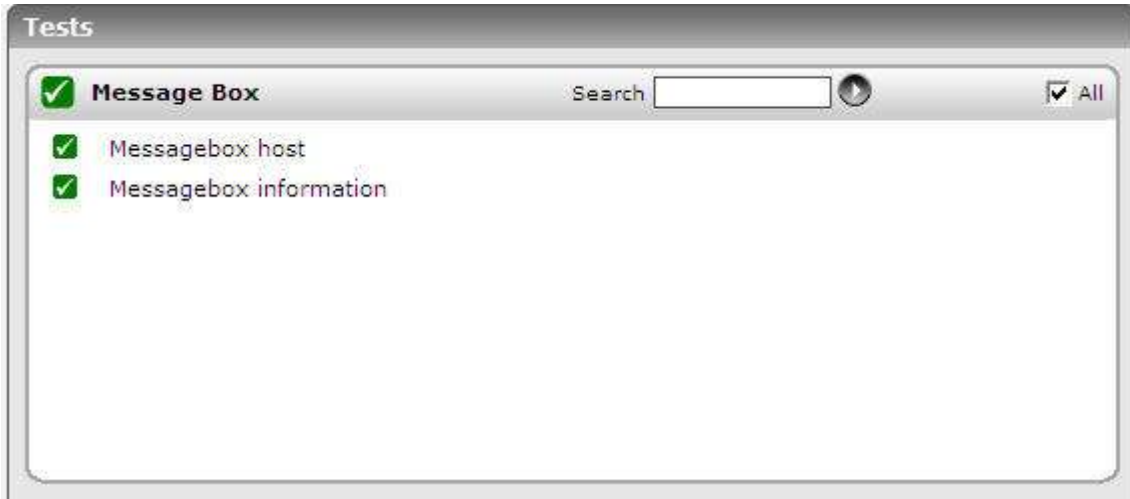


Figure 4.8: The tests mapped to the Message Box layer

4.2.2.1 BT Messagebox Host Test

The first time you configure a BizTalk server, the following set of tables are created in the MessageBox database for a BizTalkServerApplicationHost:

- The BizTalkApplicationQ
- The BizTalkServerApplicationQ_Suspended
- The BizTalkServerApplicationQ_Scheduled
- The InstanceStateMessageReferences_BizTalkServerApplication

BizTalk uses these tables to keep **references** of all the messages that are “live” in the system . That is: Messages with active subscriptions, suspended messages, and *awaiting messages* associated to each host.

The word **references** implies that the host tables are only pointers to the **Spool** table, but the real messages itself are saved in another set of tables (messageparts, parts and fragments).

This test monitors the number of message references in the host queue tables, and proactively alerts administrators to the following:

- A sudden/consistent increase in the length of the host queues
- Too many message references in the suspended queue

| | |
|---------------------------------|---|
| Purpose | This test monitors the number of message references in the host queue tables, and proactively alerts administrators to the following: <ul style="list-style-type: none"> • A sudden/consistent increase in the length of the host queues • Too many message references in the suspended queue |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Message references in Instance Queue: Indicates the number of message references in the instance state queue of this host instance. | Number | The State Queue table holds the list of messages that have been processed by an instance but will be needed later . When an orchestration uses the State Queue, it is usually because the orchestration performed some operations on a message, persisted the message, and might need the message later. This is normal operation, and you should take this into account when determining correct sizing of the State Queue |
| | Instances of Host Queue: Indicates the number of instances of the host queue for this host instance. | Number | |
| | Messages in Host Queue: Indicates the number of messages in the host queue of this host instance. | Number | Generally, this queue should not grow too large. The length of the queue indicates the number of messages waiting to be processed. A large number means you could have a backlog. |
| | Suspended Messages in Host Queue: Indicates the number of suspended messages for this host instance. | Number | When a message gets suspended it remains in the messagebox until resume or terminate actions occurs. So, if the suspended queue keeps growing, the performance of the BizTalk server will continue to get affected. A suspended message can be due, for example, to parsing errors, serialization errors, failed transmissions, or the inability to find a subscription. |

4.2.2.2 BT Messagebox Information Test

The BizTalk server includes certain SQL agent jobs to assist administrators in managing the BizTalk server databases.

MONITORING THE BIZTALK SERVER

Using this test, you can monitor the time taken to perform each of these SQL agent jobs so that, jobs that took too long to complete can be instantly identified.

| | | | |
|---|---|-------------------------|--|
| Purpose | Monitors the time taken to perform each of these SQL agent jobs so that, jobs that took too long to complete can be instantly identified | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Dead Processes Cleanup: Indicates the time taken by the MessageBox_DeadProcesses_Cleanup_BizTalkMsgBoxDb job for this host instance to complete. | Secs | This job detects when a BizTalk Server host instance (NT service) has stopped and releases all work that was being done by that host instance so that it can be worked on by another host instance. |
| | Cleanup Messages: Indicates the time taken by the MessageBox_Message_Cleanup_BizTalkMsgBoxDb job for this host instance to complete its work. | Secs | This job removes all messages that are no longer being referenced by any subscribers in the BizTalk MessageBox (BizTalkMsgBoxDb) database tables. Note: This is an unscheduled job which is started by the MessageBox_Message_ManageRefCountLog_BizTalkMsgBoxDb job. Do not manually start this job. |
| | Total Instances: Indicates the total number of host instances that exist within a message box. | Number | |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|---------------|--|
| | <p>Cleanup Message Parts: Indicates the time taken by the MessageBox_Parts_Cleanup_BizTalkMsgBoxDb job for this host instance to complete its work.</p> | <p>Secs</p> | <p>This job removes all message parts that are no longer being referenced by any messages in the BizTalk MessageBox (BizTalkMsgBoxDb) database tables. All messages are made up of one or more message parts, which contain the actual message data.</p> |
| | <p>Spool Size: Indicates the size of the spool that is available on a particular message box in this host instance.</p> | <p>Number</p> | <p>The primary measure of sustainability over time is that a backlog is not allowed to grow indefinitely. In other words, over time, there must be a balance between the high and low peak throughput levels so that the MessageBox database is able to maintain a constant and manageable average backlog. The primary measure of backlog is the depth of the spool table.</p> <p>The message bodies are handled via a set of tables represented by the spool table.</p> <p>The Spool can start growing for multiple reasons. One reason for Spool growth is if the application queues are growing. Application queues host in-flight transition data. They could grow due to various reasons like downstream bottlenecks and/or resource contention.</p> <p>If the application queues are small and the Spool is still large, verify that the purge jobs are keeping up. Ensure that the SQL-Agent Service is running and then verify that the following jobs are successfully completing:</p> <ul style="list-style-type: none"> • MessageBox_Message_Cleanup_BizTalkMsgBoxDb • MessageBox_Parts_Cleanup_BizTalkMsgBoxDb <p>One reason for this is if the SQL-Server machine is experiencing severe CPU contention, impacting the ability of the purge jobs to keep up due to CPU starvation.</p> |
| | <p>Tracked Messages: Indicates the time taken by the DTA Purge and Archive job of this host instance to complete its execution.</p> | <p>Secs</p> | <p>This job automatically archives data in the BizTalk Tracking (BizTalkDTADB) database and purges obsolete data.</p> |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|--------|---|
| | Tracking Data Size: Indicates the size of the data table that is tracked from the message available for this host instance. | Number | As BizTalk Server processes more and more data on your system, the BizTalk Tracking (BizTalkDTADB) database continues to grow in size. Unchecked growth decreases system performance and may generate errors in the Tracking Data Decode Service (TDDS). In addition to general tracking data, tracked messages can also accumulate in the MessageBox database, causing poor disk performance. This implies that ideally the value of this measure should be low. By archiving and purging data from the BizTalk Tracking database, you can maintain a healthy system, as well as keep your tracking data archived for future use. |
| | Tracking Pool Cleanup: Indicates the time taken to purge inactive pools in the tracking database tables so as to free database space. | Secs | |

4.2.3 The Orchestration Engine Layer

An orchestration is a flexible, powerful tool for representing an executable business process based on XLANG/s language. At run time, the BizTalk Orchestration Engine executes XLANG/s files that are produced by BizTalk Orchestration Designer. Orchestration Designer is a rich graphical tool for visually designing business processes. It generates XLANG/s files that have an .odx extension and contain additional visualization information in their headers and custom attribute information in their bodies.

The primary functions of the orchestration engine are:

- Persistence
- Hosting the .NET components
- Transactions
- Large message support
- Runtime validation
- Load throttling

Using the tests mapped to the **Orchestration Engine** layer you can monitor the orchestrations, the BAM interceptor, and the tracking data decode service offered by the Orchestration engine.

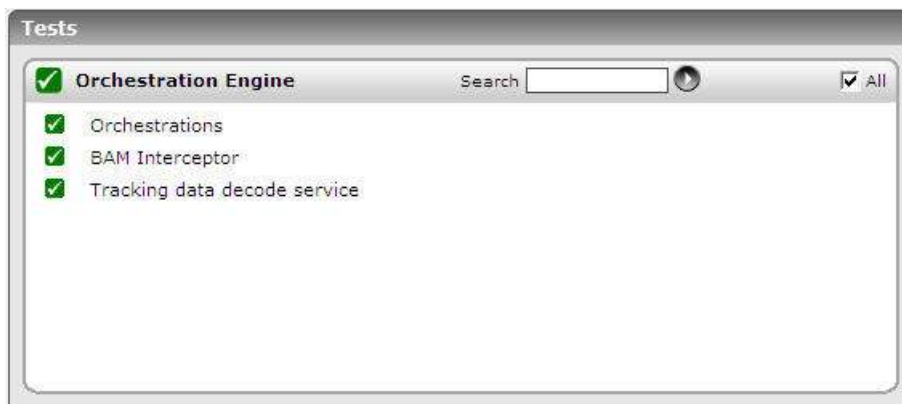


Figure 4.9: The tests mapped to the Orchestration Engine layer

4.2.3.1 BT Orchestrations Test

Orchestrations are executable business processes that can subscribe to (receive) and publish (send) messages through the MessageBox database. In addition, orchestrations can construct new messages. Messages are received using the subscription and routing infrastructure.

When subscriptions are filled for orchestrations, a new instance is activated and the message is delivered, or in the case of instance subscriptions, the instance is rehydrated if necessary and the message is then delivered. When messages are sent from an orchestration, they are published to the MessageBox in the same manner as a message arriving on a receive location with the appropriate properties getting inserted into the database for use in routing.

Messages that are constructed in an orchestration must be placed in the MessageBox database and referenced by the orchestration instance, but they should not be published because they have not yet been sent. The XLANG/s subservice makes calls to the Message Agent API to insert messages directly. This allows the orchestration engine to insert the message body into the MessageBox and have it directly associated with the running orchestration instance. The persistence of the constructed message in the MessageBox database is coordinated with persistence points in the orchestration as an additional optimization of database operations.

This test helps you determine the number of orchestrations that were created on each host instance, and also tracks the status of these orchestrations over time, thereby promptly alerting you when too many orchestrations are suspended or discarded. The test also tracks the memory usage of the orchestrations, and alerts you if excessive memory is being consumed.

| | |
|---------------------------------|---|
| Purpose | Helps you determine the number of orchestrations that were created on each host instance, and also tracks the status of these orchestrations over time, thereby promptly alerting you when too many orchestrations are suspended or discarded |
| Target of the test | A BizTalk Server 2010 |
| Agent deploying the test | An internal agent |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance on the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Idle orchestrations: Indicates the number of idle orchestration instances currently hosted by this host instance. | Number | This refers to orchestrations that are not making progress but are not dehydratable, as when the orchestration is blocked waiting for a receive, listen, or delay in an atomic transaction. |
| | Orchestrations created: Indicates the number of orchestration instances that were created since this host instance was started. | Number | |
| | Orchestrations created: Indicates the rate at which the orchestration instances were created on this host instance. | Orchestrations / Sec | |
| | Running orchestrations: Indicates the number of orchestration instances that are currently executing on this host instance. | Number | |
| | Orchestrations completed: Indicates the number of orchestration instances that were completed since this host instance was started. | Number | |
| | Orchestration completion rate: Indicates the rate at which the orchestration instances are completed. | Orchestrations/Sec | A high value is desired for this measure. A low value or a steady decline in the value of this measure could indicate an execution bottleneck. |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|--------------------|--|
| | <p>Orchestrations discarded:</p> <p>Indicates the number of orchestration instances discarded from memory since this host instance was started.</p> | Number | An orchestration can be discarded if the engine fails to persist in its state. |
| | <p>Orchestrations discarded:</p> <p>Indicates the rate at which orchestrations instances were discarded from the memory of this host instance.</p> | Orchestrations/Sec | |
| | <p>Orchestrations suspended:</p> <p>Indicates the number of orchestration instances that are suspended since this host instance was started.</p> | Number | <p>All failures encountered in orchestrations appear as exceptions.</p> <p>If an orchestration does not include any CatchException shape for an exception, the exception causes the orchestration to be Suspended, but not resumable. This means that message and service instance tracking, or a WMI script, cannot recover the instance. However, you can save all messages associated with the Suspended (not Resumable) instance using tracking (or WMI script) for diagnostic and manual retry.</p> <p>To diagnose the problem, use the Orchestration Debugger to see the last shape executed before the instance was suspended. You can also view exception details using the Orchestration Debugger.</p> |
| | <p>Orchestrations suspended:</p> <p>Indicates the rate at which orchestrations were suspended on this host instance.</p> | Orchestrations/Sec | |
| | <p>Orchestrations rehydrated:</p> <p>Indicates the number of orchestration instances that were rehydrated since this host instance was started.</p> | Number | <p>Rehydration is the process of deserializing the last running state of an orchestration from the database.</p> <p>The orchestration engine can be triggered to rehydrate an orchestration instance by the receipt of a message or by the expiration of a</p> |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|---------------------------|---|
| | <p>Orchestrations rehydrated:</p> <p>Indicates the rate at which orchestrations instances were rehydrated on this host instance.</p> | <p>Orchestrations/Sec</p> | <p>time-out specified in a Delay shape. It then loads the saved orchestration instance into memory, restores its state, and runs it from the point where it left off.</p> |
| | <p>Orchestrations dehydrated:</p> <p>Indicates the number of orchestration instances that were dehydrated since this host instance was started.</p> | <p>Number</p> | <p>Dehydration is the process of serializing the state of an orchestration into a SQL Server database.</p> <p>The orchestration engine might determine that an orchestration instance has been idle for a relatively long period of time. It calculates thresholds to determine how long it will wait for various actions to take place, and if those thresholds are exceeded, it dehydrates the instance. This can occur under the following circumstances:</p> |
| | <p>Orchestrations dehydrated:</p> <p>Indicates the rate at which orchestration instances were dehydrated on this host instance.</p> | <p>Orchestrations/Sec</p> | <ul style="list-style-type: none"> • When the orchestration is waiting to receive a message, and the wait is longer than a threshold determined by the engine. • When the orchestration is "listening" for a message, as it does when you use a Listen shape, and no branch is triggered before a threshold determined by the engine. The only exception to this is when the Listen shape contains an activation receive. • When a delay in the orchestration is longer than a threshold determined by the engine. <p>The engine dehydrates the instance by saving the state, and frees up the memory required by the instance. By dehydrating dormant orchestration instances, the engine makes it possible for a large number of long-running business processes to run concurrently on the same computer. This implies that the larger the number and rate of dehydrations minimal will be the use of system resources.</p> |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|-----------|--|
| | <p>Pending messages:</p> <p>Indicates the number of received messages for which receipt has not yet been acknowledged to the message box from the orchestration.</p> | Number | A very large value could indicate a processing bottleneck. |
| | <p>Pending work items:</p> <p>Indicates the number of code execution blocks that are scheduled for execution in the orchestration.</p> | Number | |
| | <p>Failure connections:</p> <p>Indicates the number of attempted database connections that has failed since this host instance was started.</p> | Number | Ideally, the value of this measure should be 0. |
| | <p>Database transactions:</p> <p>Indicates the number of database transactions performed since the host instance was started.</p> | Number | |
| | <p>Transactions / Sec:</p> <p>Indicates the rate of database transactions performed by the orchestrations hosted by this host instance.</p> | Trans/Sec | |
| | <p>Current Orchestrations Instances:</p> <p>Indicates the number of orchestration instances currently hosted by this host instance.</p> | Number | |
| | <p>Private memory:</p> <p>Indicates the allocated private memory for this host instance.</p> | MB | This is the current size of memory that this process has allocated that cannot be shared with other processes. |
| | <p>Virtual memory:</p> <p>Indicates the reserved virtual memory for this host instance.</p> | MB | This is the current size of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries. |

MONITORING THE BIZTALK SERVER

| | | | |
|--|--|---------|--|
| | Total physical memory: Indicates the percentage of total physical memory used on this host instance. | Percent | The dehydration behavior of BizTalk Server depends entirely on how much memory is available and how much memory is being used. The dehydration behavior is different with different amounts of memory and differences in memory use between 32-bit and 64-bit hosts. |
|--|--|---------|--|

4.2.3.2 BT BAM Interceptor Test

Information workers need flexibility in looking at and evaluating business processes. A purchasing manager might need to see how many POs are approved and denied each day, for example, while a sales manager might want an hourly update on what products are being ordered. Meeting these diverse needs requires a general framework for tracking what's going on with a particular business process. This is exactly what the Business Activity Monitoring (BAM) component in Microsoft BizTalk Server provides.

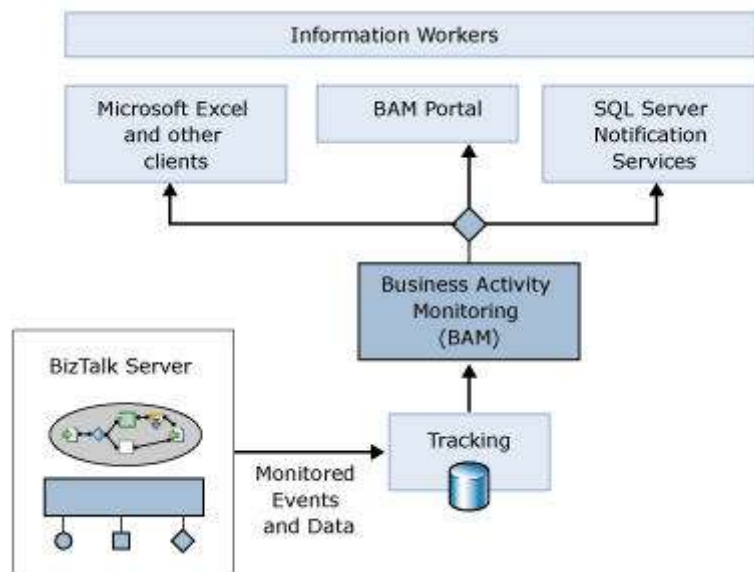


Figure 4.10: How does BAM work?

As the figure above illustrates, the BAM component allows monitoring of events and data produced by a BizTalk application. This information is made accessible using SOAP-callable Web services, and it can be accessed in several different ways, as follows:

- Through Microsoft Excel or other desktop clients, such as a custom dashboard application.
- Using a BAM portal, a component in BizTalk Server that enables business users to examine and configure BAM information.
- Through SQL Server Notification Services, allowing BAM information to be delivered as notifications.

The BAM Interceptor is an object that lets you instrument your application to capture data of interest. Using this test,

MONITORING THE BIZTALK SERVER

you can monitor the BAM interceptors, and swiftly spot the failure of BAM events.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | You can monitor the BAM interceptors, and swiftly spot the failure of BAM events | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total Failed Events: Indicates the total number of failed BAM events that occurred during data flush. | Number | |

4.2.3.3 BT Tracking Data Decode Service Test

The BAM Event Bus Service, also known as the Tracking Data Decode Service (TDDS), processes tracking data (streams) stored in a source database and persists that data in such a way that it is easy to query it at a later date. The BAM Event Bus service moves Business intelligence data to the BAM Primary Import database and BizTalk Health Monitoring data to the DTA database.

This test reveals the processing power of the TDDS by reporting the number of batches, events, and records it processes, and also sheds light on failures experienced by the TDDS while processing.

| | | | |
|----------------------------|--|--|--|
| Purpose | Reveals the processing power of the TDDS by reporting the number of batches, events, and records it processes, and also sheds light on failures experienced by the TDDS while processing | | |
| Target of the test | A BizTalk Server 2010 | | |
| Agent deploying the | An internal agent | | |

MONITORING THE BIZTALK SERVER

| | | | |
|---|---|-------------------------|----------------------------------|
| test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. ISPASSIVE - If the ISPASSIVE parameter is set to Yes, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up. | | |
| Outputs of the test | One set of results for each host instance of the BizTalk server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total Failed Batches: Indicates the total number of batches that the TDDS has failed to process on this host instance. | Number | |
| | Total Failed Events: Indicates the total number of batches that the TDDS has failed to process on this host instance. | Number | Ideally, this value should be 0. |
| | Total Events: Indicates the total number of events that are processed by the TDDS since you started it on this host instance. | Number | |

MONITORING THE BIZTALK SERVER

| | | | |
|--|---|--------|--|
| | Total Records: Indicates the total number of records that are processed by the TDDS since you started it on this host instance. | Number | |
|--|---|--------|--|

Monitoring DHCP Servers

The Microsoft® Windows® 2000 Server network operating system builds on the Microsoft support for Dynamic Host Configuration Protocol (DHCP).

Each host computer connected to a TCP/IP network must be assigned a unique IP address. The Microsoft DHCP server allows the network administrator to dynamically assign network settings to the clients that connect to a network.

The DHCP server offers the following features:

- Integration of DHCP with DNS.
- Dynamic assignment of IP addresses allows address reuse through leases.
- Multicast address allocation.
- Automatic pushdown of configurations to clients allows configuration changes to be applied transparently.

If the DHCP server experiences an overload or a slowdown while processing requests, it is bound to delay the automatic discovery of additions (client / server) to the network and the assignment of identification (i.e., IP address) to them; consequently, users may be denied timely access to critical clients or servers. Continuous monitoring of the DHCP server can alone help administrators in promptly identifying and resolving such problem conditions.

eG Enterprise prescribes a unique *DHCP* monitoring model (see Figure 5.1) for the DHCP server, which keeps a watchful eye on the requests received and acknowledgements sent by the server to help administrators determine the following:

- How quickly is the DHCP server processing request packets? Were too many requests enqueued? Have too many packets expired?
- Is the hardware on the DHCP server adequately sized to facilitate swift processing of the request packets?
- Were any negative acknowledgement messages sent by the DHCP server?
- Were any DHCP decline messages received by the server?
- Have enough IP addresses been configured on the server for assignment to clients?

MONITORING DHCP SERVERS



Figure 5.1: Layer model of a DHCP server

Every layer of Figure 5.1 above is mapped to a set of tests. The eG agent executing on the DHCP server runs these tests on the server, and extracts the metrics of interest.

Since the bottom 5 layers of Figure 5.1 have already been discussed in the *Monitoring Unix and Windows Servers* document, the section to come will discuss the **DHCP Services** layer only.

5.1 The DHCP Services Layer

The tests associated with this layer do the following:

- Track the overall responsiveness of the DHCP server to requests received from clients
- Verify the availability of free IP addresses on the server for assignment to clients



Figure 5.2: Tests associated with the DHCP Services layer

5.1.1 DHCP Performance Test

This test reports the performance statistics of the Microsoft 2000 DHCP server running on the network.

| | |
|----------------------------|--|
| Purpose | Reports the performance statistics of the DHCP server on a Windows 2000 network. |
| Target of the test | Any DHCP server |
| Agent deploying the | An internal agent |

MONITORING DHCP SERVERS

| | | | |
|---|---|-------------------------|--|
| test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The variable name of the host for which the test is to be configured. PORT – Refers to the port used by the DHCP server | | |
| Outputs of the test | One set of results for server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Avg packet rate: Refers to the average time in seconds used by the DHCP server to process each packet it receives. | Pkts/sec | This measure can vary depending on the server hardware and its I/O subsystem. A sudden or unusual increase might indicate a problem, either with the I/O subsystem becoming slower or because of an intrinsic processing overhead on the server computer. |
| | Current message queue length: Refers to the current length of the internal message queue of the DHCP server. | Number | A large value in this measure might indicate heavy server traffic. |
| | Request rate: Refers to the number of DHCP request messages received per second by the DHCP server from clients. | Reqs/sec | A sudden or unusual increase in this measure indicates a large number of clients trying to renew their leases with the DHCP server |
| | Request acks rate: Refers to the number of DHCP acknowledgement messages sent per second by the DHCP server to clients. | Reqs/sec | A sudden or unusual increase in this measure indicates that a large number of clients are being renewed by the DHCP server |
| | Request nacks rate: Refers to the number of negative acknowledgement messages sent per second by the DHCP server to clients. | Reqs/sec | A very high value might indicate potential network trouble in the form of misconfiguration of either the server or clients. When servers are misconfigured, one possible cause is a deactivated scope. For clients, a very high value could be caused by computers moving between subnets, such as laptop portables or other mobile devices. |
| | Request declines rate: Refers to the number of DHCP decline messages received per second by the DHCP server from clients. | Reqs/sec | A high value indicates that several clients have found their address to be in conflict, possibly indicating network trouble. |

MONITORING DHCP SERVERS

| | | | |
|--|---|----------|---|
| | Packets expired rate: Refers to the number of packets per second that expire and are dropped by the DHCP server. | Pkts/sec | A large value in this measure indicates that the server is either taking too long to process some packets while other packets are queued and becoming stale, or traffic on the network is too high for the server to manage. |
| | Packet drop rate: Refers to the number of duplicate packets per second dropped by the DHCP server. | Pkts/sec | This measure can be affected by multiple clients or network interfaces forwarding the same packet to the server. A large value in this measure indicates that either clients are probably timing out too fast or the server is not responding fast enough. |
| | Requests release rate: Refers to the number of DHCP release messages received per second by the DHCP server from clients. | Reqs/sec | This measure only exists if a DHCP client sends a release message to the server. This measure remains low for many DHCP network configurations . |
| | | | |

5.1.2 DHCP Utilization Test

This test reports general statistics pertaining to the Microsoft 2000 DHCP server running on the network.

| | | | |
|---|---|-------------------------|--|
| Purpose | Reports the statistics of the DHCP server on Windows 2000 network. | | |
| Target of the test | Any DHCP server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured. PORT - Refers to the port used by the DHCP server. | | |
| Outputs of the test | One set of results for server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Current addresses in use: Refers to the number of IP addresses in use in the target network. | Number | This measure indicates the number of IP addresses assigned to clients in the target network. |

MONITORING DHCP SERVERS

| | | | |
|--|--|--------|--|
| | Free addresses: Refers to the number of free IP addresses available in the target network. | Number | This measure indicates the number of IP addresses available for allocation to clients in the target network. |
|--|--|--------|--|

Monitoring the Windows Internet Name Service (WINS)

The Windows Internet Name Service (WINS) provides a distributed database for registering and querying dynamic mappings of NetBIOS names for computers and groups used on your network. WINS maps NetBIOS names to IP addresses and was designed to solve the problems arising from NetBIOS name resolution in routed environments. The main benefit of a WINS server is that it avoids the need for broadcasts to resolve computer names to IP addresses.

Typically, WINS servers use UDP port 137. **This port should be provided when you manually add a WINS server for monitoring.** The steps below highlight how WINS works:

- **Name Registration:** When a WINS client initializes, it registers its NetBIOS name by sending a name request to the configured WINS server. All services get registered as they are initialized in the WINS server database. If the WINS server is available and the name is not registered by another machine, the WINS server returns a successful registration message.
- If the NetBIOS name is already registered in the WINS database, the WINS server will send a challenge to the current registered owner. This request will be sent 3 times at 500ms intervals. If the current owner responds the WINS server will send a negative name resolution response to the WINS client attempting to register the name. If there is no response the registering client will receive a Name Registration response.
- **Name Renewal:** To continue using the same NetBIOS name, a client must renew its lease before it expires. If the client does not renew the lease, the WINS server makes it available to another WINS client. A WINS client will first attempt to refresh its name registration request after 1/8 of the TTL is completed. If the client is successful subsequent name registration requests will occur when 1/2 the TTL is expired.

If the client is unsuccessful with lease renewal on the initial attempt the client will try every 2 minutes until 1/2 TTL is remaining. At 1/2 of TTL the client will revert to the secondary WINS server if configured in 1/8 TTL intervals. At completion of TTL lease, the WINS client will revert back to the primary WINS server and start the process all over again.

- **Name Release:** Before the expiry of its lease, a client can send an explicit request to release the name assigned to it.

If even one of these steps experience latencies, it could cause a significant delay in the entire process of resolving an IP address to its corresponding NetBIOS name. This could be much worse in large environments where the WINS server might have to handle hundreds of concurrent 'name resolution' requests; here, even a seemingly insignificant drop in the processing rate of the WINS server can grow in severity within minutes, and can bring the whole environment to a virtual standstill!

If such adverse consequences are to be prevented, it is recommended that you continuously monitor the processing ability of the WINS server, so that you are promptly alerted when there is any threat to its normal functioning.

MONITORING THE WINDOWS INTERNET NAME SERVICE (WINS)

eG Enterprise offers a 100% web-based *WINS* monitoring model (see Figure 6.1) that closely observes the performance of the WINS server in relation to real-time changes in load.

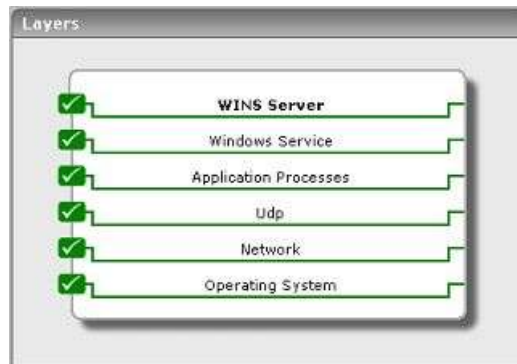


Figure 6.1: Layer model of a WINS server

Figure 6.1 comprises of a set of hierarchical layers, each of which is associated with one/more tests. The eG agent on the WINS server periodically executes these tests on the server, extracts performance data from the server, and instantly alerts administrators of an impending overload or a probable dip in the processing speed of the server.

The sections to come discuss the top layer of Figure 6.1 alone, as all other layers have been discussed in the *Monitoring Unix and Windows Servers* document.

6.1 The WINS Server Layer

Using the **Wins** test associated with it, this layer measures the rate at which the WINS server processes requests.



Figure 6.2: Test associated with the WINS server layer

6.1.1 Wins Test

This test reports general statistics pertaining to the Windows Internet Name Service (WINS).

| | |
|----------------------------|--|
| Purpose | Reports general statistics pertaining to the WINS server |
| Target of the test | A WINS server |
| Agent deploying the | An internal agent |

MONITORING THE WINDOWS INTERNET NAME SERVICE (WINS)

| | | | |
|---|---|-------------------------|---|
| test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the WINS server | | |
| Outputs of the test | One set of results for every server | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Queries: The total number of queries received by the WINS server | Queries/sec | This indicates the server workload. It is useful for capacity planning and to detect unusual usage situations. |
| | Failed queries: Total number of failed queries/sec | Failures/sec | The percentage of failed queries should be low. An unusually high number of failed queries can indicate a configuration problem, or a fault in the WINS server. |
| | Releases: The rate at which release requests are received and processed by the WINS server | Releases/sec | |
| | Failed releases: The rate of release failures | Failures/sec | Release failures could result in many names being unused for a period of time, and hence, should be minimized. |
| | Conflicts: The total rate of conflicts seen by the WINS server. This value includes both Unique and Group conflicts. | Conflicts/sec | |
| | Renewals: The total rate of renewal requests received by the WINS server. This value includes both Unique and Group renewals. | Renewals/sec | |

Monitoring MS Print Servers

Print servers are a popular mode of sharing printing resources in IT infrastructures. The Microsoft Windows operating system allows for specific servers to be designated and managed as print servers. Some of the key reasons for why IT administrators configure and use print servers include centralized management of print drivers, access control and prioritization of print jobs, central auditing capability or charging, etc. Since print servers are common resources for all the users of an IT infrastructure, IT administrators must continuously monitor the print servers to ensure high uptime, good performance, and scalability.

The eG Enterprise suite includes specialized monitoring capability for Microsoft Windows-based print servers. The layer model of a print server is given below (see Figure 7.1)

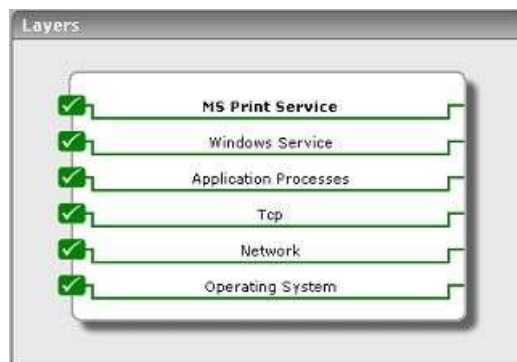


Figure 7.1: Layer model of an MS Print server

The section that follows discusses the **MS Print Service** layer only, as all other layers have been extensively discussed in the *Monitoring Unix and Windows Servers* document.

7.1 The MS Print Service Layer

This layer (see Figure 7.2) monitors the print queues on the print server and reports on their availability and overall health.



Figure 7.2: Tests associated with the MS Print Service layer

7.1.1 Print Server Test

This test auto-discovers the print queues of a print server and continuously tracks various key metrics relating to the availability and performance of each of the print queues.

| | | | |
|--|---|-------------------------|---|
| Purpose | Tracks various key metrics relating to the availability and performance of each of the print queues of a print server | | |
| Target | An MS Print server | | |
| Agent deploying this test | An internal agent | | |
| Configurable parameters for this test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST - The host for which the test is to be configured. 3. PORT - The port to which the specified HOST listens 4. USEWMI - If the USEWMI flag is Yes, then the test uses WMI to extract the statistics of interest. This option is provided because on some Windows 2000 systems (especially ones with service pack 3 or lower), the use of WMI access can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the USEWMI parameter value to No. The default is No. | | |
| Outputs of the test | One set of results for every print queue monitored | | |
| Measurements of the test | Measurement | Measurement Unit | Interpretation |
| | Availability: Indicates whether or not the Print server is currently available. | Boolean | If the value of this measure is <i>1</i> , it indicates that the print server is available. The value <i>0</i> on the other hand, indicates that the print server is unavailable. |
| | Jobs services: The rate at which users' jobs are being processed over a print queue | Jobs/Sec | The value of this metric is a key indicator of a print queue's workload. |

MONITORING MS PRINT SERVERS

| | | | |
|--|---|------------|--|
| | <p>Pages printed: The number of pages printed through a print queue during the last measurement period</p> | Number | This is another key indicator of the workload of a print queue. |
| | <p>Print traffic: Indicates the rate at which data is transmitted to a print queue for printing</p> | KBytes/Sec | |
| | <p>Current jobs: Shows the current number of jobs in a print queue.</p> | Number | Use this counter to identify excessive use of a print queue. |
| | <p>Print errors: The number of jobs to a print queue that resulted in errors during the last measurement period.</p> | Number | This value includes the number of out of paper errors and printer not ready errors. Job errors can occur even if the connection to the printer has errors due to network problems. |
| | <p>Spooled jobs: The current number of spooling jobs in a print queue</p> | Number | |
| | <p>Paper errors: The total number of out of paper errors that occurred in a print queue during the last measurement period</p> | Number | |
| | <p>Not ready errors: The total number of out of printer not ready errors that occurred in a print queue during the last measurement period</p> | Number | |

Monitoring MS Proxy Servers

Microsoft Proxy Server 2.0 is an extensible firewall and content cache server, providing Internet security while improving network response time and efficiency by 50%, on average, for businesses of all sizes. It is the first firewall product to include high-performance content caching. Similarly, it is the first content cache server to provide firewall support. Microsoft Proxy Server 2.0 offers distributed (hierarchical and array-based) Web caching, providing unbeaten scalability, fault-tolerance and load balancing to meet even the rigorous demands of large enterprises and Internet Service Providers. MS Proxy Server acts as a gateway with firewall-class security between a LAN and the Internet. The product also blocks access to undesirable sites and provides other easy-to-use management features. It works with existing networks, including IPX networks, and supports several Internet protocols and services. It is therefore imperative that the MS Proxy server is continuously monitored, so that security risks to your environment are minimized, and business is transacted smoothly and efficiently.

eG Enterprise provides a specialized *Microsoft Proxy* monitoring model (see Figure 8.1 that monitors the internal health and external availability and responsiveness of the Microsoft Proxy server, and alerts administrators to potential performance issues.

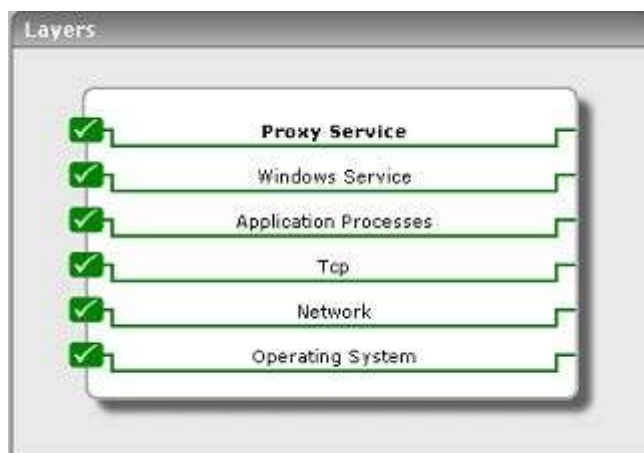


Figure 8.1: Layer model of an MS Proxy server

8.1 The Proxy Service Layer

The tests mapped to the **Proxy Service** layer monitors the performance of the following services executing on an MS Proxy server:

- The WinSock Proxy Service
- The Web Proxy Service
- The Caching service



Figure 8.2: Tests associated with the Proxy Service layer

8.1.1 Win Sock Test

The WinSock Proxy service supports Microsoft Windows operating systems using Windows Sockets. Windows Sockets is an interprocess communication mechanism derived from the Berkeley Sockets interface (originally designed for Unix systems). The Sockets interface was extended to support Windows-based clients running Microsoft implementations of TCP/IP. The name given to this Sockets interface for Windows was WinSock (for Windows Sockets). The WinSock Proxy Service support is available for both Transmission Control Protocol/Internet Protocol (TCP/IP) and Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocols. The WinSock Proxy service applies mainly to Windows clients, including Windows 3.x, Windows 95, and Windows NT.

This test reports the performance statistics pertaining to this WinSock Proxy Service.

| | |
|---|--|
| Purpose | Reports the performance statistics pertaining to this WinSock Proxy service |
| Target of the test | An MS Proxy Server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the MS Proxy server |

MONITORING MS PROXY SERVERS

| | | | |
|--|--|---|---|
| Outputs of the test | One set of results for every WinSock monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Accepting TCP connections: The number of TCP connection objects that will wait for TCP connections from WinSock proxy clients | Percent | A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted. |
| | Active sessions: The number of active sessions for the WinSock proxy service | Number | |
| | Active TCP connections: The total number of TCP connections that are currently transmitting data | Number | |
| | Active UDP connections: The number of active UDP connections | Number | |
| | Available worker threads: The number of available WinSock worker threads | Number | The high increase in the number may affect the performance of the host / applications. |
| | Data received: The rate at which data is received | KB/sec | A low value could indicate a network bottleneck |
| | Data transmitted: The rate at which data is submitted | KB/sec | A high value of this measure could result in a network congestion |
| | Failed DNS resolutions: The number of calls that have failed to resolve DNS domain name and IP address for WinSock proxy connections | Number | This value must be low; a high value indicates that there may be a network / WinSock service problem on the host. |
| Pending DNS requests: The number of calls awaiting DNS domain name and IP address resolution for WinSock proxy connections | Number | This value must be low; a high value indicates that there may be a network / WinSock service problem on the host. | |

| | | | |
|--|---|--------|---|
| | <p>Worker threads:</p> <p>The number of WinSock worker threads that are currently available or alive</p> | Number | An increase in this value may affect the performance of the host / application. |
|--|---|--------|---|

8.1.2 Proxy Server Test

The Web Proxy service provides support for HTTP (a.k.a. Web publishing), FTP, Gopher, and secure (SSL) communications. The Web Proxy service works with any CERN-compliant Web browser, such as Internet Explorer or Netscape Navigator. Because the Web Proxy supports only these widely adopted Internet standard communication methods, it isn't operating system dependent. Clients running Unix, Macintosh, or Windows operating systems can communicate with the Web Proxy service as long as they're configured with a CERN-compliant Web browser.

This test reports the performance statistics pertaining to this Web Proxy service running on an MS Proxy server.

| | | | |
|---|--|-------------------------|---|
| Purpose | Reports performance statistics pertaining to the Web Proxy service running on an MS Proxy server | | |
| Target of the test | An MS Proxy Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the MS Proxy server | | |
| Outputs of the test | One set of results for every web proxy service monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Cache hit ratio:</p> <p>The percentage of requests that have used cached data, to the total number of requests to the web proxy service</p> | Percent | A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted. |
| | <p>Client data receive rate:</p> <p>The number of active sessions for the web proxy service</p> | Number | A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications. |
| | <p>Client data transmit rate:</p> <p>The rate at which the data bytes are sent by the proxy server to the web proxy clients</p> | Kb/sec | A high value could indicate a high data transfer from the proxy server to the web proxy client, which may result in congestion in network traffic |

MONITORING MS PROXY SERVERS

| | | | |
|--|--|----------|--|
| | <p>Avg response time: The mean response time in seconds to service a request</p> | Secs/req | High network traffic, low server performance are some of the factors that cause this measure to increase. |
| | <p>Current users: The current number of users connected to the web proxy service.</p> | Number | A high value can indicate an increase in the load on the web proxy service. |
| | <p>DNS cache hits: This measure give the percentage of DNS domain names served from the proxy server cache, from the total DNS entries that are retrieved by the web proxy service.</p> | Percent | A high value can indicate an increase in load on web proxy service. |
| | <p>Failing requests: The rate of request that have completed with some error.</p> | Reqs/Sec | The high value indicates possible problems in the web proxy service. |
| | <p>FTP requests: The number of ftp requests that have been made to the web proxy service</p> | Number | A high value can indicate an increase in the load on the web proxy service. |
| | <p>HTTP requests: The number of http requests that have been made to the web proxy service.</p> | Number | A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications. |
| | <p>HTTPS sessions: The total number of HTTP-Secured sessions serviced by the SSL tunnel</p> | Number | A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications on the server. |
| | <p>Thread pool active sessions: The number of sessions being actively served by the pool of threads</p> | Number | A high value can indicate an increase in the load on the web proxy service. |
| | <p>Thread pool failures: The number of requests rejected, since the thread pool was overcommitted</p> | Number | The high value indicates a possible problem in the thread pool of the web proxy service. |

MONITORING MS PROXY SERVERS

| | | | |
|--|---|--------|---|
| | Upstream receive rate: The rate at which the data is received by the web proxy service from remote servers on the internet/proxy servers surrounding the current proxy server | Kb/sec | A high value can indicate an increase in the load on the web proxy service from one or more remote servers. |
| | Upstream transmit rate: The rate at which the data is sent by the web proxy service to remote servers on the internet/proxy servers surrounding the current proxy server | Kb/sec | A high value can indicate an increase in the load of one or more remote servers. |

8.1.3 Proxy Cache Test

Web site caching is an efficient use of resources and another benefit of the MS proxy server. Since you can use the MS proxy server as a common connection point to the Internet, you can also use it to cache frequently accessed resources. The proxy server allocates a portion of the server's hard disk space to store frequently accessed objects. Internet requests are more efficiently responded to through the use of fresh-cached data, which in the long run, helps in minimizing internet response times.

Caching can either be passive or active. Passive caching just stores objects as they are requested, so the cache is updated only when users request information. Active caching directs the server to refresh objects in the cache automatically.

You can selectively control the proxy server caching so that you can limit the size of cached objects, change the expiration limits (control the freshness of objects), and determine whether the server always caches, or always excludes from cache, certain content.

This test reports the performance statistics pertaining to this caching activity of the MS Proxy server.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the performance statistics pertaining to this caching activity of the MS Proxy server | | |
| Target of the test | An MS Proxy Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the MS Proxy server | | |
| Outputs of the test | One set of results for every web proxy server cache monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MS PROXY SERVERS

| | | | |
|--|---|----------|---|
| | Active refreshes: The rate at which data is retrieved from the Internet to refresh popular URLs in the URL cache. | Kb/sec | A low value indicates low refresh rate and a possible network problem. |
| | Active URL refreshes: The rate at which the URLs in the URL cache are refreshed from the internet | URSS/sec | A low or 0(zero) indicates the non-availability of URLs or DNS servers from the internet. |
| | Cache size: The total number of bytes currently available in the URL Cache | Kb | A high value indicates possible high usage of virtual memory on web proxy cache. |
| | URL commits: The rate at which the URLs are committed to the URLs cache | URLs/sec | The low value or 0 (zero) indicates low URL commits, low network resource availability. |
| | URLs retrieved: The rate at which the URLs are retrieved from the URL cache. | URLs/sec | A low value indicates the low availability of the URLs from the proxy cache. |
| | URLs in cache The current number of URLs in the URL cache | Number | A high value indicates possible low availability of virtual memory. |

8.1.4 Proxy Svc Test

This test can be executed from a location external to the proxy server, and presents an unbiased external perspective of the state of the server. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Microsoft Proxy* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

| | |
|--|---|
| Purpose | This test measures the state of an MS proxy server |
| Target | An MS Proxy server |
| Agent deploying this test | An external agent executing on an eG server |
| Configurable parameters for this test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed URL – The web page being accessed. While multiple URLs (separated by commas) can be provided, each URL should be of the format URL name:URL value. URL name is a unique name assigned to the URL, and the URL value is the value of the URL. For example, a URL can be specified as HomePage:http://192.168.10.12:7077/, where HomePage is the URL name and http://192.168.10.12:7077/ is the URL value. |

MONITORING MS PROXY SERVERS

| | | | |
|---------------------------------|--|-------------------------|--|
| | <ol style="list-style-type: none"> 3. HOST - The host for which the test is to be configured. 4. PORT - The port to which the specified HOST listens 5. COOKIEFILE – Whether any cookies being returned by the MS Proxy server need to be saved locally and returned with subsequent requests 6. PROXYHOST – The host on which a web proxy server is running (in case a proxy server is to be used) 7. PROXYPORT – The port number on which the web proxy server is listening 8. PROXYUSERNAME – The user name of the proxy server 9. PROXYPASSWORD – The password of the proxy server 10. CONFIRM PASSWORD – Confirm the password by retyping it here. 11. CONTENT – Is a set of instruction:value pairs that are used to validate the content being returned by the test. If the CONTENT value is <i>none:none</i>, no validation is performed. The number of pairs specified in this text box, must be equal to the number of URLs being monitored. The instruction should be one of <i>Inc</i> or <i>Exc</i>. <i>Inc</i> tells the test that for the content returned by the MS Proxy server to be valid, the content must include the specified value (a simple string search is done in this case). An instruction of <i>Exc</i> instructs the test that the server's output is valid if it does not contain the specified value. 12. CREDENTIALS – The HttpTest supports HTTP authentication. The CREDENTIALS parameter is to be set if a specific user name / password has to be specified to login to a page. This parameter is a comma separated list of user name:password pairs, one pair for each URL being monitored. A value of none:none indicates that user authorization is not required. Please be sure to check if your web site requires HTTP authentication while configuring this parameter. HTTP authentication typically involves a separate pop-up window when you try to access the page. Many sites uses HTTP POST for obtaining the user name and password and validating the user login. In such cases, the username and password have to be provided as part of the POST information and NOT as part of the CREDENTIALS specification for the HttpTest. | | |
| Outputs of the test | One set of outputs for every URL being monitored | | |
| Measurements of the test | Measurement | Measurement Unit | Interpretation |
| | <p>Proxy service availability:</p> <p>This measurement indicates whether the server was able to respond successfully to the query made by the test.</p> | Percent | <p>Availability failures could be caused by several factors such as the MS Proxy process(es) being down, the MS Proxy servers being misconfigured, a network failure, etc. Temporary unavailability may also occur if the proxy server is overloaded. Availability is determined based on the response code returned by the server. A response code between 200 to 300 indicates that the server is available.</p> |

MONITORING MS PROXY SERVERS

| | | | |
|--|---|---------|--|
| | <p>Total response time:</p> <p>This measurement indicates the time taken by the server to respond to the requests it receives.</p> | Secs | <p>Response time being high denotes a problem. Poor response times may be due to the server being overloaded or misconfigured. If the URL accessed involves the generation of dynamic content by the server, backend problems (e.g., an overload at the application server or a database failure) can also result in an increase in response time.</p> |
| | <p>TCP connection availability:</p> <p>This measure indicates whether the test managed to establish a TCP connection to the server.</p> | Percent | <p>Failure to establish a TCP connection may imply that either the MS proxy 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.</p> |
| | <p>TCP connection time:</p> <p>This measure quantifies the time for establishing a TCP connection to the MS proxy server host.</p> | Secs | <p>Typically, the TCP connection establishment must be very small (of the order of a few milliseconds). Since TCP connection establishment is handled at the OS-level, rather than by the application, an increase in this value signifies a system-level bottleneck on the host that supports the MS proxy server.</p> |
| | <p>Server response time:</p> <p>This measure indicates the time period between when the connection was established and when the server sent back a HTTP response header to the client.</p> | Secs | <p>While the total response time may depend on several factors, the server response time is typically, a very good indicator of a server bottleneck (e.g., because all the available server threads or processes are in use).</p> |
| | <p>Response code:</p> <p>The response code returned by the server for the simulated request</p> | Number | <p>A value between 200 and 300 indicates a good response. A 4xx value indicates a problem with the requested content (eg., page not found). A 5xx value indicates a server error.</p> |
| | <p>Content length:</p> <p>The size of the content returned by the server</p> | Kbytes | <p>Typically the content length returned by the server for a specific URL should be the same across time. Any change in this metric may indicate the need for further investigation on the server side.</p> |

MONITORING MS PROXY SERVERS

| | | | |
|--|--|---------|--|
| | <p>Content validity:</p> <p>This measure validates whether the server was successful in executing the request made to it.</p> | Percent | <p>A value of 100% indicates that the content returned by the test is valid. A value of 0% indicates that the content may not be valid. This capability for content validation is especially important for multi-tier web applications. For example, a user may not be able to login to the web site but the server may reply back with a valid HTML page where in the error message, say, "Invalid Login" is reported. In this case, the availability will be 100 % (since we got a valid HTML response). If the test is configured such that the content parameter should exclude the string "Invalid Login," in the above scenario content validity would have a value 0.</p> |
|--|--|---------|--|

Monitoring Windows Domain Controllers

Windows Domain Controllers are critical components of IT infrastructures. Users accessing resources in a Windows domain have to first be authenticated by the Domain Controller in order to get access. Any slowdown or failure of the domain controllers can severely impact users. Hence, 24x7 monitoring of domain controllers is critical.

The eG Enterprise suite provides a specialized *Domain Controller* monitoring model for the Windows domain controller (DC) (see Figure 9.1), using which key performance parameters related to the DC can be continuously monitored, and anomalies, instantly detected.

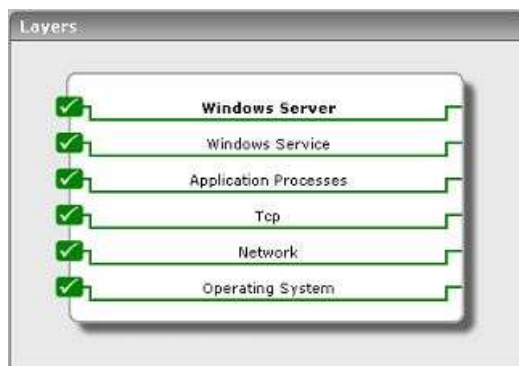


Figure 9.1: Layer model of a Windows Domain Controller

Each of the layers in this specialized model (see Figure 9.1) executes a wide variety of tests on the DC and extracts critical metrics, which help quantify the performance level achieved by the DC, and simplifies problem identification.

The *Monitoring Unix and Windows Servers* document deals extensively with the bottom 5 layers of Figure 9.1. In the section that follows, the **Windows Server** layer will be discussed.

9.1 The Windows Server Layer

Using the tests associated with this layer, administrators can gauge how effectively the DC authenticates login requests it receives.



Figure 9.2: Tests associated with the Windows Server layer

9.1.1 Windows Access Test

This test monitors the accesses to a Windows server.

| | | | |
|---|--|-------------------------|---|
| Purpose | Monitors the accesses to the Windows server | | |
| Target of the test | A Windows server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the Windows server | | |
| Outputs of the test | One set of results for every Windows server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Blocking request rejects: The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items | Reqs/sec | If the number of blocking request rejects is high, you may need to adjust the <code>MaxWorkItem</code> or <code>MinFreeWorkItems</code> server parameters |
| | Permission errors: The number of times opens on behalf of clients have failed with <code>STATUS_ACCESS_DENIED</code> in the last measurement period | Number | Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected. |

MONITORING WINDOWS DOMAIN CONTROLLERS

| | | | |
|--|--|------------|--|
| | File access denied errors: The number of times accesses to files opened successfully were denied in the last measurement period | Number | This number indicates attempts to access files without proper access authorization. |
| | Internal server errors: This value indicates the number of times an internal server error was detected in the last measurement period. | Number | Unexpected errors usually indicate a problem with the server. |
| | Data received: The rate at which the server has received data from the network | Kbytes/sec | This metric indicates how busy the server is. |
| | Data transmitted: The rate at which the server has sent data over the network | Kbytes/sec | This metric indicates how busy the server is. |
| | Resource shortage errors: The number of times STATUS_DATA_NOT_ACCEPTED was returned to clients in the last measurement period | Number | A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the InitWorkItems or MaxWorkItems server parameters might need to be adjusted. |
| | Avg response time: Average time taken by the server to respond to client requests | Secs | This is a critical measure of server health. |

9.1.2 Windows Sessions Test

This test reports various session-related statistics for a Windows server.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports various session-related statistics for a Windows server | | |
| Target of the test | A Windows Domain Controller | | |
| 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 Windows server being monitored | | |
| Measurements | Measurement | Measurement Unit | Interpretation |

MONITORING WINDOWS DOMAIN CONTROLLERS

| | | | |
|------------------|--|----------|--|
| made by the test | Logons: Rate of logons to the server | Reqs/sec | This measure reports the rate of all interactive, network, and service logons to a windows server. The measure includes both successful and failed logons. |
| | Logon errors: Number of logons in the last measurement period that had errors | Number | This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server. |
| | Current sessions: The number of sessions currently active in a server | Number | This measure is one of the indicators of current server activity. |
| | Sessions with errors: The number of sessions in the last measurement period that were closed to unexpected error conditions | Number | Sessions can be closed with errors if the session duration reaches the autodisconnect timeout. |
| | Sessions forced off: The number of sessions in the last measurement period that have been forced to logoff | Number | This value indicates how many sessions were forced to logoff due to logon time constraints. |
| | Sessions logged off: The number of sessions in the last measurement period that were terminated normally | Number | Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low. |
| | Sessions timed out: The number of sessions that have been closed in the last measurement period due to their idle time exceeding the AutoDisconnect parameter for the server | Number | The number of session timed out gives an indication of whether the AutoDisconnect setting is helping to conserve server resources |

9.1.3 Window Authentication Test

This test emulates a user logging into a Windows domain or local host and reports whether the login succeeded and how long it took.

| | |
|----------------|--|
| Purpose | Emulates a user logging into a windows domain or a local host and reports whether the login succeeded and how long it took |
|----------------|--|

MONITORING WINDOWS DOMAIN CONTROLLERS

| | | | |
|---|--|-------------------------|---|
| Target of the test | A Windows Domain Controller | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows server 4. USERNAME - This test emulates a user logging into a Microsoft Windows domain or a local host. Therefore, specify the login name of the user here. 5. PASSWORD - Enter the password that corresponds to the specified USERNAME. 6. DOMAIN - Specify the name of the domain to which the test will try to login. If the test is to login to a local host, specify 'none' here. <div style="border: 1px solid black; padding: 10px; margin-top: 10px;"> <p>Note:</p> <p>If users are spread across multiple domains, then, you can configure this test with multiple DOMAIN specifications; in this case, for every DOMAIN, a USER-PASSWORD pair might also have to be configured. Sometimes, you might want the test to login as specific users from the same domain, to check how long each user login takes. Both these scenarios require the configuration of multiple DOMAINS and/or multiple USER names and PASSWORDS. In order to enable users to specify these details with ease, eG Enterprise provides a special page; to access this page, click on the Click here hyperlink at the top of the parameters in the test configuration page. To know how to use this page, refer to the Configuring Multiple Users for the Citrix Authentication Test section in the <i>Monitoring Citrix Environments</i> document.</p> </div> | | |
| Outputs of the test | One set of results for every user account being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Authentication status: Indicates whether the login was successful or not | Percent | A value of 100 % indicates that the login has succeeded. The value 0 is indicative of a failed login. |
| | Authentication time: Indicates the time it took to login | Secs | If this value is very high then it could be owing to a configuration issue (i.e. the domain might not be configured properly) or a slow-down/unavailability of the primary domain server. |

Monitoring MS File Servers

In the client/server model, a file server is a computer responsible for the central storage and management of data files so that other computers on the same network can access the files. A file server allows users to share information over a network without having to physically transfer files. Any computer can be configured to be a host and act as a file server. In its simplest form, a file server may be an ordinary PC that handles requests for files and sends them over the network. In a more sophisticated network, a file server might be a dedicated network-attached storage device that also serves as a remote hard disk drive for other computers, allowing anyone on the network to store files on it as if to their own hard drive.

The true indicator of the efficiency of a File server is the speed with which it serves concurrent file requests. If users are unable to access important files stored on the file server as and when they need due to a temporary break in connection to the server or because of a long request queue, it might severely hamper the productivity of the users, and might unnecessarily delay critical operations. If such a problem situation is to be averted, the file server needs to be monitored, and administrators promptly warned about probable performance issues.

eG Enterprise provides out-of-the-box a specialized *Microsoft File* server model (see Figure 10.1) that periodically runs diagnostic tests on the file server to ensure that it performs to peak capacity at all times.

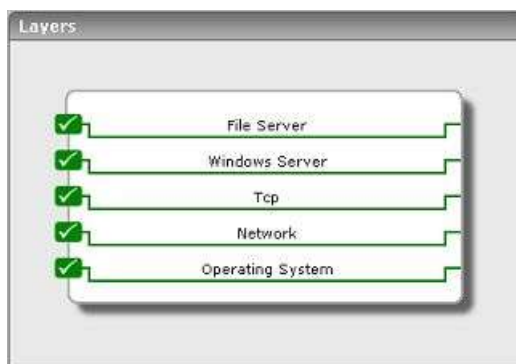


Figure 10.1: Layer model of an MS File server

The sections to come discuss the top 2 layers of Figure 10.1, since the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

10.1 The Windows Server Layer

Using the tests associated with the **Windows Server** layer, administrators can closely observe the user logins to and session behavior on the MS File server.



Figure 10.2: Tests associated with the Windows Server layer

10.1.1 Windows Access Test

This test monitors the accesses to the MS File server.

| | | | |
|---|--|-------------------------|---|
| Purpose | Monitors the accesses to the MS File server | | |
| Target of the test | An MS File server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the File server | | |
| Outputs of the test | One set of results for every File server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Blocking request rejects: The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items | Reqs/sec | If the number of blocking request rejects is high, you may need to adjust the <code>MaxWorkItem</code> or <code>MinFreeWorkItems</code> server parameters |

MONITORING MS FILE SERVERS

| | | | |
|--|--|------------|--|
| | Permission errors: The number of times opens on behalf of clients have failed with STATUS_ACCESS_DENIED in the last measurement period | Number | Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected. |
| | File access denied errors: The number of times accesses to files opened successfully were denied in the last measurement period | Number | This number indicates attempts to access files without proper access authorization. |
| | Internal server errors: This value indicates the number of times an internal server error was detected in the last measurement period. | Number | Unexpected errors usually indicate a problem with the server. |
| | Data received: The rate at which the server has received data from the network | Kbytes/sec | This metric indicates how busy the server is. |
| | Data transmitted: The rate at which the server has sent data over the network | Kbytes/sec | This metric indicates how busy the server is. |
| | Resource shortage errors: The number of times STATUS_DATA_NOT_ACCEPTED was returned to clients in the last measurement period | Number | A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the InitWorkItems or MaxWorkItems server parameters might need to be adjusted. |
| | Avg response time: Average time taken by the server to respond to client requests | Secs | This is a critical measure of server health. |

10.1.2 Windows Sessions Test

This test reports various session-related statistics for an MS File server.

| | |
|---------------------------------|--|
| Purpose | Reports various session-related statistics for an MS File server |
| Target of the test | An MS File server |
| Agent deploying the test | An internal agent |

MONITORING MS FILE SERVERS

| | | | |
|--|--|---|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the MS File server | | |
| Outputs of the test | One set of results for every MS File server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Logons: Rate of logons to the server | Reqs/sec | This measure reports the rate of all interactive, network, and service logons to an MS File server. The measure includes both successful and failed logons. |
| | Logon errors: Number of logons in the last measurement period that had errors | Number | This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server. |
| | Current sessions: The number of sessions currently active in a server | Number | This measure is one of the indicators of current server activity. |
| | Sessions with errors: The number of sessions in the last measurement period that were closed to unexpected error conditions | Number | Sessions can be closed with errors if the session duration reaches the autodisconnect timeout. |
| | Sessions forced off: The number of sessions in the last measurement period that have been forced to logoff | Number | This value indicates how many sessions were forced to logoff due to logon time constraints. |
| | Sessions logged off: The number of sessions in the last measurement period that were terminated normally | Number | Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low. |
| Sessions timed out: The number of sessions that have been closed in the last measurement period due to their idle time exceeding the AutoDisconnect parameter for the server | Number | The number of session timed out gives an indication of whether the AutoDisconnect setting is helping to conserve server resources | |

10.2 The File Server Layer

With the help of the tests associated with this layer, administrators can:

- Accurately determine the current work load on the server in terms of the number of files currently accessed on the server and the current user traffic to the server
- Quickly identify locked files and the users who have acquired a lock on those files



Figure 10.3: Tests associated with the File server layer

10.2.1 MS File Stats Test

The MsFileTest tracks various statistics pertaining to open file connections at the host.

| | |
|---|--|
| Purpose | Tracks various statistics pertaining to open file connections at the host |
| Target of the test | An MS File server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. 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: <ul style="list-style-type: none"> ➤ The eG manager license should allow the detailed diagnosis capability ➤ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
| Outputs of the test | One set of results for every MS file server monitored |

MONITORING MS FILE SERVERS

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|--|------------------|---|
| | File locks count: The number of files locked at the host | Number | A high value can indicate too many files being opened at the host. The detailed diagnosis of this measure, if enabled, lists the files that have been locked, the user who holds the lock, and the number of locks on the file. |
| | Unique users count: A unique count of users who have opened files at this host | Number | A high value can indicate too many users connected to the host. |

10.2.2 Windows Usage Test

This test tracks various statistics pertaining to sessions open at the host.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Tracks various statistics pertaining to sessions open at the host | | |
| Target of the test | An MS File server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured 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: <ul style="list-style-type: none"> ➤ The eG manager license should allow the detailed diagnosis capability ➤ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for every MS file server monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING MS FILE SERVERS

| | | | |
|------|--|--------|--|
| test | Open files: The number of files opened over the network by users connecting to the file server | Number | This measurement is an indicator of the workload on the file server. The detailed diagnosis of this measure, if enabled, provides the number of open sessions for every user, and the time for which the sessions have been idle. If the idle time displayed here is very high, then measures for closing the inactive open sessions can be initiated. |
| | Unique users: A unique count of users who have opened sessions at this host | Number | A high value can indicate too many users connected to the host. |

Monitoring ISA Proxy Servers

Microsoft Internet Security and Acceleration (ISA) Server can be deployed as a dedicated firewall that acts as the secure gateway to the Internet for internal clients. ISA Server protects all communication between internal computers and the Internet. In a simple firewall scenario, the ISA Server computer has two network interface cards, one connected to the local network and one connected to the Internet. By setting the security access policies, you prevent unauthorized access and malicious content from entering the network. You can also restrict what traffic is allowed for each user and group, application, destination, content type, and schedule.

To assure users of safe and secure access to the Internet, and to shield the network from malicious attacks, the availability and internal health of the ISA Proxy server should be constantly monitored.

The eG Enterprise suite's unique *ISA Proxy* monitoring model (see Figure 11.1) executes a wide variety of tests on the proxy server to enable administrators to determine the following:

- Does the server take too much time to service firewall requests?
- Is the Web Proxy server functioning optimally?
- Is the Web Proxy Cache utilized effectively?

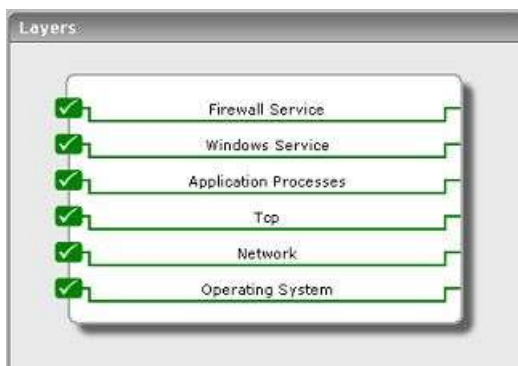


Figure 11.1: Layer model of an ISA Proxy server

The sections to come will discuss the tests associated with the **Firewall Service** layer only, since the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

11.1 The Firewall Service Layer

The tests associated with the **Firewall Service** layer monitor various critical firewall services provided by the ISA Proxy server.



Figure 11.2: The tests associated with the Firewall Service layer

11.1.1 ISA Cache Test

This test reports statistics pertaining to the ISA Proxy server cache.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports statistics pertaining to the ISA Proxy server cache | | |
| Target of the test | An ISA Proxy server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Data received from disk cache: Indicates the rate at which data is retrieved from the disk cache. | KB/Sec | |

MONITORING ISA PROXY SERVERS

| | | | |
|--|---|------------|---|
| | Data received from memory cache: Indicates the rate at which data is retrieved from the memory cache. | KB/Sec | |
| | Disk failures: Indicates the rate at which I/O failures occurred since the Firewall service started. | Fails/Sec | An I/O failure occurs when the ISA server fails to read from or write to disk cache. |
| | Disk writes: Indicates the rate at which data was written to the disk cache. | Writes/Sec | |
| | Memory cache util: Indicates the percentage of fetches made from the memory. | Percent | A high percentage may indicate that it is worthwhile allocating more available memory resources to the cache. |
| | URLs in cache: Indicates the number of URLs currently stored in the cache. | Number | |

11.1.2 ISA Firewall Test

This test reports statistics pertaining to the Firewall service of the ISA Proxy server 2004.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Measures the firewall protection of the ISA proxy server | | |
| Target of the test | An ISA Proxy server 2004 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | DNS cache hit ratio: Indicates the percentage of DNS domain names retrieved from the DNS cache. | Percent | |

MONITORING ISA PROXY SERVERS

| | | | |
|--|---|--------|--|
| | Pending DNS resolutions: Indicates the number of gethostbyname and gethostbyaddr API calls pending resolution. These are calls used to resolve host DNS domain names and IP addresses for Firewall Service connections. | Number | |
| | Worker threads: Indicates the number of Firewall Service worker threads that are currently alive. | Number | A high value indicates that the current workload of the ISA Proxy Server is very high. |

11.1.3 ISA Web Proxy Test

This test monitors the performance of the Web proxy service of the ISA Proxy server 2004.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Measures the firewall protection of the ISA proxy server | | |
| Target of the test | An ISA Proxy server 2004 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active web sessions: Indicates the number of active Web sessions currently connected to the ISA proxy Server. | Number | |
| | Connect errors: Indicates the total number of errors that occurred while connecting. | Number | |

MONITORING ISA PROXY SERVERS

| | | | |
|--|---|--------------|--|
| | DNS cache hit ratio: Indicates the percentage of DNS domain names served from the DNS cache. | Percent | |
| | Failed requests: Indicates the rate of requests that have failed because of some type of error. | Conns/Sec | A high failure rate, in comparison to the rate of incoming requests, will suggest that the ISA Proxy server is having difficulty in coping with all incoming requests. Connection settings for incoming Web requests may be incorrectly configured, or connection bandwidth may be insufficient. |
| | Inbound connections: Indicates the rate of incoming connections. | Conns/Sec | |
| | Outbound connections: Indicates the rate of outgoing connections. | Conns/Sec | |
| | Requests rate: Indicates the rate of requests to the Web Proxy filter. | Requests/Sec | A higher value means that more ISA Proxy server resources will be required to service incoming requests. |

11.1.4 Packet Engine Test

The PacketEngine test reports statistics relating to the firewall packet engine of the ISA Proxy server 2004.

| | | | |
|---|--|-------------------------|--|
| Purpose | Reports statistics relating to the firewall packet engine of the ISA Proxy server 2004 | | |
| Target of the test | An ISA Proxy server 2004 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active connections: Indicates the number of active connections currently transmitting data. | Number | A high value indicates that the current workload of the ISA Proxy Server is very high. |

MONITORING ISA PROXY SERVERS

| | | | |
|--|---|-------------|--|
| | Allowed packets rate: Indicates the number of packets allowed per second. | Packets/Sec | |
| | Data sent rate: Indicates the rate at which data was transmitted by the firewall packet engine driver. | KB/Sec | |
| | Dropped packets rate: Indicates the rate at which packets were dropped. | Packets/Sec | |
| | New connections rate: Indicates the rate at which connections were created. | Conns/Sec | |
| | Packets inspected rate: Indicates the rate at which the firewall packet engine driver inspects the packets. | Packets/Sec | |

11.1.5 Proxy Server Test

The Web Proxy service provides support for HTTP (a.k.a. Web publishing), FTP, Gopher, and secure (SSL) communications. The Web Proxy service works with any CERN-compliant Web browser, such as Internet Explorer or Netscape Navigator. Because the Web Proxy supports only these widely adopted Internet standard communication methods, it isn't operating system dependent. Clients running Unix, Macintosh, or Windows operating systems can communicate with the Web Proxy service as long as they're configured with a CERN-compliant Web browser.

This test reports the performance statistics pertaining to this Web Proxy service running on an ISA Proxy server.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Reports performance statistics pertaining to the Web Proxy service running on an ISA Proxy server | | |
| Target of the test | An ISA Proxy Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every web proxy service monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING ISA PROXY SERVERS

| | | | |
|--|---|----------|--|
| | <p>Cache hit ratio</p> <p>The percentage of requests that have used cached data, to the total number of requests to the web proxy service</p> | Percent | A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted. |
| | <p>Client data receive rate:</p> <p>The number of active sessions for the web proxy service</p> | Number | A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications. |
| | <p>Client data transmit rate:</p> <p>The rate at which the data bytes are sent by the proxy server to the web proxy clients</p> | Kb/sec | A high value could indicate a high data transfer from the proxy server to the web proxy client, which may result in congestion in network traffic |
| | <p>Avg response time:</p> <p>The mean response time in seconds to service a request</p> | Secs/req | High network traffic, low server performance are some of the factors that cause this measure to increase. |
| | <p>Current users:</p> <p>The current number of users connected to the web proxy service.</p> | Number | A high value can indicate an increase in the load on the web proxy service. |
| | <p>DNS cache hits:</p> <p>This measure give the percentage of DNS domain names served from the proxy server cache, from the total DNS entries that are retrieved by the web proxy service.</p> | Percent | A high value can indicate an increase in load on web proxy service. |
| | <p>Failing requests:</p> <p>The rate of request that have completed with some error.</p> | Reqs/Sec | The high value indicates possible problems in the web proxy service. |
| | <p>FTP requests:</p> <p>The number of ftp requests that have been made to the web proxy service</p> | Number | A high value can indicate an increase in the load on the web proxy service. |
| | <p>HTTP requests:</p> <p>The number of http requests that have been made to the web proxy service.</p> | Number | A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications. |
| | <p>HTTPS sessions:</p> <p>The total number of HTTP-Secured sessions serviced by the SSL tunnel</p> | Number | A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications on the server. |

MONITORING ISA PROXY SERVERS

| | | | |
|--|--|--------|---|
| | <p>Thread pool active sessions:</p> <p>The number of sessions being actively served by the pool of threads</p> | Number | A high value can indicate an increase in the load on the web proxy service. |
| | <p>Thread pool failures:</p> <p>The number of requests rejected, since the thread pool was overcommitted</p> | Number | The high value indicates a possible problem in the thread pool of the web proxy service. |
| | <p>Upstream receive rate:</p> <p>The rate at which the data is received by the web proxy service from remote servers on the internet/proxy servers surrounding the current proxy server</p> | Kb/sec | A high value can indicate an increase in the load on the web proxy service from one or more remote servers. |
| | <p>Upstream transmit rate:</p> <p>The rate at which the data is sent by the web proxy service to remote servers on the internet/proxy servers surrounding the current proxy server</p> | Kb/sec | A high value can indicate an increase in the load of one or more remote servers. |

11.1.6 Tests that are Disabled by Default

In addition to the tests discussed above, the **Firewall Service** layer is also mapped to a few tests that are disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *ISA Proxy* as the **Component type**, *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the >> button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

The tests that are disabled by default have been discussed in the following sections.

11.1.6.1 Firewall Service Test

The FirewallService test measures the firewall protection of the ISA proxy server.

| | |
|---------------------------------|--|
| Purpose | Measures the firewall protection of the ISA proxy server |
| Target of the test | An ISA Proxy server |
| Agent deploying the test | An internal agent |

MONITORING ISA PROXY SERVERS

| | | | |
|--|--|---|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active sessions: Indicates the number of active sessions for the firewall service. | Number | Comparing this measure at both peak and off-peak times will provide you with valuable insight into the usage patterns of the ISA server. |
| | Active TCP connections: Indicates the number of active TCP connections transmitting data. | Number | |
| | Active UDP connections: Indicates the number of active UDP connections for the firewall service. | Number | |
| | Active threads: Indicates the number of firewall worker threads that are currently active. | Number | |
| | Read rate: Indicates the number of kilobytes read by the data-pump per second. | KB/Sec | A consistent decrease in the value of this measure may indicate a delay in servicing firewall requests. |
| Write rate: Indicates the number of kilobytes written by the data-pump per second. | KB/Sec | A consistent decrease in the value of this measure may indicate a delay in servicing firewall requests. | |

11.1.6.2 Web Proxy Service Test

This test monitors the Web Proxy service. Requests from Web Proxy clients are directed to the Web Proxy service on the ISA server to determine if access is allowed.

| | |
|---------------------------|--------------------------------|
| Purpose | Monitors the Web Proxy service |
| Target of the test | An ISA Proxy server |
| Agent | An internal agent |

MONITORING ISA PROXY SERVERS

| | | | |
|--|---|-------------------------|---|
| deploying the test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Cache hit ratio: The percentage of successful web proxy client requests to the ISA Server. | Percent | This measure is a good indicator of the effectiveness of the cache. A higher percentage indicates that a number of requests are being serviced from the cache. This in turn is indicative of faster responsiveness. A zero value indicates that caching is not enabled, and a low value may indicate a configuration problem. |
| | Current users: Indicates the number of clients that are currently running the web proxy service. | Number | Monitoring this measure at both peak and off-peak times will enable users to assess the extent of server usage. This measure may also be useful if you need to temporarily stop ISA Server services. |
| | Read rate: Indicates the rate at which data bytes are received from Web Proxy clients. | KB/Sec | A consistent decrease in the value of this measure may indicate a delay in servicing requests. |
| | Active threads: Indicates the rate at which data bytes are sent to Web Proxy clients. | KB/Sec | A consistent decrease in the value of this measure may indicate a delay in servicing requests. |
| | Avg requests sec: Indicates the number of kilobytes read by the data-pump per second. | KB/Sec | A consistent decrease in the value of this measure may indicate a delay in servicing firewall requests. |
| | Write rate: Indicates the average amount of time required by the ISA server to process a request. | Secs/Request | This measure can be monitored at peak and off-peak times to receive a clear idea about how fast client requests are being serviced. A very high value of this measure might indicate that the ISA Server is having difficulty in handling all requests. |
| Thread pool size: Indicates the number of threads in the thread pool | Number | | |

MONITORING ISA PROXY SERVERS

| | | | |
|--|--|--------|--|
| | Thread pool sessions: Indicates the number of sessions being actively serviced by thread pool threads. | Number | |
| | Thread pool failures: Indicates the number of requests rejected because the thread pool was full. | Number | |

11.1.6.3 Web Proxy Cache Test

This test monitors the Web Proxy cache. The ISA server implements a cache of frequently-requested objects to improve network performance. You can configure the cache to ensure that it contains the data that is most frequently used by the organization or accessed by your Internet clients.

| | | | |
|---|--|-------------------------|--|
| Purpose | Monitors the Web Proxy cache | | |
| Target of the test | An ISA Proxy server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT – Refers to the port used by the ISA Proxy server | | |
| Outputs of the test | One set of results for every ISA Proxy server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Disk cache space: Indicates the amount of space used by the disk cache. | KB | If the value of this measure grows closer to or equal to the allocated disk cache space, it would indicate that subsequent cache requests might be rejected due to non-availability of adequate cache space. This, in turn, would increase the rate of direct disk accesses, which will consequently degrade system performance. |
| | Memory cache space: Indicates the amount of space used by the memory cache | KB | An excessive consumption of the memory cache space would result in slow-down of the system. This is because the lack of sufficient space in the memory cache would cause the real memory (RAM) to directly service the requests for objects. |

MONITORING ISA PROXY SERVERS

| | | | |
|--|--|----------|--|
| | URL commit rate: Indicates the speed at which URLs are being written to the cache. | URLs/Sec | |
|--|--|----------|--|

Monitoring Microsoft Radius Servers

Internet Authentication Service (IAS) in Microsoft® Windows Server™ is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access, including wireless, authenticating switch, dial-up and virtual private network (VPN) remote access, and router-to-router connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers.

The following illustration shows IAS as a RADIUS server for a variety of access clients and a RADIUS proxy. IAS uses an Active Directory domain for user credential authentication of incoming RADIUS Access-Request messages.

When IAS is used as a RADIUS server, RADIUS messages provide authentication, authorization, and accounting for network access connections in the following way:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the IAS server.
3. The IAS server evaluates the Access-Request message.
4. If required, the IAS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the IAS server.
5. The user credentials are checked and the dial-in properties of the user account are obtained by using a secure connection to a domain controller.
6. The connection attempt is authorized with both the dial-in properties of the user account and remote access policies.
7. If the connection attempt is both authenticated and authorized, the IAS server sends an Access-Accept message to the access server. If the connection attempt is either not authenticated or not authorized, the IAS server sends an Access-Reject message to the access server.
8. The access server completes the connection process with the access client and sends an Accounting-Request message to the IAS server, where the message is logged.
9. The IAS server sends an Accounting-Response to the access server.

MONITORING MICROSOFT RADIUS SERVERS

Issues in the functioning of IAS, if not promptly isolated and resolved, might result in the complete collapse of the remote authentication and authorization service provided by the Windows server. 24x7 monitoring of IAS, hence becomes imperative.

The eG Enterprise suite provides out-of-the-box monitoring support to the Windows Internet Authentication Service, and proactively alerts administrators of authentication, authorization, or accounting bottlenecks encountered by the IAS server. The specialized *Microsoft Radius* monitoring model (see Figure 12.1) offered by the eG Enterprise suite executes a variety of tests on the IAS server; these tests, in turn, use the perfmon utility of Windows to extract critical performance statistics pertaining to the services offered by the IAS server.

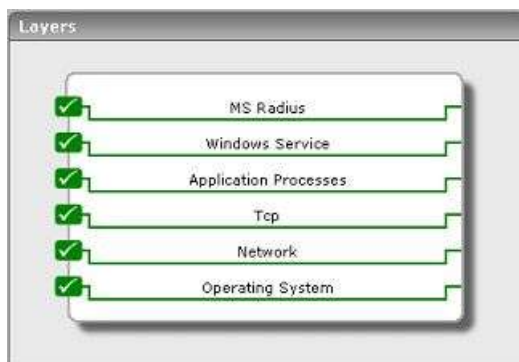


Figure 12.1: The layer model of the MS Radius server

This section will discuss the **MS Radius** layer alone, as all the other layers have been elaborately discussed in the *Monitoring Unix and Windows Servers* document.

12.1 The MS Radius Layer

This layer monitors the authentication, authorization, and accounting activities performed by the IAS server and clients.



Figure 12.2: The tests associated with the MS Radius layer

12.1.1 IAS Acc Server Test

Besides providing remote authentication services to RADIUS clients, the IAS server also provides a central accounting recording service for all accounting requests that are sent by the RADIUS clients. Once the IAS server completes the connection process initiated by a RADIUS client, the access server which processed the connection request sends an Accounting-Request message to the IAS server, where the message is logged. The IAS server then sends an Accounting-Response to the access server. In addition, the access server also sends Accounting-Request messages for the following:

- During the time in which the connection is established
- When the access client connection is closed
- When the access server is started and stopped

The IasAccSvr test monitors the accounting-requests received and accounting-responses sent by the IAS server to RADIUS clients.

| | | | |
|---|--|-------------------------|---|
| Purpose | Monitors the accounting-requests received and accounting-responses sent by the IAS server to RADIUS clients | | |
| Target of the test | An IAS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT - The port at which the IAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for every IAS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Packets sent: Indicates the rate at which packets were sent by the IAS server. | Packets/Sec | |
| | Packets received: Indicates the rate at which the IAS server received packets. | Packets/Sec | When viewed along with the <i>Packets sent</i> measure, this measure serves as a good indicator of the traffic on the server. |

MONITORING MICROSOFT RADIUS SERVERS

| | | | |
|--|---|-------------|--|
| | <p>Packets dropped:</p> <p>Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.</p> | Packets/Sec | A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation. |
| | <p>Invalid requests:</p> <p>Indicates the rate at which packets were received from an unknown address.</p> | Reqs/Sec | |
| | <p>Malformed packets:</p> <p>Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.</p> | Packets/Sec | |
| | <p>Unknown packets:</p> <p>Indicates the rate at which packets of an unknown type were received.</p> | Packets/Sec | |
| | <p>No record packets:</p> <p>Indicates the rate at which RADIUS Accounting-Request packets were received and responded to but not recorded.</p> | Records/Sec | |
| | <p>Accounting requests:</p> <p>Indicates the rate at which RADIUS Accounting-Requests were received from this client on the accounting port.</p> | Reqs/Sec | |
| | <p>Accounting responses:</p> <p>Indicates the rate at which RADIUS Accounting-Response packets were sent to this client on the accounting port.</p> | Reqs/Sec | The <i>Accounting requests</i> and <i>Accounting responses</i> measures serve as effective indicators of the workload on the IAS server. |
| | <p>Duplicate requests:</p> <p>Indicates the rate at which duplicate RADIUS Accounting-Request packets were received from this client.</p> | Reqs/Sec | |

MONITORING MICROSOFT RADIUS SERVERS

| | | | |
|--|--|----------|--|
| | Bad authenticators: Indicates the rate at which Accounting-Requests containing invalid signature attributes were received. | Reqs/Sec | |
|--|--|----------|--|

12.1.2 IAS Acc Client Test

The IasAccClient test monitors the accounting-requests sent and accounting-responses received by the RADIUS clients from the IAS servers.

| | | | |
|---|---|-------------------------|---|
| Purpose | Monitors the accounting-requests sent and accounting-responses received by the RADIUS clients from the IAS servers | | |
| Target of the test | An IAS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT - The port at which the IAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for every IAS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Packets sent: Indicates the rate at which packets were sent by this client. | Packets/Sec | |
| | Packets received: Indicates the rate at which this RADIUS client received packets. | Packets/Sec | When viewed along with the <i>Packets sent</i> measure, this measure serves as a good indicator of the traffic that originated from a client. |
| | Packets dropped: Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | Packets/Sec | A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation. |

MONITORING MICROSOFT RADIUS SERVERS

| | | | |
|--|--|-------------|--|
| | <p>Malformed packets:</p> <p>Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.</p> | Packets/Sec | |
| | <p>Unknown packets:</p> <p>Indicates the rate at which packets of an unknown type were received.</p> | Packets/Sec | |
| | <p>No record packets:</p> <p>Indicates the rate at which RADIUS Accounting-Request packets were received and responded to but not recorded.</p> | Records/Sec | |
| | <p>Accounting requests:</p> <p>Indicates the rate at which RADIUS Accounting-Requests were sent by this client on the accounting port.</p> | Reqs/Sec | |
| | <p>Accounting responses:</p> <p>Indicates the rate at which RADIUS Accounting-Response packets were sent to this client on the accounting port.</p> | Reqs/Sec | The <i>Accounting requests</i> and <i>Accounting responses</i> measures serve as effective indicators of the workload on the client. |
| | <p>Duplicate requests:</p> <p>Indicates the rate at which duplicate RADIUS Accounting-Request packets were received from this client.</p> | Reqs/Sec | |
| | <p>Bad authenticators:</p> <p>Indicates the rate at which Accounting-Requests containing invalid signature attributes were received.</p> | Reqs/Sec | |

12.1.3 IAS Auth Server Test

Internet Authentication Service (IAS) can be used as a RADIUS server to perform authentication, authorization, and accounting for RADIUS clients. When Internet Authentication Service (IAS) is used as a RADIUS server, it provides the a central authentication and authorization service for all access requests that are sent by RADIUS clients. IAS uses either a Microsoft® Windows NT® Server 4.0 domain, an Active Directory® domain, or the local Security Accounts Manager (SAM) to authenticate user credentials for a connection attempt. IAS uses the dial-in properties of the user account and remote access policies to authorize a connection.

MONITORING MICROSOFT RADIUS SERVERS

The IasAuthSvr test measures how well the IAS server performs remote authentication and authorization.

| | | | |
|---|---|-------------------------|---|
| Purpose | Measures how well the IAS server performs remote authentication and authorization | | |
| Target of the test | An IAS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT - The port at which the IAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for every IAS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Packets sent: Indicates the rate at which packets were sent by the IAS server. | Packets/Sec | |
| | Packets received: Indicates the rate at which the IAS server received packets. | Packets/Sec | When viewed along with the <i>Packets sent</i> measure, this measure serves as a good indicator of the traffic on the server. |
| | Packets dropped: Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | Packets/Sec | A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation. |
| | Invalid requests: Indicates the rate at which packets were received from an unknown address. | Reqs/Sec | |
| | Malformed packets: Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count. | Packets/Sec | |
| | Unknown packets: Indicates the rate at which packets of an unknown type were received. | Packets/Sec | |

MONITORING MICROSOFT RADIUS SERVERS

| | | | |
|--|---|----------------|---|
| | Access accepts: Indicates the rate at which RADIUS Access-Accept packets were sent by the IAS server to this client. | Accepts/Sec | |
| | Access challenges: Indicates the rate at which Access-Challenge messages are being processed. | Challenges/Sec | |
| | Access rejects: Indicates the rate at which Access-Reject messages are being processed. | Rejects/Sec | A very high value of this measure could warrant a review of the remote access policies. |
| | Access requests: Indicates the rate at which packets were received on an authentication port from this client. | Reqs/Sec | |
| | Duplicate requests: Indicates the rate at which duplicate RADIUS Access-Request packets were received from this client. | Reqs/Sec | |

12.1.4 IAS Auth Client Test

The IasAuthClient test monitors the access-requests sent by access-clients to the IAS server, and indicates how many requests were accepted/rejected by the IAS server.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the access-requests sent by access-clients to the IAS server, and indicates how many requests were accepted/rejected by the IAS server | | |
| Target of the test | An IAS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT - The port at which the IAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for every IAS server monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT RADIUS SERVERS

| | | | |
|------|---|----------------|---|
| test | <p>Packets sent:</p> <p>Indicates the rate at which packets were sent by this client to the IAS server.</p> | Packets/Sec | |
| | <p>Packets received:</p> <p>Indicates the rate at which this client received packets from the IAS server.</p> | Packets/Sec | When viewed along with the <i>Packets sent</i> measure, this measure serves as a good indicator of the traffic on the client. |
| | <p>Packets dropped:</p> <p>Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.</p> | Packets/Sec | A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation. |
| | <p>Malformed packets:</p> <p>Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.</p> | Packets/Sec | |
| | <p>Unknown packets:</p> <p>Indicates the rate at which packets of an unknown type were received.</p> | Packets/Sec | |
| | <p>Access accepts:</p> <p>Indicates the rate at which RADIUS Access-Accept packets were sent to this client.</p> | Accepts/Sec | |
| | <p>Access challenges:</p> <p>Indicates the rate at which Access-Challenge messages are being processed.</p> | Challenges/Sec | |
| | <p>Access rejects:</p> <p>Indicates the rate at which Access-Reject messages are being processed.</p> | Rejects/Sec | A very high value of this measure could warrant a review of the remote access policies. |
| | <p>Access requests:</p> <p>Indicates the rate at which packets were received on an authentication port from this client.</p> | Reqs/Sec | |

MONITORING MICROSOFT RADIUS SERVERS

| | | | |
|--|---|----------|--|
| | Bad authenticators: Indicates the rate at which packets containing invalid signature attributes were received. | Reqs/Sec | |
| | Duplicate requests: Indicates the rate at which duplicate RADIUS Access-Request packets were received from this client. | Reqs/Sec | |

Monitoring the Microsoft RAS Server

Microsoft Remote Access Service (RAS) is a feature in the Windows Server family, including Windows Server 2003, Windows 2000 Server, and , NT4 Server. A Limited version of RAS is also included in Windows XP Professional. RAS allows remote dial-up clients to connect to a Local Area Network using analog phone lines or ISDN lines. A typical use would be by an ISP (Internet Service Provider) to allow users to dial in to their LAN, or by a corporate network administrator to allow their users to connect to the corporate LAN from remote sites. The remote clients connect to RAS using the TCP/IP protocol encapsulated in the Point-to-Point (PPP) protocol, which allows the remote client to access the LAN as if they were plugged directly into it.

Needless to say, even a brief non-availability of RAS can cause critical services to go out of the reach of remote clients. Continuous monitoring of the RAS server can alone ensure a higher uptime of the RAS service. Using the *Microsoft RAS* monitoring model (see Figure 13.1) presented by the eG Enterprise suite, administrators can closely observe RAS operations 24x7, be forewarned of probable issues, and quickly attend to the issues before any permanent damage occurs.



Figure 13.1: Layer model of the MS RAS server

The sections to come will deal with the tests mapped to the **MS RAS Service** layer only, as the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

13.1 The MS RAS Service Layer

Using the tests depicted by Figure 13.2, the **MS RAS Service** layer enables administrators to assess the effectiveness of the dial-up communication service provided by the RAS device.

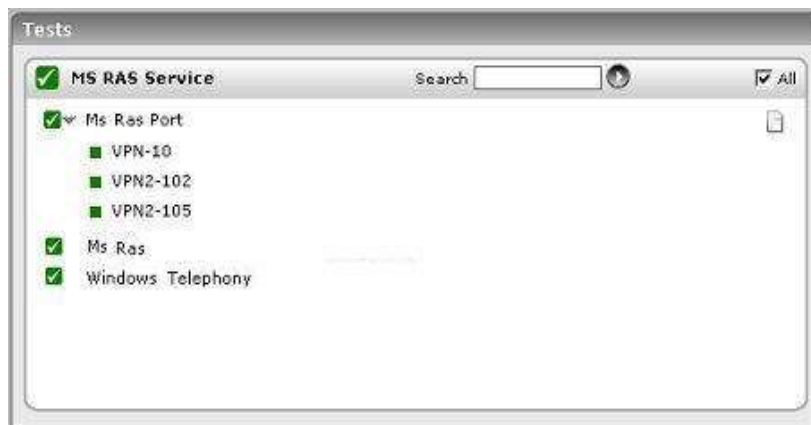


Figure 13.2: The tests associated with the MSRAS_SERVICE layer

13.1.1 Microsoft RAS Port Test

The MsRasPort test reports the performance statistics pertaining to every port of the Remote Access Service (RAS) device on the computer.

| | | | |
|---|---|-------------------------|--|
| Purpose | Reports the performance statistics pertaining to every port of the Remote Access Service (RAS) device on the computer | | |
| Target of the test | A Microsoft RAS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT - The TCP port at which the RAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for every RAS port | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes transmitted: Indicates the rate at which bytes were transmitted. | Bytes/Sec | |
| | Bytes received: Indicates the rate at which bytes were received. | Bytes/Sec | When viewed along with the Bytes_transmitted measure, this measure serves as a good indicator of the traffic on the network. |

MONITORING THE MICROSOFT RAS SERVER

| | | | |
|--|--|------------|---|
| | <p>Frames transmitted: Indicates the number of frames transmitted per second.</p> | Frames/Sec | |
| | <p>Frames received: Indicates the number of frames received per second.</p> | Frames/Sec | |
| | <p>Total errors: Indicates the number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors per second.</p> | Number | |
| | <p>Compression ratio for bytes sent: Indicates the compression ratio for the bytes being transmitted.</p> | Percent | |
| | <p>Compression ratio for bytes received: Indicates the compression ratio for the bytes being received.</p> | Percent | |
| | <p>Total connections: Indicates the number of remote access connections.</p> | Number | |
| | <p>CRC errors: Indicates the current number of CRC errors for this port.</p> | Number | CRC errors occur when the frame received contains erroneous data. |
| | <p>Timeout errors: Indicates the current number of timeout errors for this port.</p> | Number | Timeout errors occur when an expected packet is not received in time. |
| | <p>Serial overrun errors: Indicates the current number of serial overrun errors for this port t.</p> | Reqs/Sec | Serial Overrun errors occur when the hardware cannot handle the rate at which data is received. |
| | <p>Alignment errors: Indicates the current number of alignment errors for this port.</p> | Number | Alignment errors occur when a received byte is different from the expected byte. |

| | | | |
|--|---|--------|---|
| | Buffer overrun errors: Indicates the current number of buffer overrun errors for this port. | Number | Buffer Overrun errors occur when the software cannot handle the rate at which data is received. |
|--|---|--------|---|

13.1.2 Microsoft RAS Test

The MsRas test reports the performance statistics that are aggregated across all the ports of the Remote Access Service (RAS) device on the computer.

| | | | |
|---|---|-------------------------|--|
| Purpose | Reports the performance statistics that are aggregated across all the ports of the Remote Access Service (RAS) device on the computer | | |
| Target of the test | A Microsoft RAS server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured PORT - The TCP port at which the RAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for the RAS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bytes transmitted: Indicates the rate at which bytes were transmitted. | Bytes/Sec | |
| | Bytes received: Indicates the rate at which bytes were received. | Bytes/Sec | When viewed along with the Bytes_transmitted measure, this measure serves as a good indicator of the traffic on the network. |
| | Frames transmitted: Indicates the number of frames transmitted per second. | Frames/Sec | |
| | Frames received: Indicates the number of frames received per second. | Frames/Sec | |
| | Total errors: Indicates the number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors per second. | Number | |

| | | | |
|--|---|----------|---|
| | Compression ratio for bytes sent: Indicates the compression ratio for the bytes being transmitted. | Percent | |
| | Compression ratio for bytes received: Indicates the compression ratio for the bytes being received. | Percent | |
| | Total connections: Indicates the number of remote access connections. | Number | |
| | CRC errors: Indicates the current number of CRC errors for this port. | Number | CRC errors occur when the frame received contains erroneous data. |
| | Timeout errors: Indicates the current number of timeout errors for this port. | Number | Timeout errors occur when an expected packet is not received in time. |
| | Serial overrun errors: Indicates the current number of serial overrun errors for this port. | Reqs/Sec | Serial Overrun errors occur when the hardware cannot handle the rate at which data is received. |
| | Alignment errors: Indicates the current number of alignment errors for this port. | Number | Alignment errors occur when a received byte is different from the expected byte. |
| | Buffer overrun errors: Indicates the current number of buffer overrun errors for this port. | Number | Buffer Overrun errors occur when the software cannot handle the rate at which data is received. |

13.1.3 Windows Telephony Test

The MsTelephony test measures the performance of the telephone-communication activity on a computer running Windows 2000 or a higher operating system.

| | |
|----------------------------|--|
| Purpose | Measures the performance of the telephone-communication activity on a computer running Windows 2000 or a higher operating system |
| Target of the test | A Microsoft RAS server |
| Agent deploying the | An internal agent |

MONITORING THE MICROSOFT RAS SERVER

| | | | |
|---|--|-------------------------|-----------------------|
| test | | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured 3. PORT - The TCP port at which the RAS server listens. The default is NULL. | | |
| Outputs of the test | One set of results for the RAS server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Telephone lines: Indicates the number of telephone lines currently serviced by this computer. | Number | |
| | Telephone devices: Indicates the number of telephone devices (telephones or speaker phones) currently serviced by this computer. | Number | |
| | Active telephone lines: Indicates the number of telephone or integrated services digital network (ISDN) lines serviced by this computer that are currently in use by applications. | Number | |
| | Active telephone devices: Indicates the number of telephone devices (telephones or speaker phones) that are currently in use by applications. | Number | |
| | Outgoing calls: Indicates the rate at which outgoing calls are made by this computer. | Calls/Sec | |
| | Incoming calls: Indicates the rate at which incoming calls are answered by this computer. | Calls/Sec | |

MONITORING THE MICROSOFT RAS SERVER

| | | | |
|--|---|--------|--|
| | Client applications using telephony services: Indicates the number of applications that are currently using telephony services. | Number | |
| | Current outgoing calls: Indicates the number of outgoing calls that are currently being serviced by this computer. | Number | |
| | Current incoming calls: Indicates the number of incoming calls that are currently being serviced by this computer. | Number | |

Monitoring Microsoft System Management Servers (SMS)

Microsoft Systems Management Server provides a comprehensive solution for change and configuration management for the Microsoft platform, enabling organizations to provide relevant software and updates to users quickly and cost-effectively.

In order to make sure that critical software updates are quickly and readily available to the users, the Microsoft SMS has to be monitored periodically for availability and optimal performance.

eG Enterprise provides administrators with an exclusive *Microsoft SMS* monitoring model that carefully examines the critical services and core functions of the Microsoft SMS, and proactively alerts them to performance aberrations that can adversely impact the user interaction with the server.

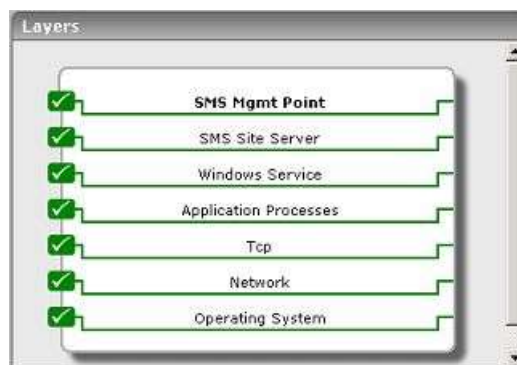


Figure 14.1: The layer model of Microsoft SMS

The sections to come discuss the top 2 layers of the hierarchical layer model depicted by Figure 14.1. The other layers have already been dealt with in the *Monitoring Unix and Windows Servers* document.

14.1 The SMS Site Server Layer

The tests mapped to this layer monitor the health of core components of the Microsoft SMS, such as:

- The Discovery Data Manager
- The Inventory Loader
- The SMS Memory Queue
- The SMS_STATUS_MANAGER
- The Software Inventory Processor



Figure 14.2: The tests associated with the SMS Site Server layer

Besides, the layer also reveals the state of threads executing on the Microsoft SMS, and the effectiveness of its Software Metering feature.

14.1.1 Data Discovery Test

This test monitors the Discovery data manager of SMS. This Data Manager discovers data about the SMS Clients (computers connected to the network and the SMS server).

| | |
|---|---|
| Purpose | Monitors the Discovery data manager of SMS |
| Target of the test | Microsoft SMS |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| | | | |
|--------------------------------------|--|-------------------------|--|
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bad data records processed: Indicates the number of bad (ill-formed or invalid) data records processed by the Discovery Data Manager. | Number | |
| | Data records waiting in the input queue: Indicates the number of SMS Discovery data records waiting in the Discovery Manager's input queue the last time the input queue was scanned minus the number of data records processed till then. | Number | When many data records are written to the input queue, this counter is too low until the Discovery Manager scans the input queue again. This means many data records have been processed in that period. |
| | Total data records processed: Indicates the number of Discovery Data records processed in the last test frequency. | Number | |

14.1.2 Inv Load Test

This test reports metrics pertaining to the Inventory Data Loader of SMS, which loads the client configuration details pertaining to the system hardware.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports metrics pertaining to the Inventory Data Loader of SMS | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |
| | | | |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| | | | |
|------|--|--------|---|
| test | Bad Management Information Files (MIFs) processed: Indicates the number of bad (ill-formed or otherwise invalid) SMS hardware inventory records (in MIF - Management Information Format files) processed by Inventory Data Loader since it was last started. | Number | |
| | MIFs enqueued: Indicates the number of MIF files (containing SMS hardware inventory records) that were waiting in the Inventory Data Loader's input queue the last time Inventory Data Loader scanned the queue, minus the MIF files processed since then | Number | When many MIF files are being written to the input queue, this measure will be too low until Inventory Data Loader scans the input queue again. |
| | MIFs processed: Indicates the number of SMS hardware inventory records (in MIF files) processed by the Inventory Data Loader since it was last started. | Number | |

14.1.3 Memory Queue Test

The MemoryQueue test monitors the health of the SMS memory queue. It is to this SMS Memory Queue thread that a component adds an object when waiting and another component picks the object for its function and removes it from the queue.

| | |
|---|---|
| Purpose | Monitors the health of the SMS memory queue |
| Target of the test | Microsoft SMS |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured |
| Outputs of the test | One set of results for every memory queue thread on the monitored SMS server |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|--|------------------|----------------|
| | Objects dequeued: Indicates the number of objects that the destination component has removed from the queue. | Number | |
| | Objects enqueued: Indicates the number of objects that the source component has added to the queue | Number | |

14.1.4 SMS Status Messages Test

The SmsStatusMsgs test tracks the status messages handled by the SMS_STATUS_MANAGER.

| Purpose | Tracks the status messages handled by the SMS_STATUS_MANAGER | | |
|---|--|------------------|----------------|
| Target of the test | Microsoft SMS | | |
| 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 | | |
| Outputs of the test | One set of results for every SMS component monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | High priority: Indicates the number of SMS status messages replicated to the parent site at high priority by the Status Manager. | Number | |
| | Low priority: Indicates the number of SMS status messages replicated to the parent site at low priority by the Status Manager. | Number | |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| | | | |
|--|---|--------|--|
| | <p>Normal priority:</p> <p>Indicates the number of SMS status messages replicated to the parent site at normal priority by the Status Manager.</p> | Number | |
| | <p>Report app evt log:</p> <p>Indicates the number of SMS status messages reported by the Status Manager to the Windows NT Application Event Log on the site server.</p> | Number | |
| | <p>Database writes:</p> <p>Indicates the number of SMS status messages queued by the Status Manager to be written to the SMS site database.</p> | Number | <p>This number equals the number of status messages actually written to the database, unless Status Manager cannot write to the database (because it is full, for example), in which case the number of queued messages (shown by this counter) will increase even though no messages are being written to the database. (Queued messages are stored as .SQL files in \SMS\Inboxes\Statmgr.box\Retry.) When the database becomes writable again, the queued messages will rapidly be written to it, and this counter will again reflect the actual number of messages written to the database.</p> |

14.1.5 SMS Threads Test

This test reports the state of the SMS threads.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the state of the SMS threads | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| | | | |
|-------------|--|--------|---|
| test | Running threads: Indicates the number of running threads in the SMS Executive (SMSEXEC.EXE) service. | Number | When this measure is associated with a single thread instead of the entire service, its value is zero (the thread is not running) or one (the thread is running). |
|-------------|--|--------|---|

14.1.6 Software Inventory Proc Test

This test reports metrics pertaining to the Software Inventory Processor of SMS. The Software Inventory Processor processes the files produced by the Software Inventory Manager.

| | | | |
|---|--|-------------------------|---|
| Purpose | Reports metrics pertaining to the Software Inventory Processor of SMS | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Bad software inventory records processed: Indicates the number of bad (ill-formed or otherwise invalid) SMS software inventory records (SINVs) processed by the Software Inventory Processor since it was last started. | Number | |
| | Software inventory records waiting in input queue: Indicates the number of SMS software inventory records (SINVs) waiting in the Software Inventory Processor's input queue the last time Software Inventory Processor scanned the queue, minus the SINVs that have been processed since the queue was last scanned. | Number | When many SINVs are being written to the input queue, this counter is too low until Software Inventory Processor scans the input queue again. |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| | | | |
|--|--|--------|--|
| | <p>Total software inventory records processed:</p> <p>Indicates the number of SMS software inventory records (SINVs) processed by Software Inventory Processor since it was last started.</p> | Number | |
|--|--|--------|--|

14.1.7 Software Metering Test

This test monitors the Software Metering feature, which allows one to monitor program usage on client computers. By using software metering, one can collect data about software usage in one's organization. Software metering data can be conveniently summarized to produce useful reports that can help one monitor licensing compliance and plan software purchases in one's organization. Software metering collects detailed information about the programs that you chose to monitor. This includes information about program usage, program users, the program start time, and the length of time it is used.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the Software Metering feature, which allows one to monitor program usage on client computers | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Bad software metering files processed:</p> <p>Indicates the number of bad (ill-formed or otherwise invalid) SMS software metering usage files processed by Software Metering Processor since it was last started.</p> | Number | |

MONITORING MICROSOFT SYSTEM MANAGEMENT SERVERS (SMS)

| | | | |
|--|--|--------|--|
| | Usage files waiting in the input queue: Indicates the number of SMS software metering usage files waiting in the Software Metering Processor's input queue, minus the number of files that have been processed since the queue was last scanned. | Number | |
| | Usage processing threads: Indicates the number of threads the Software Metering Processor is currently using to process incoming SMS software metering usage files. | Number | |
| | Total usage records processed: Indicates the number of software metering records processed by the SWM Processor. | Number | |

14.2 The SMS Mgmt Point Layer

This layer tracks the health of the SMS Management Point components.



Figure 14.3: The tests associated with the SMS Mgmt Point layer

14.2.1 Management Point Data Loader Test

The MgmtPointDataLoader test reports metrics pertaining to the Management Point Data Loader object, which monitors the SMS interactions with the database.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports metrics pertaining to the Management Point Data Loader object, which monitors the SMS interactions with the database | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Connections created: Indicates the number of database connections created by the management point. | Number | |

| | | | |
|--|--|-----------|--|
| | <p>Connections create rate:</p> <p>Indicates the number of database connections created by the management point per second.</p> | Conns/Sec | |
|--|--|-----------|--|

14.2.2 MgmtPointHwInv Test

This test reports metrics pertaining to the Hardware Inventory Manager. The SMS hardware inventory feature automatically collects detailed information about the hardware characteristics of clients in an SMS hierarchy. By using this feature, you can collect a wide variety of information about client computers such as memory, operating system, peripherals, services, and processes that are running on the client computer.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports metrics pertaining to the Hardware Inventory Manager | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Delta reports:</p> <p>Indicates the number of hardware inventory reports marked as Delta.</p> | Number | |
| | <p>Reports data generated:</p> <p>Indicates the size of generated reports.</p> | MB | |
| | <p>Reports processed:</p> <p>Indicates the number of reports processed, successfully or unsuccessfully.</p> | Number | |
| | <p>Reports process rate:</p> <p>Indicates the number of reports processed per second.</p> | Reports/Sec | |

14.2.3 Management Point Policy Manager Test

This test monitors the responses of the SMS Policy Manager to the policy requests of clients.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the responses of the SMS Policy Manager to the policy requests of clients | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Request arrival rate: Indicates the rate at which Policy Assignment requests are arriving at the Policy Manager. | Requests/Sec | |

14.2.4 Management Point Policy Test

This test reports the results of the client requests to the SMS Policy Manager. There are certain SMS policies which download in the client system. This is controlled by the SMS Policy Manager.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the results of the client requests to the SMS Policy Manager. There are certain SMS policies which download in the client system | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | | | |

| | | | |
|-------------|---|--------------|--|
| test | Cache hit rate: Indicates the rate of requests to the Get Policy component that resulted in the policy being served from a cache. | Hits/Sec | |
| | Requests process rate: Indicates the rate of requests to the Get Policy component. | Requests/Sec | |

14.2.5 Management Point Status Manager Test

This test reports metrics pertaining to the Status Manager of SMS. SMS generates status messages to report the activity of components on site systems and clients. A status message is a text string, generated by a component, describing a specific activity performed by the component. In addition, each status message contains important information such as which component generated the message, the exact time that the message was generated, and the severity of the message. Status messages are sent from clients and site systems to the site server and are stored in the SMS site database. You can then view status messages in the SMS Administrator console. Viewing status messages in the SMS Administrator console helps you monitor the activity of the various components, determine the health of SMS, and identify issues that might require your attention.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports metrics pertaining to the Status Manager of SMS. SMS generates status messages to report the activity of components on site systems and clients | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Events processed: Indicates the number of events (i.e. status messages) processed, successfully or unsuccessfully. | Number | |
| | Events process rate: Indicates the number of events (i.e. status messages) processed per second. | Number | |

14.2.6 Management Point Software Inventory Test

This test reports metrics pertaining to the reports generated by the Software Inventory manager of SMS. With the SMS Software Inventory Manager one can collect information about the applications listed in Add or Remove Programs in Control Panel. By using software inventory, one can collect a significantly larger amount of information about client's software.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports metrics pertaining to the reports generated by the Software Inventory manager of SMS | | |
| Target of the test | Microsoft SMS | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – The host for which the test is to be configured | | |
| Outputs of the test | One set of results for every Microsoft SMS server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Delta reports: Indicates the number of Software Inventory reports marked as Delta. | Number | |
| | Reports data generated: Indicates the size of generated reports. | MB | |
| | Reports processed: Indicates the number of reports processed successfully or unsuccessfully. | Number | |
| | Reports process rate: Indicates the number of reports processed per second. | Reports/Sec | |

Externally Monitoring the Active Directory Server

The *Active Directory* server model discussed in Chapter 0 of this document, performs in-depth internal monitoring of the health of an Active Directory (AD) server. However, sometimes, administrators might be denied access to the AD servers to be monitored, and hence might be unable to install agents on them. Such administrators might still want to monitor the availability and responsiveness of the Active Directory. To cater to the needs of these administrators, eG Enterprise offers the *External AD* model (see Figure 15.1).

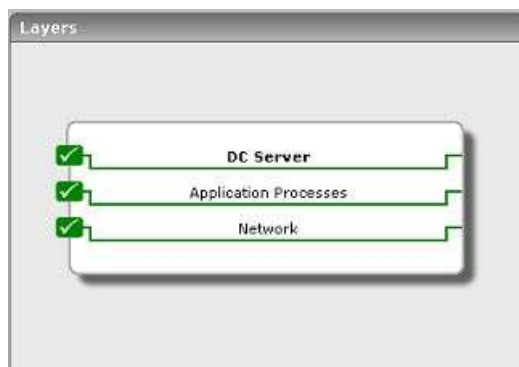


Figure 15.1: Layer model of the External AD server

Besides detecting the non-availability or slow responsiveness of an AD server, this model also runs port availability and network health checks, so as to ensure that all external performance parameters are functioning properly.

15.1 The Network Layer

The **Network** test (see Figure 15.2) associated with this layer performs network-level pings to assess the overall health of the network connection to the AD server.

EXTERNALLY MONITORING THE ACTIVE DIRECTORY SERVER



Figure 15.2: The test associated with the Network layer

For details on the **Network** test, refer to the *Monitoring Unix and Windows Servers* document.

15.2 The Application Processes Layer

Using the **TcpPortStatus** test depicted by Figure 15.3, administrators can externally monitor the availability and responsiveness of the AD server port.



Figure 15.3: The tests associated with the Application Processes layer

Please refer to the *Monitoring Unix and Windows Servers* document for a discussion on the **TcpPortStatus** test.

15.3 The DC Server Layer

By emulating a user request to the AD server, the **ADServer** test associated with this layer (see Figure 15.4) helps determine the availability and responsiveness of the AD server.

EXTERNALLY MONITORING THE ACTIVE DIRECTORY SERVER



Figure 15.4: The tests associated with the DC Server layer

For more details, refer to Chapter 0 of this document.

Monitoring the AD Cluster Service

An active directory (AD) cluster service is a collection of physical AD servers that can act as a single logical server. Requests to a cluster are routed through a virtual cluster server that is assigned a cluster IP address and TCP port. Requests to this server can be handled by any of the individual nodes in the cluster at any given point in time, depending on which node is active at that time.

Since clusters are deployed in environments where 24*7 availability and responsiveness are critical, it is imperative that the performance of the clusters is monitored all the time.

To monitor an Active Directory cluster, an eG external agent is deployed, which emulates a user login to the cluster to determine the availability of the cluster and the speed with which the cluster responds to the emulated request. The emulated requests are directed at the virtual cluster server. Therefore, you need to manage the virtual cluster server as an *AD Cluster* using the eG administrative interface.

Note:

For more details on how eG Enterprise monitors clusters, refer to Chapter 0 of the *eG User Manual*.

The layer model of the *AD Cluster* has been depicted by Figure 16.1 below.

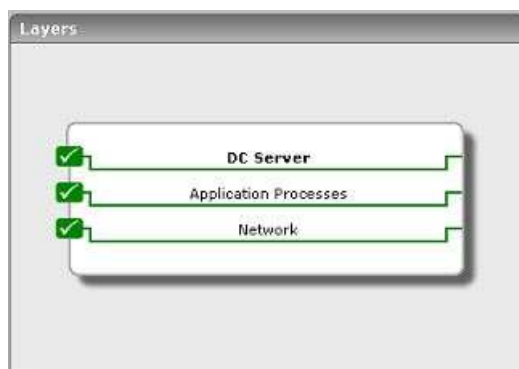


Figure 16.1: Layer model of the AD cluster service

MONITORING THE AD CLUSTER SERVICE

The following section deals only with the first layer of Figure 16.1, as the other layers and the external metrics they report have been dealt with in the previous chapter.

16.1 The DC Server Layer

The **ADServer** test associated with this layer emulates a user login to the cluster to determine its availability and responsiveness. The test sends the emulated request to the virtual cluster server (i.e., the *AD Cluster*), which will promptly forward the request to that node in the cluster that currently owns the cluster server. If at least one node in the cluster is currently active, then the login request will succeed, indicating the good health of the cluster. On the other hand, if none of the nodes in the cluster are active, then the emulated request will fail, indicating the non-availability of the cluster.




Figure 16.2: The tests associated with the DC_SERVER layer

In Chapter 0 of this document, the **ADServer** test has been elaborately discussed. Refer to it for further details.

Monitoring Windows Clusters

Microsoft Cluster Server (MSCS) is software designed to allow servers to work together as computer cluster, to provide failover and increased availability of applications, or parallel calculating power in case of high-performance computing (HPC) clusters (as in supercomputing).

To monitor Windows clusters, eG Enterprise provides a specialized *Windows Cluster* monitoring model. Cluster monitoring enables you to determine the following:

- How many servers (i.e., nodes) are within the cluster? What is their current state?
- How many resource groups have been configured for the cluster? Is any resource group offline currently?
- How many resources have been configured for the cluster, and what is their state?
- Are the network interfaces connecting cluster nodes to each other, available?
- Is any cluster network down?

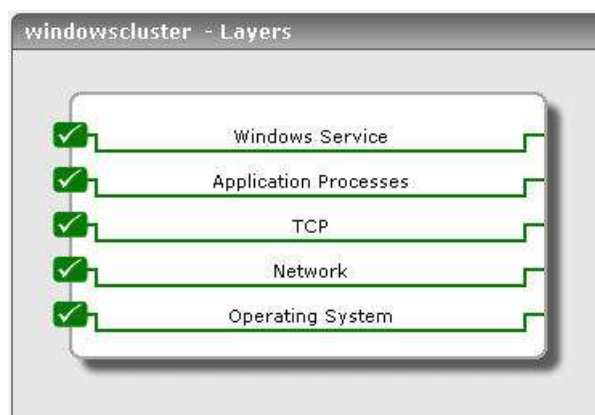


Figure 17.1: The layer model of the IIS web server with the Windows Cluster layer

The tests mapped to the **Windows Service** layer, upon execution, connects to the Windows cluster service that implements the Windows clustering and provides information pertaining to the cluster status.

Note

17.1 The Windows Service Layer

Using the tests mapped to this layer, you can determine the status of the Windows cluster service.



Figure 17.2: The tests mapped to the Windows Service layer

17.1.1.1 Cluster Groups Test

A resource group is a collection of resources, managed by the Cluster service as a single, logical unit. This logical unit is often referred to as a failover unit, because the entire group moves as a single unit between nodes. Resources and cluster elements are grouped logically according to the resources added to a resource group. When a Cluster service operation is performed on a resource group, the operation affects all individual resources contained in the group. Typically, a resource group is created that contains the individual resources required by the clustered program.

Cluster resources may include physical hardware devices, such as disk drives and network cards, and logical items such as IP addresses, network names, and application components.

The ClusterGroups test indicates the current status of the resource groups.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the current status of the resource groups | | |
| Target of the test | | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows application. 4. EXCLUDE CLUSTER OFFLINE GROUPS - Provide a comma-separated list of cluster groups to be excluded from the monitoring scope of this test. | | |
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING WINDOWS CLUSTERS

| | | | |
|------|---|--------|--|
| test | Groups online: Indicates the number of resource groups in the cluster that are currently online. | Number | |
| | Groups offline: Indicates the number of groups in the cluster that are currently offline. | Number | This count includes only those groups in which all resources are offline. |
| | Groups partially online: Indicates the number of groups in the cluster that are partially online. | Number | This count includes only those groups in which some resources are online and some offline. |

17.1.1.2 Cluster Nodes Test

Every server in a cluster is referred to as a *Node*. Using the ClusterNodes test, you can identify how many nodes in the cluster are currently not available.

| | | | |
|---|--|-------------------------|--|
| Purpose | Reports the current status of the nodes in the cluster | | |
| Target of the test | | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the Windows application. | | |
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Is cluster node up?: Indicates the number of nodes that are currently up and running in the cluster. | Number | An online node or a node whose status is 'Up', is an active member of the cluster. It adheres to cluster database updates, contributes input into the quorum algorithm, maintains cluster network and storage heartbeats, and can own and run resource groups. |
| | Is cluster node down?: Indicates the number of nodes in the cluster that are currently offline. | Number | A node whose status is 'down' or 'offline' is considered to be an inactive member of the cluster. The node and its Cluster service might or might not be running. |

MONITORING WINDOWS CLUSTERS

| | | | |
|--|---|--------|--|
| | <p>Is cluster node paused?:</p> <p>Indicates the number of nodes in the cluster that are paused.</p> | Number | <p>This refers to nodes that are active members of the cluster. The node adheres to cluster database updates, contributes input into the quorum algorithm, and maintains network and storage heartbeats, but it cannot accept resource groups. It can support only those resource groups that it currently owns. The paused state enables maintenance to be performed. Online and paused states are treated as equivalent states by the majority of the server cluster components.</p> |
|--|---|--------|--|

17.1.1.3 Cluster Network Interfaces Test

The network adapter or adapters (also known as network interface cards or NICs) on each node in a cluster enables the nodes to communicate with each other and with clients. This test reveals whether/not there are any network interfaces in the cluster that are not functioning properly.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reveals whether/not there are any network interfaces in the cluster that are not functioning properly | | |
| Target of the test | | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the Windows application. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING WINDOWS CLUSTERS

| | | | |
|------|--|--------|---|
| test | Network interfaces that are up: Indicates the number of network interfaces that are currently up and running in the cluster. | Number | |
| | Network interfaces that are down: Indicates the number of network interfaces in the cluster that are currently not running. | Number | Each node in a failover cluster requires network connectivity with the other nodes. Problems with a network adapter or other network device (either physical problems or configuration problems) can interfere with connectivity. Therefore, ideally, the value of this measure should be 0. |

17.1.1.4 Cluster Resources Test

A resource is a physical or logical entity that is capable of being managed by a cluster, brought online, taken offline, and moved between nodes. A resource can be owned only by a single node at any point in time.

This test reports the number of resources in the cluster and their current states.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the number of resources in the cluster and their current states | | |
| Target of the test | | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the Windows application. 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: <ul style="list-style-type: none"> The eG manager license should allow the detailed diagnosis capability Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING WINDOWS CLUSTERS

| | | | |
|--|--|--------|--|
| | <p>Online resources: Indicates the number of resources that are currently online.</p> | Number | |
| | <p>Offline resources: Indicates the number of resources that are currently offline.</p> | Number | |
| | <p>Failed resources: Indicates the number of resources that have failed.</p> | Number | <p>Typically, a resource failure triggers a recovery action, which could be a resource restart or a transfer of the resource to another node.</p> <p>Typically, the Failover Manager and Resource Monitor work together to detect and recover from resource failures. Resource Monitors keep track of resource status by using the resource DLLs to periodically poll resources. Polling involves two steps, a cursory LooksAlive query and a longer, more definitive, IsAlive query. When Resource Monitor detects a resource failure, it notifies Failover Manager and continues to monitor the resource. Failover Manager maintains resources and resource group status. It also performs recovery when a resource fails and invokes Resource Monitors in response to user actions or failures. After a resource failure is detected, Failover Manager performs recovery actions that include restarting a resource and its dependent resources, or moving the entire resource group to another node. The recovery action that is taken is determined by resource and resource group properties, in addition to node availability. During failover, the resource group is treated as the unit of failover. This ensures that resource dependencies are correctly recovered. When a resource recovers from a failure, Resource Monitor notifies Failover Manager. Failover Manager then performs automatic failback of the resource group, based on the configuration of the resource group failback properties.</p> |

17.1.1.5 Cluster Networks Test

A network (sometimes called an interconnect) performs one of the following roles in a cluster:

- A *private network* carries internal cluster communication. The Cluster service authenticates all internal communication, but administrators who are particularly concerned about security can restrict internal communication to physically secure networks.
- A *public network* provides client systems with access to cluster application services. IP Address resources are created on networks that provide clients with access to cluster services.
- A *mixed (public-and-private) network* carries internal cluster communication and connects client systems to cluster application services.
- A network that is not enabled for use by the cluster (that is, neither public nor private) carries traffic unrelated to cluster operation.

Regardless of the role that a network performs, its availability is critical to the smooth functioning of the cluster, as without the network, communication between cluster nodes and between clients and cluster nodes become impossible.

The ClusterNetworks test indicates whether the cluster networks are up or down.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Indicates whether the cluster networks are up or down | | |
| Target of the test | | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows application. 4. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the | Measurement | Measurement Unit | Interpretation |

MONITORING WINDOWS CLUSTERS

| | | | |
|-------------|---|--------|---|
| test | <p>Networks that are up:</p> <p>Indicates the number of cluster networks that are currently available.</p> | Number | |
| | <p>Networks that are down:</p> <p>Indicates the number of cluster networks that are currently unavailable.</p> | Number | <p>If there is only one cluster network available in a cluster and it goes down, the cluster nodes stop communicating with each other. When two nodes are unable to communicate, they are said to be partitioned. After two nodes become partitioned, the Cluster service automatically shuts down on one node to guarantee the consistency of application data and the cluster configuration. This can lead to the unavailability of all cluster resources.</p> <p>Therefore, it is recommended that you configure multiple networks as private or mixed to protect the cluster from a single network failure. For instance, if each node has at least two networks configured, and both are say, mixed networks, the Cluster service can tolerate network failures. In this scenario, the Cluster service can detect a public network adapter failure and fail over all resources that depend on that adapter (through its IP address) to a node where this network is available. This is accomplished because the private network is still functioning properly. If, on the other hand, an adapter on the private network fails, the Cluster service can use the public network for internal communication. This is accomplished because the public network is mixed, allowing both internal and client traffic.</p> |

Monitoring Microsoft Sharepoint

Microsoft Sharepoint is a collection of products and software elements that include, Internet Explorer based collaboration functions, process management modules, search modules and a document-management platform. Sharepoint can be used to host web sites that access shared workspaces, information stores and documents, as well as host defined applications such as wikis and blogs. All users can manipulate proprietary controls called "web parts" or interact with pieces of content such as lists and document libraries.

If any of the services offered by Microsoft Sharepoint malfunction, it could deny users access to critical organizational data, thereby hampering their productivity and obstructing the achievement of business goals. It is therefore imperative that the Microsoft Sharepoint server is monitored 24x7 for performance deficiencies.

eG Enterprise offers two specialized monitoring models - one for each of the Sharepoint versions - *Microsoft Sharepoint 2007* and *Microsoft Sharepoint 2010*.

This chapter discusses both these models in great detail.

18.1 Monitoring Sharepoint 2007

The *Microsoft Sharepoint 2007* monitoring model continuously monitors the performance of the Sharepoint 2007 server, and proactively alerts administrators to issues.

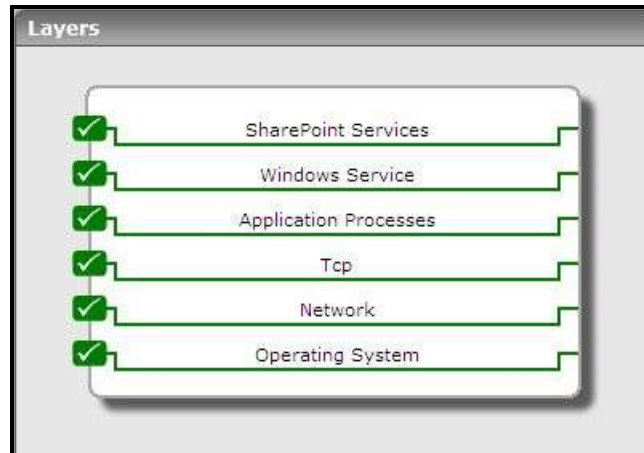


Figure 18.1: The layer model of Sharepoint

Each layer of Figure 18.1 is mapped to a wide variety of tests that report a number of metrics related to the health of the Sharepoint server in question. Using these metrics, the administrators can find quick and accurate answers for the following performance queries:

- Are there too many documents in the first and second queues of the archival plugin? Do these numbers indicate that the crawler is in a starved state?
- Were any error documents returned by the archival plugin?
- How well is the document converter functioning? Are too many conversion requests pending on the converter?
- How is the Excel calculation service performing? Is it responding to requests quickly? How effectively is the service using its cached charts? Are its workbook caches adequately sized?
- Are the Excel Web Access and Excel Web Services components experiencing any slowdowns in request processing?
- Is the content managed by Sharepoint adequately indexed? Are search queries being successfully executed or are too many queries failing?
- Is the gatherer service in a back-off state? If so, why?
- Are your site hit frequency rules very rigid? Are they creating too many delayed documents?
- Are too many threads waiting for documents?
- Are too many threads waiting for a response from the filter process? Is it owing to a network issue or is it because they are bound to a hungry-host?
- Was the gatherer unable to access any documents? If so, how many times?
- Are there too many unprocessed documents on the gatherer?
- Is the Sharepoint Publishing Cache well-tuned? Is the cache hit ratio high?

The sections to come discuss the tests associated with the **Sharepoint Services** layer only, as the remaining layers

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

18.1.1 The Sharepoint Services Layer

Using the tests mapped to this layer, administrators can periodically audit the service levels achieved by the components engaged in the searching and indexing of content managed by Sharepoint. These components include:

- g. The Office Server Search Archival Plugin
- h. The Office Server Search Schema Plugin
- i. The Office Server Search Indexer Catalogs
- j. The Office Server Search Gatherer

Similarly, the layer also sheds light on the core components of the Sharepoint Excel Services – namely, the Excel Calculation Service, the Excel Web Access, and the Excel Web Service.

In addition, the layer monitors the health of the object caches and the document converters on Sharepoint 2007.

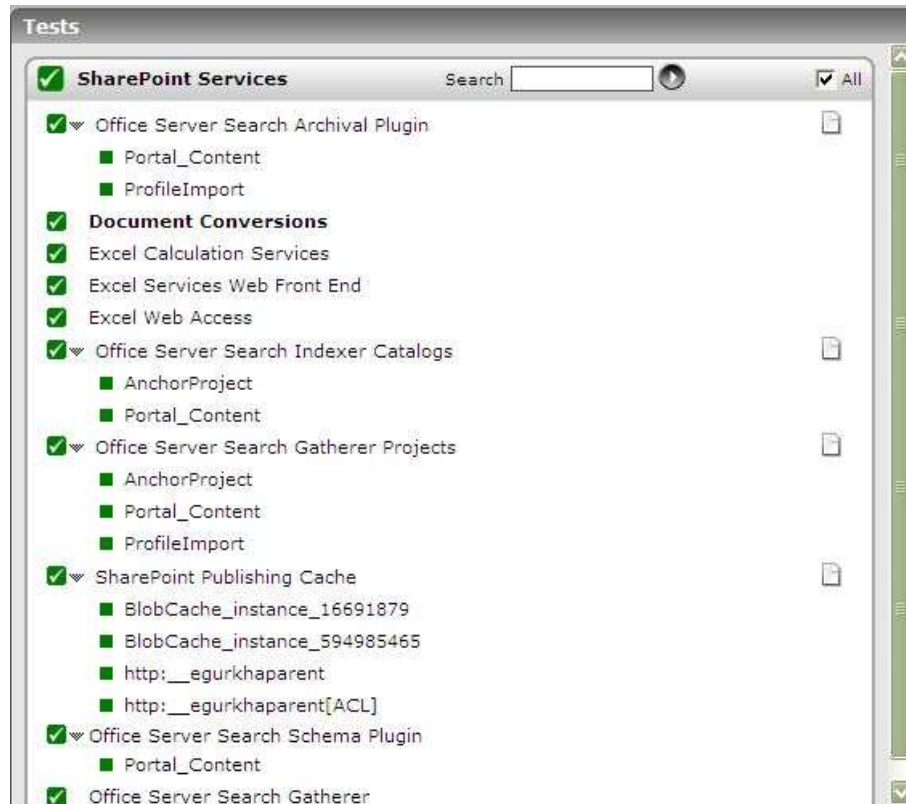


Figure 18.2: The tests mapped to the Sharepoint Services layer

18.1.1.1 Office Server Search Archival Plugin Test

The Search feature of the MOSS 2007 not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites. This is what makes MOSS 2007 that much more valuable to users.

MONITORING MICROSOFT SHAREPOINT

The **Archival** and **Schema** plugins are internal components of the MOSS Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the MOSS search feature is functioning, identify irregularities early, and fine-tune the MOSS server to ensure peak performance of the search engine.

The **Office Server Search Archival Plugin** focuses on the archival plugin component, and helps assess its processing ability.

| | | | |
|---|---|-------------------------|--|
| Purpose | Helps assess the processing ability of the archival plugin component | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed <ol style="list-style-type: none"> 1. HOST - The host for which the test is to be configured 2. PORT – Refers to the port used by the Windows application. | | |
| Outputs of the test | One set of results each for the <i>ProfileImport</i> and <i>Portal_Content</i> instances | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active documents in first queue: Indicates the number of documents that are actively using the first queue of the plugin. | Number | One of the more difficult tasks that a Search admin faces is figuring out how to build out the myriad of crawl schedules needed to keep the content on the Sharepoint server freshly indexed. When you are building out these schedules you will want to keep a close eye on the system and slowly add new schedules to minimize starving the crawl of resources while maxing out the utilization of the crawler. Starvation for Enterprise Search is defined as the crawlers inability to allocate another thread to retrieve the next document in the queue of work. This can be caused by resource (I/O) contention on the SQL machine, too many hosts concurrently participating in the crawl, "hungry" hosts that do not quickly relinquish a thread and finally back-ups (since crawls are paused during this time). The values of these measures typically help determine whether the crawler is in a starved state or not. If they are both consistently at 500 for the Portal_Content instance or 50 for the ProfileImport instance, then you are in a starved state and you are likely to be bottle-necked in SQL for I/O on the Crawl DB drive. Look into tuning SQL for better I/O. |
| | Active documents in second queue: Indicates the number of documents actively using the second queue of the plugin. | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|------------------------------------|
| | Error documents: Indicates the number of documents which currently returned errors from the plugin. | Number | Ideally, this value should be low. |
| | Bulk insert sessions: Indicates the number of active bulk insert sessions to the database server. | Number | |
| | Active queue length: Indicates the number of documents currently available in the active queue. | Number | |
| | Blocked documents: Indicates the number of documents currently waiting for a queue. | Number | |

18.1.1.2 Document Conversions Test

A document converter is a custom executable file that takes a document of one file type, and generates a copy of that file in another file type. For example, a document converter might take a Microsoft Office Excel file and use it to generate a Microsoft Office PowerPoint file. Using document converters, you can transform your content into different versions to suit your business needs.

Because document conversions can be resource intensive, Office Sharepoint Server 2007 relies on two services, DocConversionLoadBalancerService and DocConversionLauncherService, to manage the load balancing, prioritizing, and scheduling of the conversions. When a user initiates a document conversion, either through the user interface or object model, Office Sharepoint Server 2007 passes the document conversion request to these two services. It is the DocConversionLauncherService service that actually calls the document converter. When called, the document converter takes the original file and generates a converted copy. Office Sharepoint Server 2007 then takes the converted copy and performs certain post-processing actions on it. These actions include:

- Adding the metadata from the original file to the converted copy.
- Adding metadata that identifies the original file and document converter used to generate the converted copy.
- Notifying the specified people that the conversion has been performed.
- Placing the converted copy into the same document library as the original file.

This test monitors the document conversion process of the Sharepoint server and enables administrators to determine how well the converter is able to process document conversion requests.

| | |
|----------------|---|
| Purpose | Monitors the document conversion process of the Sharepoint server and enables administrators to determine how well the coverter is able to process document conversion requests |
|----------------|---|

MONITORING MICROSOFT SHAREPOINT

| | | | |
|---|---|-------------------------|--|
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows application. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Incoming E-mail messages processed: Indicates the rate at which e-mail messages have been received and processed by Sharepoint. | E-mails/Sec | |
| | Pending conversions: Indicates the number of document conversions that are currently pending. | Number | Ideally, the value of this measure should be low. A high value for the measure could indicate a processing bottleneck. |

18.1.1.3 Excel Calculation Services Test

Excel Services is built on the Sharepoint products and technologies platform. There are three core components of Excel Services:

- Excel Calculation Service
- Excel Web Access
- Excel Web Service

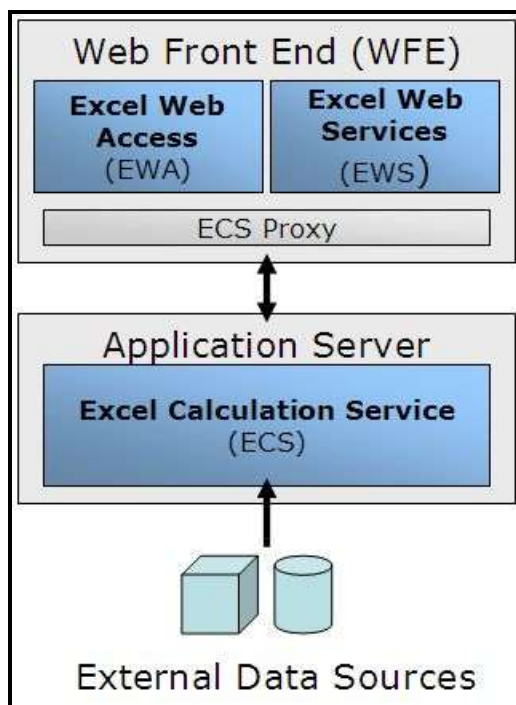


Figure 18.3: Excel services architecture

The role of Excel Calculation Service is to load workbooks, calculate them, call custom code (user-defined functions) and refresh external data. It also maintains the session state for interactivity. Excel Calculation Services maintains a session for the duration of interactions with the same workbook by a user or caller. A session is closed when the caller explicitly closes it or when the session times out on the server. Excel Services caches the opened Excel workbooks, calculation states, and external data query results, for improved performance when multiple users access the same set of workbooks.

In order to determine the quality of the user experience with the Excel Calculation Service, it is essential to know how smooth the user-service interaction is, how quickly the service is able to process the requests, and how effectively the service utilizes its caches. The **Excel Calculation Services** test closely monitors the aforesaid performance parameters, and accurately gauges the health of the service.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Accurately gauges the health of the Excel Calculation Service | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the Windows application. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------------|---|
| | <p>Requests with errors:</p> <p>Indicates the number of requests to the Excel Calculation Service that are returned with errors per second.</p> | Requests/Sec | Ideally, the value of this measure should be low. |
| | <p>Average number of sessions opened:</p> <p>Indicates the average number of sessions opened per second.</p> | Sessions/Sec | c. |
| | <p>Cached charts requested:</p> <p>Indicates the number of charts per second that were provided from a cached image.</p> | Charts/Sec | d. A high value is generally desired for this measure, as it indicates the existence of a well-tuned cache. Such a cache goes a long way in reducing processing overheads. |
| | <p>Active sessions:</p> <p>Indicates the number of currently active sessions on Excel Calculation Services.</p> | Number | e. This value is a good indicator of the current workload on the service. |
| | <p>Average processing time for a request:</p> <p>Indicates the average processing time for a request on Excel Calculation Services.</p> | Secs | f. A high value for this measure or a gradual increase in this value could be indicative of a processing bottleneck on the service. |
| | <p>Average session time:</p> <p>Indicates the average session time.</p> | Secs | g. |
| | <p>Current size of memory cache:</p> <p>Indicates the current size of unused items of the excel calculation service manager in bytes.</p> | MB | h. |
| | <p>Excel calculation service workbook cache size:</p> <p>Indicates the current size of the Excel Calculation Services workbook cache.</p> | MB | i. A high value for this measure indicates that the cache is adequately sized. A poorly-sized cache can adversely impact service performance, especially when multiple users try to access the same set of workbooks. |

| | | | |
|--|--|--------------|----|
| | <p>Rendered charts requested:</p> <p>Indicates the number of chart requests per second.</p> | Charts/Sec | j. |
| | <p>Requests received:</p> <p>Indicates the number of requests received per second on Excel Calculation Services.</p> | Received/Sec | k. |
| | <p>Active requests:</p> <p>Indicates the number of requests being actively processed on Excel Calculation Services.</p> | Number | l. |

18.1.1.4 Excel Services Web Front End Test

The core components of Excel Services - the Excel Web Access, Excel Services, and Excel Calculation Services components - can be divided into components on the Web front-end server and those that live on a back-end application server. The **Web front end** includes **Excel Web Access** and **Excel Web Services**.

Excel Web Services is the Excel Services component that provides programmatic access to its Web service. You can develop applications that call Excel Web Services to calculate, set, and extract values from workbooks, as well as refresh external data connections. Using Excel Web Services, you can incorporate server-side workbook logic into an application, automate the updating of Excel workbooks and create application-specific user interfaces around server-side Excel calculation.

Using the **Excel Services Web Front End** test, you can track the number and rate of requests to the Excel Web Services component.

| | |
|---|--|
| Purpose | Tracks the number and rate of requests to the Excel Web Services component |
| Target of the test | A Sharepoint Server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. |
| Outputs of the test | One set of results for the Sharepoint server monitored |

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|--|------------------|----------------|
| | Active requests: Indicates the current number of requests to the Excel Web Services component. | Number | |
| | Requests rate: Indicates the rate at which requests were received by the Excel Web Services component. | Requests/Sec | m. |

18.1.1.5 Excel Web Access Test

Excel Web Access is an Excel Services Web Part in Office Sharepoint Server 2007 that renders (in other words, creates the HTML for) live Excel workbooks on a Web page, and allows the user to interact with those workbooks and explore them. Excel Web Access is the visible Excel Services component for the user.

This test measures the responsiveness of the Excel Web Access component to user requests.

| Purpose | Measures the responsiveness of the Excel Web Access component to user requests | | |
|---|--|------------------|--|
| Target of the test | A Sharepoint 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 – Refers to the port used by the HOST . | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Average chart image request time: Indicates the average time taken between the request for a chart image and the issuance of the response to the web browser by Excel Web Access. | Secs | An unusually high value for this measure is a cause for concern, as it indicates a slowdown in the responsiveness of the Excel Web Access component. |

| | | | |
|--|---|--------------|----|
| | <p>Chart image request:</p> <p>Indicates the number of requests for chart images that are served by Excel Web Access per second.</p> | Requests/Sec | n. |
| | <p>Excel web access average request time:</p> <p>Indicates the excel web access average request time.</p> | Secs | o. |

18.1.1.6 Office Server Search Indexer Catalogs Test

The MOSS 2007 Search feature is implemented using two MOSS services:

- **Indexing:** Responsible for crawling content sources and building index files.
- **Searching:** Responsible for finding all information matching the search query by searching the index files.

All searching is performed against the index files; if these files do not contain what the user is looking for, there will not be a match. So, the index files are critical to the success of the search feature of MOSS. The search functionality can be described in its simplest form as a Web page where the user defines his or her search query. The index service works together with the searching service to let you search Office Sharepoint Server content.

This test monitors the search queries to every content index on the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by MOSS 2007.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the search queries to every content index on the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by MOSS 2007 | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|---|
| | Failed queries: Indicates the number of queries to the content index that currently failed. | Number | Ideally, this value should be 0. |
| | Succeeded queries: Indicates the number of queries to the content index that succeeded. | Number | p. A high number of successful queries serves as a good indicator of the efficiency of the index and query services provided by Sharepoint. |
| | Queries: Indicates the number of queries currently executing on the content index. | Number | q. |
| | Documents filtered: Indicates the number of documents currently filtered in the content index. | Number | r. |
| | Index size: Indicates the current size of the content index. | Number | s. |

18.1.1.7 Office Server Search Gatherer Test

The MOSS 2007 Search feature is implemented using two MOSS services:

- **Indexing:** Responsible for crawling content sources and building index files.
- **Searching:** Responsible for finding all information matching the search query by searching the index files.

All searching is performed against the index files; if these files do not contain what the user is looking for, there will not be a match. So, the index files are critical to the success of the search feature of MOSS. The search functionality can be described in its simplest form as a Web page where the user defines his or her search query.

The index role can be configured to run on its own MOSS server, or run together with all the other roles, such as the Web service, Excel Services and Forms Services. It performs its indexing tasks following this general workflow:

1. Sharepoint stores all configuration settings for the indexing in its database.
2. When activated, the index will look in Sharepoint's databases to see what content sources to index, and what type of indexing to perform, such as a full or incremental indexing.
3. The index service will start a program called the **Gatherer**, which is a program that will try to open the content that should be indexed.
4. For each information type, the **Gatherer** will need an Index Filter, or **IFilter**, that knows how to read text inside this particular type of information. For example, to read a MS Word file, an IFilter for .DOC is needed.
5. The Gatherer will receive a stream of Unicode characters from the IFilter. It will now use a small program called a Word Breaker; its job is to convert the stream of Unicode characters into words.
6. However, some words are not interesting to store in the index, such as "the", "a", and numbers; the Gatherer

MONITORING MICROSOFT SHAREPOINT

will now compare each word found against a list of Noise Words. This is a text file that contains all words that will be removed from the stream of words.

7. The remaining words are stored in an index file, together with a link to the source. If that word already exists, only the source will be added, so one word can point to multiple sources.
8. If the source was information stored in Sharepoint, or a file in the file system, the index will also store the security settings for this source. This will prevent a user from getting search results that he or she is not allowed to open.
9. Since the success of an indexing operation also depends upon how the **Gatherer** program functions, administrators need to keep their eyes open for irregularities in the functioning of the gatherer, so that such anomalies are detected instantly, and corrected before they can stall the indexing process.

This test monitors the gatherer, and reports issues in its performance (if any).

| | | | |
|---|---|-------------------------|--|
| Purpose | Monitors the gatherer, and reports issues in its performance (if any) | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Documents filtered: Indicates the number of documents filtered per second. | Documents/Sec | If this rate is decreasing over time, you should perform some troubleshooting to find out why your server is not filtering documents. Look for memory issues, processor issues, network issues, or site hit frequency rules that slow the gatherer process. |
| | Filtering threads: Indicates the current number of filtering threads in the system. | Number | t. |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|---------------|---|
| | <p>Threads accessing the network:</p> <p>Indicates the number of threads currently waiting for a response from the filter process.</p> | <p>Number</p> | <p>These threads have sent or are sending their request off to the remote data store and are either waiting for a response or consuming the response and filtering it. You can distinguish the difference between actually waiting on the network versus filtering the document by looking at a combination of CPU usage and Network usage counters.</p> <p>If this number is consistently high then you are either network bound or you are bound by a "hungry" host. If you are <u>not</u> meeting your crawl freshness goals, you can either change your crawl schedules to minimize overlapping crawls or look the remote repositories you are crawling to optimize them for more throughput.</p> |
| | <p>Active queue length:</p> <p>Indicates the number of documents currently waiting for robot threads.</p> | <p>Number</p> | <p>u. If the value of this measure is not 0, then all threads should be filtered.</p> |
| | <p>Admin clients:</p> <p>Indicates the number of currently connected administrative clients.</p> | <p>Number</p> | <p>v.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|--|
| | <p>Reason to back off: A code describing why the gatherer service went into back-off state.</p> | Number | <p>The values that this measure can take and the states they denote are available below:</p> <ul style="list-style-type: none"> 0 - Up and Running. 1 - High system IO traffic. 2 - High notifications rate. 3 - Delayed recovery in progress. 4 - Due to user activity. 5 - Battery low. 6 - Memory low. 99 - Some internal reason. <p>During a back-off period, indexing is suspended. To manually back off the gatherer service, pause the search service. If the search service itself generates the back-off, an event will be recorded and the search service will be paused automatically. There is no automatic restart, so you must manually start the search service in order to end a back-off state. Note that there is little reason to start the search service until you have solved the problem that caused the back-off in the first place.</p> |
| | <p>Threads waiting for plug-ins: Indicates the number of threads currently waiting for plug-ins to complete an operation</p> | Number | <p>These threads have the filtered documents and are processing it in one of several plug-ins. This is when the index and property store are created.</p> <p>If you have a consistently high number for this counter, check the metrics reported by the Office Server Search Archival Plugin test for problem pointers.</p> |
| | <p>Delayed documents: Indicates the number of documents that were currently delayed due to site hit frequency rules.</p> | Number | <p>If you have a plethora of rules and this number is steadily increasing over time, consider relaxing or simplifying your site hit frequency rules.</p> <p>A very high number may indicate a conflict in the rules that the gatherer cannot resolve or follow with efficiency.</p> |

| | | | |
|--|--|---------------|--|
| | <p>Idle threads:</p> <p>Indicates the number of threads that are currently waiting for documents.</p> | Number | <p>These threads are not currently doing any work and will eventually be terminated. If you consistently have a more than <i>Max Threads/Hosts</i> idle threads you can schedule an additional crawl. If this number is 0 then you are starved. Do not schedule another crawl in this time period and analyze the durations of your crawls during this time to see if they are meeting your freshness goals. If your goals are not being met you should reduce the number of crawls.</p> |
| | <p>Hearbeats:</p> <p>Indicates the number of hearbeats per second.</p> | Hearbeats/Sec | <p>A heartbeat occurs once every 10 seconds while the service is running. If the service is not running there will be no heartbeat.</p> |

18.1.1.8 Sharepoint Publishing Cache Test

Object caching Office Sharepoint Server 2007 supports caching of certain page items, such as navigation data and data accessed through cross-list queries. Caching page items reduces the requirement to retrieve field data from the database every time a page is rendered. The caching system also caches complete field data for a page, excluding data for any Web Part controls on the page.

Using the statistics provided by this test, you can fine-tune your cache size, so as to maximize cache hits and minimize object discards.

| | | | |
|---|--|-------------------------|-----------------------|
| Purpose | Helps you fine-tune your cache size, so as to maximize cache hits and minimize object discards | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

| | | | |
|--|---|---------|---|
| | <p>Publishing cache hit ratio:</p> <p>Indicates the ratio of hits to misses on the publishing cache.</p> | Percent | <p>A hit ratio greater than 90% and a low object discard rate are generally good signs that the current size is satisfactory. However, you should also measure user response time for key operations to adjust this setting.</p> |
| | <p>Object discards:</p> <p>Indicates the total number of items that have been removed from the publishing cache since the last measurement period due to cache compaction.</p> | Number | <p>If you set the size too large, you might waste valuable memory for the other caches, such as the ASP.NET output cache if it is used. Certain Web Parts, such as the Content Query Web Part, stores their XSLT stylesheets in the output cache. If the object cache size is set too large, ASP.NET might flush output cache memory to make room for it. CPU usage might increase after the flushing. This is especially important for a system that is running on a 32-bit operating system because each worker process is limited to 2 GB application memory space. If you set the object cache size limit too large, the IIS worker process (w3wp) can run out of memory.</p> |

18.1.1.9 Office Server Search Schema Plugin Test

The Search feature of the MOSS 2007 not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites. This is what makes MOSS 2007 that much more valuable to users.

The **Archival** and **Schema** plugins are internal components of the MOSS Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the MOSS search feature is functioning, identify irregularities early, and fine-tune the MOSS server to ensure peak performance of the search engine.

The **Office Server Search Schema Plugin** test focuses on the schema plugin component, and helps assess its processing ability.

| | |
|---|---|
| Purpose | Focuses on the schema plugin component, and helps assess its processing ability |
| Target of the test | A Sharepoint Server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. |
| Outputs of the test | One set of results for the Sharepoint server monitored |

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|--|------------------|----------------------------------|
| | Aliases mapped: Indicates the total number of aliases which have been currently mapped to the schema. | Number | |
| | Duplicate aliases: Indicates the number of aliases that the schema currently ignored as they are duplicates. | Number | |
| | Refresh count: Indicates the number of aliases that have been refreshed from the database, currently. | Number | |
| | Error documents: Indicates the number of documents that have currently returned errors from the plug-in. | Number | Ideally, this value should be 0. |

18.1.1.10 Office Server Search Gatherer Projects Test

As already mentioned, the indexing service will start a program called the **Gatherer**, which is a program that will try to open the content that should be indexed. Using an iFilter, the **Gatherer** reads the content as Unicode characters, converts the characters into words, identifies words that are worth indexing, and stores them in the content indexes.

For each content index, this test reports critical performance statistics revealing the content processing ability of the gatherer.

| | |
|---|---|
| Purpose | Reports critical performance statistics revealing the health of the gatherer |
| Target of the test | A Sharepoint Server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--------------------------------------|---|-------------------------|--|
| Outputs of the test | One set of results for each content index on the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Documents added: Indicates the number of document additions per second. | Documents/Sec | |
| | Error: Indicates the number of filtered documents which returned an error per second. | Documents/Sec | A low value is typically desired for this measure. |
| | Retries: Indicates the total number of times that access to a document has been retried. | Number | A high value of this measure indicates that the gatherer is attempting to access a document numerous times, without success. You should check the gatherer logs and identify the problem document. Then ensure that it has the correct extension and that you have the correct IFilter for it. |
| | Incremental crawls: Indicates the number of incremental crawls currently in progress. | Number | |
| | Waiting documents: Indicates the current queue size of unprocessed documents in the gatherer. | Number | A high value of this measure could indicate a processing bottleneck on the gatherer. If this measure returns the value 0 on the other hand, it could indicate that the gatherer is idle. |

18.2 Monitoring Sharepoint 2010/2013

Figure 18.4 depicts the *Sharepoint 2010/2013* monitoring model.

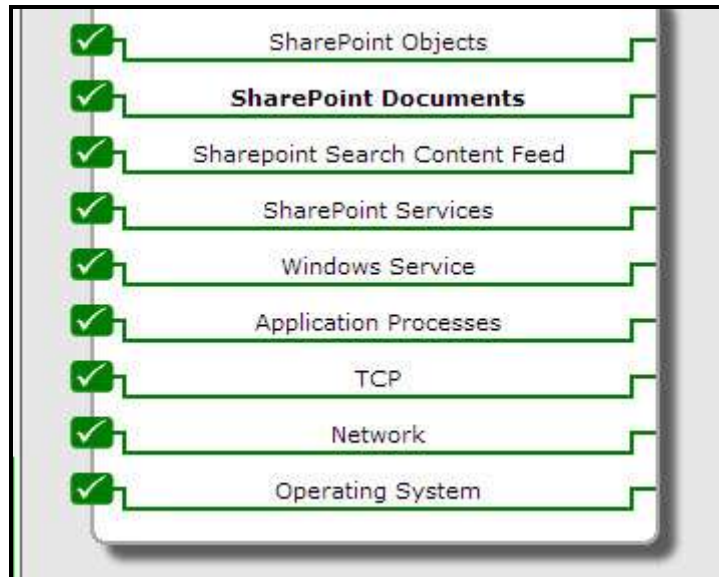


Figure 18.4: The layer model of Microsoft Sharepoint 2010

Each layer of Figure 18.4 is mapped to a variety of tests that periodically check the health of the core components and services of the Sharepoint 2010 server. Using the metrics reported by these tests, administrators can find quick and accurate answers for the following performance queries:

- Has the archival plugin marked too many documents for retry?
- Are too many documents in the archival plugin waiting for a queue?
- Have any errors occurred in index propagation?
- Is index reception error-free?
- Did any search query fail?
- Is query execution taking too long? If so, where is the query spending maximum time?
- Is the query CPU-intensive? If so, where is the query spending the maximum CPU time?
- Is any Sharepoint Foundation process overloaded? If so, which one is it?
- Is any Sharepoint Foundation process taking too long to execute requests? Which process is it?
- Which process is taking too much time to execute queries?
- Is the schema plugin able to process documents and properties quickly?

MONITORING MICROSOFT SHAREPOINT

- Are there too many idle threads on the Sharepoint server?
- Is any thread waiting for a network response from the filter process?
- Have too many servers timed out?
- Was any slowdown noticed in document filtering? Is it due to site hit frequency rules? If so, how many documents were affected as a result?
- Is filtering failing for any document?

The sections that follow will only discuss the **Sharepoint Services** layer of Figure 18.4. The other layers have already been dealt with in the *Monitoring Unix and Windows Servers* document.

18.2.1 The Sharepoint Documents Layer

Using the tests mapped to this layer, you can closely monitor the growth in the number and size of document libraries, documents, and lists.




Figure 19.1: The tests mapped to the Sharepoint Documents Layer

18.2.1.1 Sharepoint Document Libraries and Lists Test

Document libraries are collections of files that you can share with team members on a Web based on Microsoft Windows SharePoint Services. For example, you can create a library of common documents for a project, and team members can use their Web browsers to find the files, read them, and make comments. Users with Microsoft Office 2003 can check out and edit the files as if they resided on a local or network drive.

A list in SharePoint is used to store data across columns in separate rows. You can think of a list as a table in a database that will have columns and rows. You can think of a list as a table in a database that will have columns and rows. You can also think of it as a spreadsheet with columns and rows. Items such as issues, software bugs, employee addresses, phone numbers, web site links or pretty much anything else can be stored.

To ensure that all the web applications deployed on the Sharepoint farm have adequate storage resources at their disposal, administrators must make sure that document libraries and lists used by the web applications do not grow

MONITORING MICROSOFT SHAREPOINT

uncontrollably, both in number and in size. For this, administrators must keep a close watch on the growth of the document libraries and lists. This is where the **Sharepoint Document Libraries and Lists** test helps! This test reports the total number of document libraries and lists created on Sharepoint, tracks the rate at which these numbers are growing, and promptly alerts administrators to an abnormal increase in the number of document libraries and lists. In addition, the test also measures the size of document libraries from time to time, and intimates administrators if the size increases unexpectedly! The detailed diagnosis of this test also reports the top-10 document libraries and lists in terms of size, thus leading administrators to those libraries and lists that could be draining the storage resources of Sharepoint.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Reports the total number of document libraries and lists created on Sharepoint, tracks the rate at which these numbers are growing, and promptly alerts administrators to an abnormal increase in the number of document libraries and lists. In addition, the test also measures the size of document libraries from time to time, and intimates administrators if the size increases unexpectedly! The detailed diagnosis of this test also reports the top-10 document libraries and lists in terms of size, thus leading administrators to those libraries and lists that could be draining the storage resources of Sharepoint. | | |
| Target of the test | A Sharepoint Server 2010/2013 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the Sharepoint Server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|---------|---|
| | <p>Number of document libraries:</p> <p>Indicates the total number of document libraries on the Sharepoint server.</p> | Number | <p>A consistent increase in the value of this measure could indicate that new document libraries are regularly created on Sharepoint. You may want to check how much space these new libraries are consuming to understand the true impact of this addition on storage resources.</p> <p>You can use the detailed diagnosis of this measure to identify the top-10 document libraries in terms of size – i.e., space usage.</p> |
| | <p>Documents in document libraries:</p> <p>Indicates the total number of documents in all document libraries on Sharepoint.</p> | Number | <p>A consistent increase in the value of this measure could indicate the influx of new documents into existing document libraries or the creation of new libraries with a new set of documents. You may want to check how much space these new documents are consuming to understand the true impact of this addition on storage resources.</p> |
| | <p>Size of document libraries:</p> <p>Indicates the total size of all the document libraries on Sharepoint.</p> | MB | <p>A consistent increase in the value of this measure could be attributed to the addition of new document libraries, new documents, and large-sized documents.</p> |
| | <p>Average number of documents per document library:</p> <p>Indicates the average number of documents per library.</p> | Number | |
| | <p>Document library growth rate:</p> <p>Indicates the percentage growth in the number of document libraries handled by Sharepoint, since the last measurement period.</p> | Percent | <p>A consistent increase in the value of this measure could indicate that new document libraries are regularly created on Sharepoint. You may want to check how much space these new libraries are consuming to understand the true impact of this addition on storage resources.</p> |
| | <p>Lists count:</p> <p>Indicates the number of lists on Sharepoint.</p> | Number | <p>A consistent increase in the value of these measures could indicate that new lists are regularly created on Sharepoint. You may want to check how much space these new lists are consuming to understand the true impact of this addition on storage resources. You can use the detailed diagnosis of the <i>Lists</i> measure to identify the top-10 Sharepoint lists in terms of size – i.e., space usage.</p> |
| | <p>Lists growth rate:</p> <p>Indicates the percentage growth in the number of lists on Sharepoint, since the last measurement period.</p> | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|--|
| | <p>Attachments:</p> <p>Indicates the number of attachments on Sharepoint.</p> | Number | |
|--|--|--------|--|

The detailed diagnosis of the *Number of document libraries* measure lists the top 10 libraries in Sharepoint with the maximum number of documents. Using this information, you can quickly identify that document library with the highest document count and also figure out the **PARENTWEBURL** of the web application with which the library is associated. If that web application grows abnormally in size or count of documents, this information will lead administrators to the exact document library that is responsible for it.

| List of Top 10 Document Library | | | | |
|---------------------------------|-----------|-------------|-----------|------------------------|
| TIME | TITLE | DESCRIPTION | ITEMCOUNT | PARENTWEBURL |
| Jan 30, 2014 06:42:11 | | | | |
| | Documents | - | 8 | / |
| | Documents | - | 4 | / |
| | Documents | - | 3 | /sites/mysites |
| | Documents | - | 2 | /sites/testcomplete |
| | Documents | - | 1 | /sites/egInnovations |
| | Documents | - | 1 | /site |
| | Documents | - | 0 | /sites/quota |
| | Documents | - | 0 | /sites/new_site_privat |
| | Documents | - | 0 | /sites/test |

Figure 19.2: The detailed diagnosis of the Number of document libraries measure

The detailed diagnosis of the *Lists count* measure displays the top 10 lists in Sharepoint with the maximum number of items. Using this information, you can quickly identify that list with is most heavily populated and also figure out the **PARENTWEBURL** of the web application with which the list is associated. If that web application grows abnormally, this information will lead administrators to the exact list that may be responsible for it.

| List of Top 10 Lists | | | | |
|-----------------------|----------------|---|-----------|--------------|
| TIME | TITLE | DESCRIPTION | ITEMCOUNT | PARENTWEBURL |
| Jan 30, 2014 06:42:11 | | | | |
| | Composed Looks | Use this list to store composed looks. These looks can be applied to this site by navigating to Site Settings and choosing Change the look. | 18 | /my |

Figure 19.3: The detailed diagnosis of the Lists count measure

18.2.1.2 Sharepoint Docs and Pages Test

Documents are stored within a document library in Sharepoint. Documents add to the size of the sites, site collections, and web applications they are associated with. Significant and rapid spikes in the number and size of documents on the Sharepoint server can hence cause sites, site collections, and ultimately, web applications to grow in size exponentially; in the long run, this may result in a severe space crunch in the content database. This is why, administrators need to keep a close watch on the number of documents handled by the Sharepoint server and the space resources they use. To achieve this, administrators can use the **Sharepoint Docs and Pages** test! This test periodically monitors the number and size of documents in the Sharepoint server, reports abnormal document growth, and thus warns administrators of potential space contentions well before they actually occur!

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Periodically monitors the number and size of documents in the Sharepoint server, reports abnormal document growth, and thus warns administrators of potential space contentions well before they actually occur | | |
| Target of the test | A Sharepoint Server 2010/2013 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for the Sharepoint Server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|---------|---|
| | <p>Number of documents in Sharepoint:</p> <p>Indicates the total number of documents in the Sharepoint server.</p> | Number | A consistent increase in the value of this measure could indicate that new documents are created in Sharepoint at regular intervals. You may want to check how much space these new documents are consuming to understand the true impact of this addition on storage resources. |
| | <p>Versions:</p> <p>Indicates the total number of document versions in Sharepoint.</p> | Number | A consistent increase in the value of this measure could indicate that newer versions of one/more existing documents are now available in Sharepoint. This in turn implies that many outdated/obsolete documents may also exist in Sharepoint. In the event of rapid growth in document count, you may want to delete the stale versions of documents so as to control the growth and make space for newer documents. |
| | <p>Size of all documents:</p> <p>Indicates the total size of all the documents in Sharepoint.</p> | MB | A consistent increase in the value of this measure could be attributed to the addition of new documents and/or large-sized documents. |
| | <p>Average size of a document:</p> <p>Indicates the average size of a document.</p> | MB | With the help of the value of this measure, you can ascertain whether/not Sharepoint is the container for documents of large sizes. |
| | <p>Documents growth rate:</p> <p>Indicates the percentage growth in the number of documents in Sharepoint, since the last measurement period.</p> | Percent | A consistent increase in the value of this measure could indicate there is a consistent addition of new documents to Sharepoint. Compare the value of this measure with that of the <i>Versions</i> measure to understand whether the addition of newer 'versions' of existing documents is in any way contributing to the growth rate. If so, you may want to delete older versions of documents and unnecessary documents to curb the growth. |
| | <p>Number of file formats stored:</p> <p>Indicates the total number of file formats stored in Sharepoint.</p> | Number | Use the detailed diagnosis of this measure to know which file formats are stored in Sharepoint. |

18.2.2 The Sharepoint Objects Layer

The tests mapped to this layer promptly capture the sporadic spikes or steady growth in the contents of the critical Sharepoint data containers such as content databases, sites and site collections, and web applications. Overgrown applications and objects responsible for the uncontrollable growth can thus be isolated.

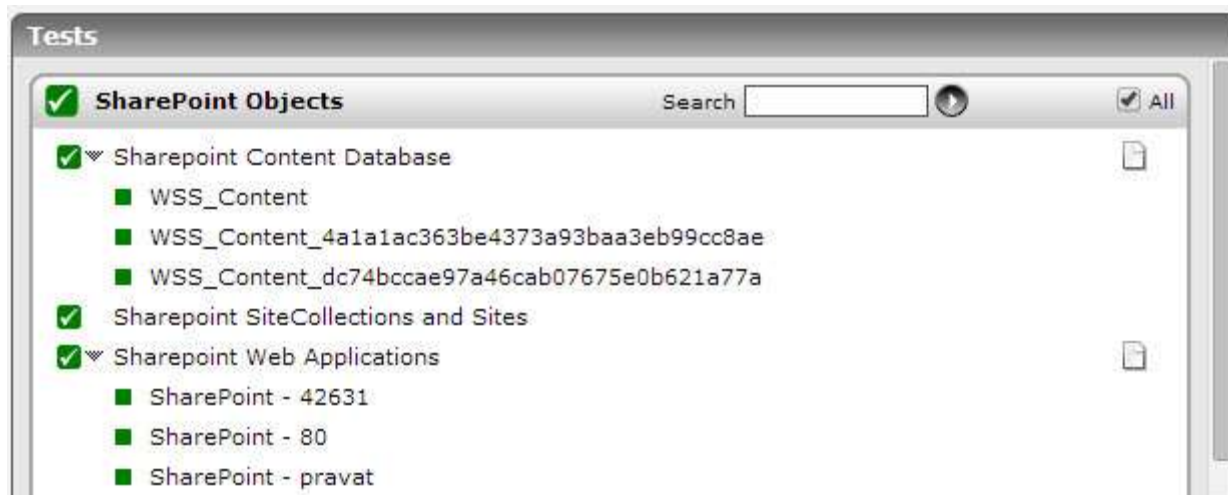


Figure 19.4: The tests mapped to the Sharepoint Objects layer

18.2.2.1 Sharepoint Content Databases Test

Content databases are used to store SharePoint data. This data consists of sites, permissions, documents, lists, etc.

A content database is closely related to two other SharePoint objects: a web application and a site collection. A web application is backed by an Internet Information Services (IIS) web site, and it contains one or more content databases which can contain one or more site collections. Think of a site collection as the "container" for data and also as the security boundary. Security is defined at the site collection level allowing administrators to control access to the sites and their data.

Content databases can grow pretty quickly, and if this growth is not tracked and controlled, users may be left with no space for Sharepoint data. Sharepoint administrators should hence prudently and proactively plan their data storage needs, accordingly size the content databases, and effectively manage the space available in the databases, so that manageability, performance, and reliability issues do not arise. This is where the **Sharepoint Content Databases** test helps!

Besides reporting the state of each content database where Sharepoint data is stored, this test also monitors the size, usage, and growth of every database, thus pointing administrators to those databases that are over-used or are exhibiting alarming growth patterns! In addition, the test provides hints for enhancing the overall performance of the content databases – will it help to cleanup the orphaned items? should the recycle bin storage space be reduced? should the content database host fewer site collections?

| | |
|---------------------------------|--|
| Purpose | Besides reporting the state of each content database where Sharepoint data is stored, this test also monitors the size, usage, and growth of every database, thus pointing administrators to those databases that are over-used or are exhibiting alarming growth patterns! In addition, the test provides hints for enhancing the overall performance of the content databases – can the recycle bin storage space be reduced? should the content database host fewer site collections? |
| Target of the test | A Sharepoint Server 2010/2013 |
| Agent deploying the test | An internal agent |

MONITORING MICROSOFT SHAREPOINT

| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | | | | | | | |
|---|---|-------------------------|---|---------------|---------------|-----|---|----|---|
| Outputs of the test | One set of results for each content database of the Sharepoint Server being monitored | | | | | | | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation | | | | | | |
| | <p>Is database in use?:</p> <p>Indicates whether/not this content database is in use.</p> | | <p>The values that this measure can report and their corresponding numeric values are listed in the table below:</p> <table border="1" data-bbox="932 695 1414 842"> <thead> <tr> <th style="text-align: center;">Measure Value</th> <th style="text-align: center;">Numeric Value</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Yes</td> <td style="text-align: center;">1</td> </tr> <tr> <td style="text-align: center;">No</td> <td style="text-align: center;">0</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the usage state of the content database. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> | Measure Value | Numeric Value | Yes | 1 | No | 0 |
| Measure Value | Numeric Value | | | | | | | | |
| Yes | 1 | | | | | | | | |
| No | 0 | | | | | | | | |

| | | | |
|--|---|-----------|---|
| | <p>Content database size: Indicates the current size of this content database.</p> | <p>GB</p> | <p>Microsoft recommends that no content database be more than 200 GB in size.</p> <p>Content databases of up to 4 TB are supported when the following requirements are met:</p> <ul style="list-style-type: none"> • Disk sub-system performance of 0.25 IOPs per GB. 2 IOPs per GB is recommended for optimal performance. • You must have developed plans for high availability, disaster recovery, future capacity, and performance testing. <p>You should also carefully consider the following factors:</p> <ul style="list-style-type: none"> • Requirements for backup and restore may not be met by the native SharePoint Server backup for content databases larger than 200 GB. It is recommended to evaluate and test SharePoint Server solutions to determine the best solution for your specific environment. • It is strongly recommended to have proactive skilled administrator management of the SharePoint Server and SQL Server installations. • The complexity of customizations and configurations on SharePoint Server may necessitate refactoring (or splitting) of data into multiple content databases. Seek advice from a skilled professional architect and perform testing to determine the optimum content database size for your implementation. Examples of complexity may include custom code deployments, use of more than 20 columns in property promotion, or features listed as not to be used in the over 4 TB section below. |
|--|---|-----------|---|

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|---------|---|
| | | | <ul style="list-style-type: none"> • Refactoring of site collections allows for scale out of a SharePoint Server implementation across multiple content databases. This permits SharePoint Server implementations to scale indefinitely. This refactoring will be easier and faster when content databases are less than 200 GB. • It is suggested that for ease of backup and restore that individual site collections within a content database be limited to 100 GB. |
| | <p>Disk space usage of content database:</p> <p>Indicates the percentage disk space in the SQL server that is used by the content database.</p> | Percent | <p>A high value for this measure is a cause for concern, as it indicates excessive disk space consumption by a content database.</p> <p>Compare the value of this measure across content databases to identify that database which is eroding the disk space of the SQL server.</p> |

| | | | |
|--|--|----------------|--|
| | <p>Content database growth rate:</p> <p>Indicates the percentage growth in the size of this content database since the last measurement period.</p> | <p>Percent</p> | <p>A consistent rise in the value of this measure is a sign that the content database is growing rapidly!</p> <p>To curb this growth, you may want to consider the following measures:</p> <ul style="list-style-type: none"> • Use an ootb Record Center as an archive for old content: The users must manually send each document to the RC using e.g. move and leave a link; note that only the latest major version with metadata is kept – all version history is lost. The information management policies supported by SharePoint for retention and disposition can be used to automate the cleanup. As the RC has its own content databases, the live collaboration databases will grow slower or even shrink as outdated information is moved to the archive. Keeping the live databases small ensures shorter recovery time; while the recovery time for the archived content can be considerable, but not business critical. Search must be configured appropriately to cover both live and archived content. • Use a third-party archiving solution for SharePoint. This has the same pros & cons as the previous option, but the functionality is probably better in relation to keeping version history and batch management of outdated content. Search must be configured appropriately to cover both live and archived content. • Use a third-party remote blob storage (RBS) solution for SharePoint so that documents are registered in the database, but not stored there. This gives smaller content databases, but more complicated backup and recovery as the content now resides both in databases and on disk. Provided that you don't lose both at the same time, the recovery time should be shorter. Search will work as before, as all content is still logically in the "database". |
| | | <p>352</p> | <p>D) Use powershell scripts or other code to implement the disposition of outdated content. The script can e.g. copy old documents to disk and delete old versions from the content database; the drawback</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|---------------|--|
| | | | <ul style="list-style-type: none"> Use powershell scripts or other code to implement the disposition of outdated content. The script can e.g. copy old documents to disk and delete old versions from the content database; the drawback being that all metadata will be lost and there is no link left in SharePoint. The databases size will shrink as data is actually deleted, and backup and recovery is more complicated as content is now both in the database and on disk. Search can be configured to also crawl and index the files on disk, but content ranking will suffer as the valuable metadata is lost. |
| | <p>Orphaned items in this content database:</p> <p>Indicates the number of orphaned sites in this content database.</p> | <p>Number</p> | <p>An Orphaned Site is where SharePoint only has partial information and not a complete set of data for a given site collection in your Windows SharePoint Services or SharePoint Portal Server content databases or configuration databases. The site may in fact still be viewable via the browser, but you may notice that many things are broken.</p> <p>If the <i>Content database growth rate</i> measure is increasing consistently, you may want to check the variations in the value of this measure over the same time period to figure out whether/not the existence of too many orphan sites is contributing to the growth in the size of the content database. If so, you may want to cleanup the orphan sites to right-size your database and to ensure optimum performance.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|----------------|--|
| | <p>Site limit of content database:</p> <p>Indicates the maximum number of site collections that this content database can host.</p> | <p>Number</p> | <p>Microsoft strongly recommends limiting the number of site collections in a content database to 5,000. However, up to 10,000 site collections in a database are supported. Note that in a content database with up to 10,000 total site collections, a maximum of 2,500 of these can be non-Personal site collections. It is possible to support 10,000 Personal site collections if they are the only site collections within the content database.</p> <p>These limits relate to speed of upgrade. The larger the number of site collections in a database, the slower the upgrade with respect to both database upgrade and site collection upgrades.</p> <p>The limit on the number of site collections in a database is subordinate to the limit on the size of a content database that has more than one site collection. Therefore, as the number of site collections in a database increases, the average size of the site collections it contains must decrease.</p> <p>Exceeding the 5,000 site collection limit puts you at risk of longer downtimes during upgrades. If you plan to exceed 5,000 site collections, Microsoft recommends that you have a clear upgrade strategy to address outage length and operations impact, and obtain additional hardware to speed up the software updates and upgrades that affect databases.</p> |
| | <p>Configured site limit usage:</p> <p>Indicates the percentage of the configured site limit that is used by the content database.</p> | <p>Percent</p> | <p>A value close to 100% indicates that the configured site limit is about to be reached.</p> <p>By comparing the value of this measure across content databases, you can easily identify the database that hosts too many site collections. You may then have to reassess the ability of that content database to handle additional site collections, and accordingly decide whether to reconfigure the site limit or reduce the number of site collections hosted by the database.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|----------------|--|
| | <p>Recycle bin storage space in this content database:</p> <p>Indicates the space used by the items present in the second stage recycle bin of this content database.</p> | <p>MB</p> | <p>Recycle Bins are used to help users protect and recover data. Microsoft SharePoint Server 2010 supports two stages of Recycle Bins: the first-stage Recycle Bin and second-stage Recycle Bin.</p> <p>When a user deletes an item, the item is automatically sent to the first-stage Recycle Bin. By default, when an item is deleted from the first-stage Recycle Bin, the item is sent to the second-stage Recycle Bin.</p> <p>A high value for this measure could indicate that a large amount of deleted data resides in the second stage recycle bin, unnecessarily consuming disk space and increasing the size of the content database.</p> |
| | <p>Recycle bin storage space growth rate:</p> <p>Indicates the percentage growth in the space used in the second stage recycle bin of this content database, since the last measurement period.</p> | <p>Percent</p> | <p>A consistent increase in the value of this measure indicates that deleted data is steadily accumulating in the recycle bin; this is a cause of concern, as data in the second stage recycle bin can add megabytes to the overall size of the content database!</p> <p>Every site collection has a second stage recycle bin and the size of this bin must not grow beyond 50 percent of the quota set for that site collection. You may want to reduce this percentage to ensure that the recycle bin does not grow too unwieldy and impact the size and performance of the content database.</p> |

18.2.2.2 Site Collections and Sites Test

A site collection is made up of one top-level site and all sites below it. As shown in the following figure, it is the top level of organization in a SharePoint 2013 web application.




Figure 18.5: Site Collections and Sites

The number of site collections you can have in a single web application depends on the capacity of your server infrastructure.

From an architecture standpoint, all the content of a site collection must be stored in a single content database. You cannot have a site collection's content spread out across multiple content databases. Content databases scale with your infrastructure capacity so site collections can share a content database. A site collection can exist in only one content database, but one content database can host the content for multiple site collections. Similarly, any given SharePoint 2013 site can only exist in one site collection, but a site collection can host a multitude of sites. A site cannot exist outside of a site collection.

The number of site collections and sites sharing a single content database can impact the size of the database and its performance; administrators should therefore exercise restraint when associating sites and site collections with a content database. In addition, the amount of content that the sites and site collections store in their content database is also a key factor influencing the size of the content database. Variations to these two parameters – count and size - hence need to be closely monitored, so that administrators can proactively detect abnormal growth in the size of the content databases, isolate the site collections and sites that may be contributing to this, and take measures to fine-tune the site and site collection configurations to ensure peak performance of the content databases. The **Site Collections and Sites** test aids administrators in this endeavor!

MONITORING MICROSOFT SHAREPOINT

This test captures the total number of site collections and sites on the Sharepoint server / farm and reports whether/not these numbers exceed the permissible limits. In addition, the test also tracks changes in the size of the site collections and sites over time, and promptly intimates administrators if the actual size is about to reach/exceed the size quota set for the site collection. In the process, the test points you to those site collections that are growing rapidly and the sites that may be contributing to their growth. If administrators initiate measures to curb the abnormal growth in the number or the size of the site collections and sites, they can once again take the help of this test to understand which sites and site collections are the least popular, so that such sites and site collections can be marked as probable targets for deletion or trimming.

| | |
|---------------------------------|---|
| Purpose | Captures the total number of site collections and sites on the Sharepoint server / farm and reports whether/not these numbers exceed the permissible limits. In addition, the test also tracks changes in the size of the site collections and sites over time, and promptly intimates administrators if the actual size is about to reach/exceed the size quota set for the site collection. In the process, the test points you to those site collections that are growing rapidly and the sites that may be contributing to their growth. If administrators initiate measures to curb the abnormal growth in the number or the size of the site collections and sites, they can once again take the help of this test to understand which sites and site collections are the least popular, so that such sites and site collections can be marked as probable targets for deletion or trimming |
| Target of the test | A Sharepoint Server 2010/2013 |
| Agent deploying the test | An internal agent |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------------------------------|------------------------------|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. LEAST ACTIVE SITE COLLECTION DAYS – If a site collection is not modified for a duration exceeding the value (in days) specified here, then this test will count that site collection as a <i>Least active site collection</i>. 5. LEAST ACTIVE SITE DAYS - If a site is not modified for a duration exceeding the value (in days) specified here, then this test will count that site as a <i>Least active site</i>. 6. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 7. 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: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| <p>Outputs of the test</p> | <p>One set of results for the Sharepoint Server being monitored</p> | | |
| <p>Measurements made by the test</p> | <p>Measurement</p> | <p>Measurement Unit</p> | <p>Interpretation</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|---------------|--|
| | <p>Site collections:</p> <p>Indicates the number of site collections in the Sharepoint environment.</p> | <p>Number</p> | <p>The maximum recommended number of site collections per farm is: Personal Sites - 500,000, Other site templates - 250,000. The Sites can all reside on one web application, or can be distributed across multiple web applications.</p> <p>Note that this limit is affected by other factors that might reduce the effective number of site collections that can be supported by a given content database. Care must be exercised to avoid exceeding supported limits when a container object, such as a content database, contains a large number of other objects. For example, if a farm contains a smaller total number of content databases, each of which contains a large number of site collections, farm performance might be adversely affected long before the supported limit for the number of site collections is reached.</p> |
| | <p>Total size of site collections:</p> <p>Indicates the total size of all site collections in the Sharepoint environment.</p> | <p>MB</p> | <p>A site collection can be as large as the content database size limit for the applicable usage scenario.</p> <p>For more information about the different content database size limits for specific usage scenarios, see the Content database limits discussed in the Interpretation column of the <i>Content database size</i> measure of the Sharepoint Content Database test.</p> <p>In general, Microsoft strongly recommends limiting the size of site collections to 100 GB for the following reasons:</p> <ul style="list-style-type: none"> • Certain site collection actions, such as site collection backup/restore, cause large SQL Server operations which can affect performance or fail if other site collections are active in the same database. • SharePoint site collection backup and restore is only supported for a maximum site collection size of 100 GB. For larger site collections, the complete content database must be backed up. If multiple site collections larger than 100 GB are contained in a single content database, backup and restore operations can take a long time and are at risk of failure. |

| | | | |
|--|--|---------------|---|
| | <p>Site collections exceeding quota limit:</p> <p>Indicates the number of site collections that are of a size that is greater than the configured quota template.</p> | <p>Number</p> | <p>A Quota Template allows Sharepoint administrators to specify the maximum amount of content that can be stored within a Site Collection. This way, administrators can exercise greater control on the amount of content that a site collection can store in the content database, which in turn, makes for better performance and a high quality user experience with Sharepoint.</p> <p>A non-zero value for this measure is indicative of the fact that one/more site collections are consuming more storage resources than they should. The detailed diagnosis of this measure will lead you to those errant site collections, so that you can figure out which sites on those collections are violating the set storage thresholds.</p> |
| | <p>Least active site collections:</p> <p>Indicates the number of site collections that are not frequently used.</p> | <p>Number</p> | <p>This measure reports the count of those sites that were not modified for a duration greater than the value of the LEAST ACTIVE SITE COLLECTION DAYS parameter. You can use the detailed diagnosis of this measure to know which site collections are seldom used.</p> <p>If the value of the Site collections measure appears to be rapidly approaching the maximum recommended site collection limit, then the detailed metrics will help you identify those site collections that are rarely used and are hence candidates for removal.</p> |
| | <p>Most active site collections:</p> <p>Indicates the number of site collections that were modified even yesterday.</p> | <p>Number</p> | <p>Use the detailed diagnosis of this measure to identify those site collections that are very actively used.</p> |
| | <p>Users in site collections:</p> <p>Indicates the number of users in site collections.</p> | <p>Number</p> | <p>Besides storage, quota templates can also restrict the number of users who can be added to the Active Directory directory service from a single site collection. When the maximum number of users for a site collection has been reached, no additional user accounts can be added unless one or more user accounts are deleted from the site collection. It is hence good practice to keep an eye on the changes to this measure, so as to proactively detect a potential user quota violation.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|---------------|---|
| | <p>Number of sites: Indicates the total number of sites in site collections.</p> | <p>Number</p> | <p>Microsoft recommends the creation of a maximum of 250,000 sites and subsites per site collection.</p> <p>You can create a very large total number of web sites by nesting subsites. For example, in a shallow hierarchy with 100 sites, each with 1,000 subsites, you would have a total of 100,000 web sites.</p> <p>Compare the value of this measure across site collections to know which collection consists of the maximum number of sites.</p> |
| | <p>Total size of sites: Indicates the total size of the sites in site collections.</p> | <p>MB</p> | <p>Typically, the value of this measure will be the same as that of the <i>Total size of site collections</i> measure.</p> <p>A site collection can be as large as the content database size limit for the applicable usage scenario.</p> <p>For more information about the different content database size limits for specific usage scenarios, see the Content database limits discussed in the Interpretation column of the <i>Content database size</i> measure of the Sharepoint Content Database test.</p> <p>In general, we strongly recommend limiting the size of site collections to 100 GB for the following reasons:</p> <ul style="list-style-type: none"> • Certain site collection actions, such as site collection backup/restore, cause large SQL Server operations which can affect performance or fail if other site collections are active in the same database. • SharePoint site collection backup and restore is only supported for a maximum site collection size of 100 GB. For larger site collections, the complete content database must be backed up. If multiple site collections larger than 100 GB are contained in a single content database, backup and restore operations can take a long time and are at risk of failure. |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|---|
| | <p>Most active sites:</p> <p>Indicates the number of sites that were accessed even yesterday.</p> | Number | Use the detailed diagnosis of this measure to identify those site collections that are very actively used. |
| | <p>Least active sites:</p> <p>Indicates the number of sites that are not used frequently.</p> | Number | <p>This measure reports the count of those sites that were not modified for a duration greater than the value of the LEAST ACTIVE SITE DAYS parameter. You can use the detailed diagnosis of this measure to know sites are seldom used.</p> <p>If the value of the <i>Number of sites</i> measure appears to be rapidly approaching the maximum recommended site limit, then the detailed metrics will help you identify those sites that are rarely used and are hence candidates for removal.</p> |

The detailed diagnosis of the *Least active site collections* measure reveals the top 10 site collections that were used the least. In times of rapid web application growth, this list will indicate those site collections that can be removed to curb the growth.

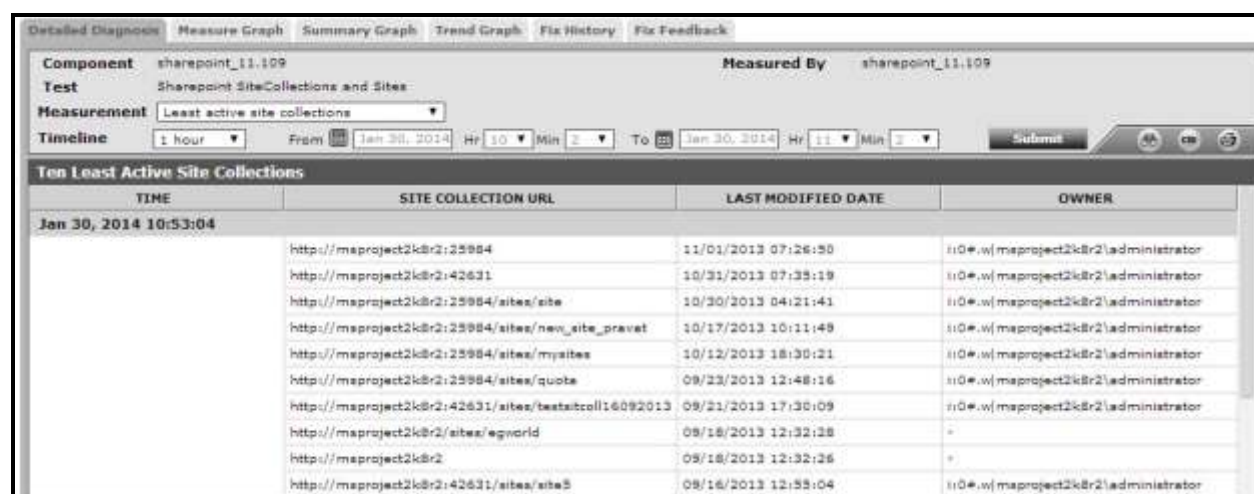


Figure 19.5: The detailed diagnosis of the Least active site collections measure

The detailed diagnosis of the *Least active sites* measure reveals the top 10 sites that were used the least. In times of rapid growth in the size of a site collection, this list will indicate those sites that can be removed to curb the growth.

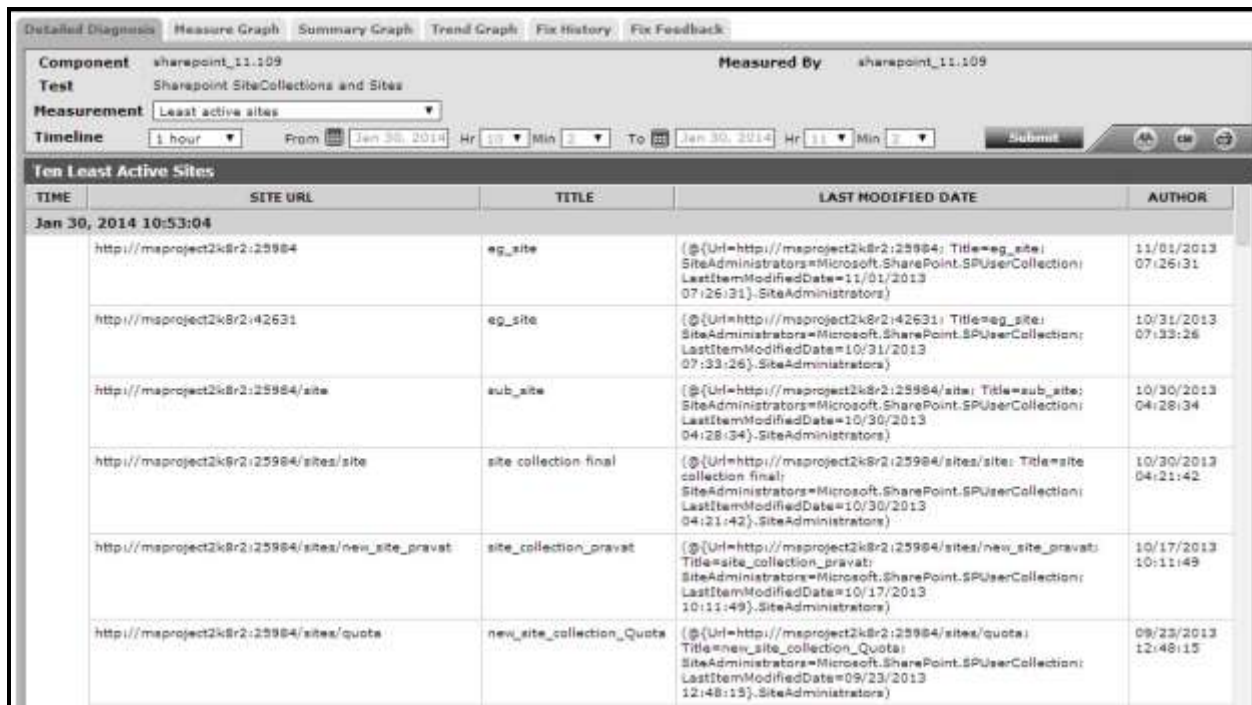


Figure 19.6: The detailed diagnosis of the Least active sites measure

18.2.3 Sharepoint Web Applications Test

Web Applications (WAs) are top-level containers for content in a SharePoint farm, and are typically the interface through which a user interacts with SharePoint - site collections, lists, and libraries come underneath the web application. A web application is associated with a set of access mappings or URLs which are defined in the SharePoint central management console, then automatically replicated into the IIS configuration of every server configured in the farm. WAs are typically independent of each other, have their own application pools, and can be restarted independently in Internet Information Services. Web Applications provide the ability to isolate content, processes, features and users. For example, you can separate the content anonymous users can see vs. what authenticated users can see by hosting the same content in different web apps.

A web application can grow in size over time! If this growth is not kept under control, then you may end up with a situation where a few web applications are hogging the storage resources provided by the Sharepoint environment, leaving the other web applications with limited to no resources! To avoid this, administrators need to be able to quickly isolate the web applications that are growing rapidly, understand their composition, and isolate the reasons for the abnormal growth. The **Sharepoint Web Applications** test helps administrators with this! For each web application deployed on a Sharepoint server, this test monitors the current size of that web application and captures a consistent increase in the size of the same, thus pointing you to those web applications that are growing in size at a steady pace and the content databases they are using. In addition, the test also leads you to the probable reasons for the abnormal size of the web application – is it because the web application is handling documents of huge sizes? or is it because the web application is storing too many versions of a document, which is in fact adding to its size? Or is it owing to the numerous sites, site collections, and document libraries that are being hosted by that web application?

| | |
|----------------|---|
| Purpose | For each web application deployed on a Sharepoint server, this test monitors the current size of that web application and captures a consistent increase in the size of the same, thus pointing you to those web applications that are growing in size at a steady pace and the content |
|----------------|---|

MONITORING MICROSOFT SHAREPOINT

| | | | |
|---|--|-------------------------|-----------------------|
| | databases they are using | | |
| Target of the test | A Sharepoint Server 2010/2013 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. DD FREQUENCY - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is <i>1:1</i>. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying <i>none</i> against DD FREQUENCY. 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 enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| Outputs of the test | One set of results for each web application on the Sharepoint Server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Size of this web application: Indicates the current size of this web application. | GB | |

| | | | |
|--|--|----------------|--|
| | <p>Web application growth rate:</p> <p>Indicates the percentage growth in the size of this web application since the last measurement period.</p> | <p>Percent</p> | <p>Compare the value of this measure across web applications to know which web application has grown the maximum since the previous measurement period.</p> <p>By closely tracking the variations in this measure for that web application over time, you can determine whether/not the web application is growing rapidly in size! If so, it is a cause for concern, as it indicates that that web application has the potential of consuming all available storage resources!</p> <p>In such a situation, you may want to reset the size limit for the site collections that are within the web application, so as to curb its growth.</p> <p>A site collection can be as large as the content database size limit for the applicable usage scenario.</p> <p>For more information about the different content database size limits for specific usage scenarios, see the Content database limits discussed in the Interpretation column of the <i>Cotent database size</i> measure of the Sharepoint Content Database test.</p> <p>In general, Microsoft strongly recommends limiting the size of site collections to 100 GB for the following reasons:</p> <ul style="list-style-type: none"> • Certain site collection actions, such as site collection backup/restore, cause large SQL Server operations which can affect performance or fail if other site collections are active in the same database. • SharePoint site collection backup and restore is only supported for a maximum site collection size of 100 GB. For larger site collections, the complete content database must be backed up. If multiple site collections larger than 100 GB are contained in a single content database, backup and restore operations can take a long time and are at risk of failure. |
|--|--|----------------|--|

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|--|
| | <p>Users in this web application:</p> <p>Indicates the number of users in this web application.</p> | Number | Compare the value of this measure across web applications to identify that application which has the maximum number of users. |
| | <p>Content databases used by this web application:</p> <p>Indicates the number of content databases that were used by this web application.</p> | Number | |
| | <p>Site collections part of this web application:</p> <p>Indicates the number of site collections in this web application.</p> | Number | <p>The maximum recommended number of site collections per farm is: Personal Sites - 500,000, Other site templates - 250,000. The Sites can all reside on one web application, or can be distributed across multiple web applications.</p> <p>Compare the value of this measure across web applications to know which application consists of the maximum number of site collections. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of site collections has anything to do with the increase in web application size.</p> |
| | <p>Sites part of this web application:</p> <p>Indicates the total number of sites in the site collections that are part of this web application.</p> | Number | <p>Microsoft recommends the creation of a maximum of 250,000 sites and subsites per site collection.</p> <p>You can create a very large total number of web sites by nesting subsites. For example, in a shallow hierarchy with 100 sites, each with 1,000 subsites, you would have a total of 100,000 web sites.</p> <p>Compare the value of this measure across web applications to know which application consists of the maximum number of sites. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of sites has anything to do with the increase in web application size.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|---------------|--|
| | <p>Number of document libraries:</p> <p>Indicates the number of document libraries in this web application.</p> | <p>Number</p> | <p>Document libraries are collections of files that you can share with team members on a Web based on Microsoft Windows SharePoint Services.</p> <p>By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of document libraries. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of document libraries has anything to do with the increase in web application size.</p> |
| | <p>Lists in this web application:</p> <p>Indicates the number of lists in this web application.</p> | <p>Number</p> | <p>A list in SharePoint is used to store data across columns in separate rows. By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of Sharepoint lists. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of lists has in any way impacted the web application size.</p> |
| | <p>Attachments:</p> <p>Indicates the number of attachments in this web application.</p> | <p>Number</p> | <p>By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of attachments. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of attachments has in any way impacted the web application size.</p> |
| | <p>Documents in this web application:</p> <p>Indicates the total number of documents in this web application.</p> | <p>Number</p> | <p>By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of documents. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of documents has in any way impacted the web application size.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|---|
| | Size of documents: Indicates the total size of all documents that are available in this web application. | GB | Compare the value of this measure across web applications to identify that application with the maximum document size. This can be attributed to the existence of one/more large-sized documents or many moderately sized documents in the web application. If that web application appears to be growing in size rapidly, you may want to keep an eye on this measure to figure out if it is owing to the increase in document size. |
| | Document versions: Indicates the number of document versions in this web application. | Number | Typically, Sharepoint can support a maximum of 40,000 major versions and 511 minor versions of documents. If this limit is exceeded basic file operations—such as file open or save, delete, and viewing the version history— may not succeed. |
| | Average number of documents per document library: Indicates the average number of documents per library in this web application. | Number | |

18.2.4 The Sharepoint Services Layer

The tests mapped to this layer shed light on the current status, overall health, and efficiency of the critical services offered by Sharepoint Foundation. This includes the Search archival and schema plugins, the search indexing mechanism, the search gatherer, and the critical Sharepoint Foundation processes.



Figure 18.6: The tests mapped to the Sharepoint Services layer

18.2.4.1 Sharepoint Search Archival Test

The Search feature of the Microsoft Sharepoint server not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites.

The **Archival** and **Schema** plugins are internal components of the Microsoft Sharepoint server Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the Microsoft Sharepoint server search feature is functioning, identify irregularities early, and fine-tune the Microsoft Sharepoint server to ensure peak performance of the search engine.

This test monitors the performance of the Sharepoint Foundation Search Archival Plugin.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the performance of the Sharepoint Foundation Search Archival Plugin | | |
| Target of the test | A Sharepoint Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results each for the Sharepoint Server | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Upload queues available to filtering threads: Indicates the number of upload queues that are available to filtering threads in this plugin. | Number | |
| | Queues committing changes and completing uploads: Indicates the number of queues that are exclusively allotted for committing the changes and completing the uploads. | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|--|
| | <p>Queues waiting to flush data to the property store:</p> <p>Indicates the number of queues that are waiting to flush data to the property store.</p> | Number | A property store is a table of properties and their values that are used and maintained by the Search service. Each row in the table corresponds to a document in the full-text index. |
| | <p>Queues being used by filtering threads:</p> <p>Indicates the number of queues that are being used by the filter threads in this plugin.</p> | Number | |
| | <p>Bulk insert sessions to the database server:</p> <p>Indicates the number of active bulk insert sessions to the database server.</p> | Number | |
| | <p>Documents processed:</p> <p>Indicates the number of documents that are processed in this plugin during the last measurement period.</p> | Number | A high value is desired for this measure. If the value decreases steadily over a period of time, it indicates a performance bottleneck. |
| | <p>Documents marked for retry by archival plugin:</p> <p>Indicates the number of documents that were marked for retry from this plugin during the last measurement period.</p> | Number | Ideally the value of this measure should be low. A higher value may indicate a performance bottleneck. |
| | <p>Documents waiting for a queue:</p> <p>Indicates the number of documents that were waiting for a queue during the last measurement period.</p> | Number | Ideally the value of this measure should be low. A higher value may indicate a performance bottleneck. |

18.2.4.2 Sharepoint Foundation Search Indexer Test

Using the Search feature of Sharepoint 2010, users can easily find the information they need in Sharepoint Foundation Sites.

The key components of Sharepoint's Search architecture are as follows:

- **Indexer:** Also referred to as the **Crawl Component** or the **Crawler**, the **Indexer** is solely responsible for

building indexes. The indexers enumerate the source content and pass text information to the relevant index partition on the query server. The indexer also indexes any metadata to the search property database and updates the crawl status in the crawl database.

- **Crawl Database:** The **Crawl Database** tracks what needs to be crawled and what has been crawled.
- **Query Component:** Commonly referred to as the **Query Server**, this component will perform a search against an index created by the indexer. The query component will apply such things as security trimming, best bets, relevancy, removes duplicates, etc.
- **Index partition:** Indexes can be split into multiple partitions called **index partitions** to improve the amount of time it takes to perform a search by the query component. For every query component there will be a single index partition that is queried by the query component.
- **Index Partition Mirror:** Mirrors can be created for the index partitions. These mirrors again provide the ability to provide redundancy and better search result performance.
- **Property Database:** This database stores metadata and security information items in the index. The property database will be associated with one or more query components and is used as part of the query process. These properties will be populated as part of the crawling process which creates the index.
- **Search Admin Database:** The **Search Administration Database** is mostly responsible for managing information associated to the configuration and topology of the Sharepoint Search service. There will only be one instance of this database for each Search Application Service instance.

Figure 18.7 depicts how these components work together to implement the search feature of Sharepoint 2010.

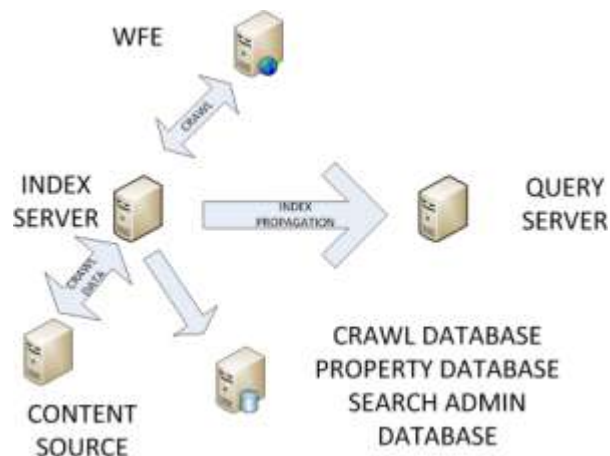


Figure 18.7: How Search works in Sharepoint 2010?

When a user enters a search query on a Web Front End (WFE) server, the query server processes the query. While processing, the query server retrieves the information that fulfills the query criteria from the index partition stored on its local file system, and also retrieves metadata information from the search property database. The index partition on the other hand, receives text information from the indexers that enumerate the source content. Once the desired query results are available, the query server packages the results, and delivers the results back to the requesting WFE server.

The success of Sharepoint Search feature therefore depends upon how quickly the query server processes the queries it receives, and how effective the index files built by the indexer are.

This test monitors the search queries to the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by Microsoft Sharepoint Server.

MONITORING MICROSOFT SHAREPOINT

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the search queries to the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by Microsoft Sharepoint Server | | |
| Target of the test | A Sharepoint Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 3. HOST - The host for which the test is to be configured 4. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active connections to the indexer plugin: Indicates the number of currently active connections to this indexer plugin. | Number | |
| | Index size: Indicates the current size of the content index that is being managed by this indexer plugin. | Number | |
| | Tasks in queue of propagation task sender: Indicates the number of tasks that were in queue of the propagation task sender. | Number | |
| | Tasks in queue of index receiver: Indicates the number of tasks that were in queue of the index receiver. | Number | |
| | Tasks in queue of index propagator: Indicates the number of tasks that were in queue of the index propagator. | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|--|
| | <p>Errors in Index propagation:</p> <p>Indicates the number of errors in index propagation during the last measurement period.</p> | Number | <p>Once the indexer builds the indexes, it propagates/pushes the index files from the index server to the query server. The indexer then waits for the query server to absorb the index, after which it acknowledges that the documents are successfully crawled.</p> <p>Ideally, no errors should occur in this process - i.e., the value of this measure should be ideally 0. The incidence of one or more errors can adversely impact the user experience with Sharepoint's Search mechanism.</p> |
| | <p>Errors in Index reception:</p> <p>Indicates the number of errors in index reception during the last measurement period.</p> | Number | <p>Ideally, no errors should occur in this process - i.e., the value of this measure should be ideally 0.</p> |
| | <p>Indexes received successfully:</p> <p>Indicates the number of indexes that were received successfully by this indexer plugin during the last measurement period.</p> | Number | <p>A high value is desired for this measure. A sudden/gradual decrease in the value of this measure may indicate a performance bottleneck in the Microsoft Server Search Indexer plugin.</p> |
| | <p>Indexes propagated successfully:</p> <p>Indicates the number of indexes that were propagated successfully by this indexer plugin during the last measurement period.</p> | Number | <p>A high value is desired for this measure. A sudden/gradual decrease in the value of this measure may indicate a performance bottleneck in the Microsoft Server Search Indexer plugin.</p> |
| | <p>Documents filtered:</p> <p>Indicates the number of documents that were filtered by this indexer plugin during the last measurement period.</p> | Number | |
| | <p>Documents in indexes that are being propagated:</p> <p>Indicates the number of documents in indexes which were being propagated by this indexer plugin during the last measurement period.</p> | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|--|
| | Queries handled: Indicates the number of queries that were handled on the content index during the last measurement period. | Number | |
| | Successful queries: Indicates the number of queries that were processed successfully during the last measurement period. | Number | A high value is desired for this measure. |
| | Failed Queries: Indicates the number of queries that failed to process during the last measurement period. | Number | Ideally, the value of this measure should be zero. |

| | | | |
|--|---|-------------|--|
| | <p>Avg latency of queries in the last minute:</p> <p>Indicates the average latency with which the queries were processed in the last minute.</p> | <p>Secs</p> | <p>Ideally, when an end user executes a query, results should be returned in less than one second. If this is not the case routinely, then end user experience with the Search feature is bound to suffer. The common reasons for poor query performance and their recommended solutions are as follows:</p> <ul style="list-style-type: none"> • One or more index partitions contain more than 10 million documents: Add an additional index partition, and if possible, an additional index partition mirror. If all query servers already contain an active and a mirrored index partition, add more query servers. • One or more query servers are memory bound and/or paging virtual memory on disk: Add additional memory to the query server. Ensure that the query server has enough RAM to store 33% of each index partition (present on the query server) in memory. • Query performance suffers during the first few queries after the server is rebooted or during crawl processing and index propagation: Ensure that the physical disk where the index partition is stored is capable of providing 2,000 IOPS for each index partition. • Query latency is high though all query servers are adequately sized: Ensure that the property database server has enough RAM available to store 33% of the property store tables in memory. Make sure that the property database server is not CPU or disk I/O bound. Additional property database servers or property databases can also be added based on need. |
|--|---|-------------|--|

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|------|--|
| | <p>Execution time to create a query restriction:</p> <p>Indicates the average execution time to create a query restriction.</p> | Secs | <p>Whenever query latency is very high - i.e., if the Avg latency of queries in the last minute measure reports a very high value - then, you can compare the values of these measures to understand where the query is spending too much time. You can thus identify the bottleneck areas and accordingly decide on the action to be taken to improve query performance.</p> |
| | <p>Execution time to resolve query:</p> <p>Indicates the average execution time to resolve a query.</p> | Secs | |
| | <p>Execution time to get row results of a query:</p> <p>Indicates the average execution time to get row results of a query.</p> | Secs | |
| | <p>Execution time spent in other parts of a query:</p> <p>Indicates the average time taken to create a query restriction.</p> | Secs | |
| | <p>CPU time to create a query restriction:</p> <p>Indicates the average CPU time that is required to create a query restriction.</p> | Secs | <p>If a query is found to be CPU-intensive, you can compare the values of these measures to determine where the query is consuming CPU excessively.</p> |
| | <p>CPU time to resolve a query:</p> <p>Indicates the average CPU time taken to resolve a query.</p> | Secs | |
| | <p>CPU time to get row results for a query:</p> <p>Indicates the average CPU time taken to get row results of a query.</p> | Secs | |
| | <p>CPU time spent in other parts of a query:</p> <p>Indicates the average CPU time taken to execute other parts of the query.</p> | Secs | |

18.2.4.3 Sharepoint Foundation Test

Microsoft Sharepoint Foundation is the essential solution for organizations that need a secure, manageable, web-based collaboration platform. It serves as the basis for Sharepoint server and offers out of the box elements such as

MONITORING MICROSOFT SHAREPOINT

blogs, wikis, team workspaces, and document libraries, providing users with the ideal way to share information and collaborate within a customized website. In addition, it provides services such as Business Data Connectivity services to integrate external data, basic search services and workflow services.

This test auto-discovers the Sharepoint processes, and for each process, reports the workload on the process and how efficiently that process handles the load. This way, the test leads you to the processes that are very busy and provides pointers to what could be keeping them busy.

| | | | |
|---|---|-------------------------|---|
| Purpose | Auto-discovers the Sharepoint processes, and for each process, reports the workload on the process and how efficiently that process handles the load | | |
| Target of the test | A Sharepoint Server 2010 | | |
| Agent deploying the test | An internal/remote agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for each Sharepoint Foundation process | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Active threads: Indicates the number of threads that are currently executing in Sharepoint code of this process. | Number | Many active threads is an indicator of a bottleneck. |
| | Incoming page requests: Indicates the number of incoming requests to access a particular page in the last second. | Number | This measure is a good indicator of the workload on this process. |
| | Requests being processed currently: Indicates the requests that are currently processed by this Sharepoint process. | Reqs | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|---|
| | <p>Avg execution time of requests processed:</p> <p>Indicates the average time taken by this process for executing the requests that are processed in the last second.</p> | Secs | Ideally, this value should be low. If the value of this measure increases steadily, then it indicates a performance bottleneck. |
| | <p>Requests rejected:</p> <p>Indicates the number of page requests that were rejected by this process during the last second.</p> | Number | Ideally, the value of this measure should be zero. |
| | <p>Requests responded to by the Sharepoint server:</p> <p>Indicates the number of page requests that were responded by this Sharepoint process during the last second.</p> | Number | |
| | <p>Throttled page requests:</p> <p>Indicates the number of page requests that have been throttled by this process during the last measurement period.</p> | Number | |
| | <p>Executing SQL queries:</p> <p>Indicates the number of SQL queries that are currently executing on this Sharepoint server.</p> | Number | |
| | <p>Query execution time:</p> <p>Indicates the average time taken by this Sharepoint server to execute the SQL queries.</p> | Secs | If the time taken to execute a query is <i>high</i> , it indicates that the query is unoptimal or there may be a database slowdown. |
| | <p>Native heaps in use:</p> <p>Indicates the number of native heaps that are currently in use by this Sharepoint process.</p> | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|--|
| | <p>Native heaps allocated by process:</p> <p>Indicates the number of native heaps that are allocated by this Sharepoint process.</p> | Number | |
| | <p>Global heap size:</p> <p>Indicates the size of the global heaps that are used by this Sharepoint process for cache related activity.</p> | MB | |
| | <p>Size of all per thread native heaps:</p> <p>Indicates the size of the native heaps that are used by all the threads that are being executed by this Sharepoint process.</p> | MB | |

18.2.4.4 Sharepoint Foundation Search Schema Test

The Search feature of the Microsoft Sharepoint server not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites.

The **Archival** and **Schema** plugins are internal components of the Microsoft Sharepoint server Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the Microsoft Sharepoint server search feature is functioning, identify irregularities early, and fine-tune the Microsoft Sharepoint server to ensure peak performance of the search engine.

This test monitors the performance of the Sharepoint Foundation Search Schema and Alias Mapping Plugin, and enables an informed assessment of its processing ability.

| | |
|---|---|
| Purpose | Monitors the performance of the Sharepoint Foundation Search Schema and Alias Mapping Plugin, and enables an informed assessment of its processing ability. |
| Target of the test | A Sharepoint Server 2010 |
| Agent deploying the test | An internal/remote agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. |
| Outputs of the test | One set of results each for the <i>ProfileImport</i> and <i>Portal_Content</i> instances |

MONITORING MICROSOFT SHAREPOINT

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|---|------------------|----------------|
| | <p>Documents processed by schema plugin:</p> <p>Indicates the number of documents that are processed by this schema plugin during the last measurement period.</p> | Number | |
| | <p>Properties processed by schema plugin:</p> <p>Indicates the number of properties that are processed by this schema plugin during the last measurement period.</p> | Number | |
| | <p>Aliases loaded:</p> <p>Indicates the number of aliases that have been currently loaded to this schema plugin.</p> | Number | |
| | <p>Aliases have been mapped:</p> <p>Indicates the total number of aliases that have been currently mapped to this schema plugin during the last measurement period.</p> | Number | |
| | <p>Aliases ignored as they are duplicates:</p> <p>Indicates the number of aliases that the schema currently ignored as they are duplicates during the last measurement period.</p> | Number | |
| | <p>Aliases refreshed from the database:</p> <p>Indicates the number of aliases that have been refreshed from the database during the last measurement period.</p> | Number | |

18.2.4.5 Sharepoint Foundation Search Gatherer Test

The search functionality can be described in its simplest form as a Web page where the user defines his or her search query. The index role can be configured to run on its own Microsoft Sharepoint server, or run together with all the other roles, such as the Web service, Excel Services and Forms Services. It performs its indexing tasks following this general workflow:

1. Sharepoint stores all configuration settings for the indexing in its database.
2. When activated, the index will look in Sharepoint's databases to see what content sources to index, and what type of indexing to perform, such as a full or incremental indexing.
3. The index service will start a program called the Gatherer, which is a program that will try to open the content that should be indexed.
4. For each information type, the Gatherer will need an Index Filter, or IFilter, that knows how to read text inside this particular type of information. For example, to read a MS Word file, an IFilter for .DOC is needed.
5. The Gatherer will receive a stream of Unicode characters from the IFilter. It will now use a small program called a Word Breaker; its job is to convert the stream of Unicode characters into words.
6. However, some words are not interesting to store in the index, such as "the", "a", and numbers; the Gatherer will now compare each word found against a list of Noise Words. This is a text file that contains all words that will be removed from the stream of words.
7. The remaining words are stored in an index file, together with a link to the source. If that word already exists, only the source will be added, so one word can point to multiple sources.
8. If the source was information stored in Sharepoint, or a file in the file system, the index will also store the security settings for this source. This will prevent a user from getting search results that he or she is not allowed to open.

Since the success of an indexing operation also depends upon how the Gatherer program functions, administrators need to keep their eyes open for irregularities in the functioning of the gatherer, so that such anomalies are detected instantly, and corrected before they can stall the indexing process.

This test monitors the performance of the Sharepoint Foundation Search Gatherer, and reports issues in its performance (if any).

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the performance of the Sharepoint Foundation Search Gatherer, and reports issues in its performance (if any) | | |
| Target of the test | A Sharepoint Server 2010 | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results each for the <i>ProfileImport</i> and <i>Portal_Content</i> instances | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

| | | | |
|--|--|--------|--|
| | <p>Filtering threads in the system:</p> <p>Indicates the current number of filtering threads in the system.</p> | Number | |
| | <p>Threads waiting for documents:</p> <p>Indicates the number of threads that are currently waiting for documents.</p> | Number | <p>These threads are not currently doing any work and will eventually be terminated. If you consistently have a more than <i>Max Threads/Hosts</i> idle threads you can schedule an additional crawl. If this number is 0 then you are starved. Do not schedule another crawl in this time period and analyze the durations of your crawls during this time to see if they are meeting your freshness goals. If your goals are not being met you should reduce the number of crawls.</p> |
| | <p>Threads waiting for network response from the filter process:</p> <p>Indicates the number of threads that were waiting for a response from the filter process.</p> | Number | <p>If you figure out that there is no activity that is taking place as far as this measure is concerned, and if the value of this measure is equal to the <i>Filtering threads in system</i> measure, it indicates a network issue or the unavailability of the server that is crawling into.</p> |
| | <p>Threads committing transactions:</p> <p>Indicates the number of threads that are committing transactions.</p> | Number | |
| | <p>Threads waiting for plug-ins to complete an operation:</p> <p>Indicates the number of threads currently waiting for plug-ins to complete an operation.</p> | Number | <p>These threads have the filtered documents and are processing it in one of several plug-ins. This is when the index and property store are created.</p> |
| | <p>Threads loading transactions from persisted crawl queue:</p> <p>Indicates the number of transactions that are loaded from the persisted crawl queue.</p> | Number | |
| | <p>Threads processing links:</p> <p>Indicates the number of threads that are processing links.</p> | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------|---|
| | <p>Filtering processes in the system:</p> <p>Indicates the number of filtering processes that are active in the system.</p> | Number | |
| | <p>Filter objects in the system:</p> <p>Indicates the number of filter objects in the system.</p> | Number | |
| | <p>Documents waiting for robot threads:</p> <p>Indicates the number of documents that are waiting for robot threads.</p> | Number | If the value of this measure is 0, then it implies that all the threads are filtering threads. |
| | <p>Currently connected admin clients:</p> <p>Indicates the number of currently connected admin clients.</p> | Number | |
| | <p>Amount of resources allowed for the Gatherer service:</p> <p>Indicates the amount of resources that the Gatherer service is allowed to use.</p> | Number | |
| | <p>Servers recently accessed by the system:</p> <p>Indicates the number of servers that were recently accessed by the system.</p> | Number | |
| | <p>Servers currently unavailable:</p> <p>Indicates the number of servers that are currently unavailable to the system.</p> | Number | A server becomes unavailable if the requests made to the server is timed out. |
| | <p>Available cached stemmer instances:</p> <p>Indicates the number of cached stemmer instances in the system.</p> | Number | Stemmers are nothing but components shared by the Search and Indexing engines that generate inflected forms for a word. Too many stemmer instances that are cached may indicate a resource usage problem. |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|---|
| | <p>System I/O rate:</p> <p>Indicates the rate at which the system IO disk traffic is detected during back off period.</p> | KB/Sec | <p>During a back-off period, indexing is suspended. To manually back off the gatherer service, pause the search service. If the search service itself generates the back-off, an event will be recorded and the search service will be paused automatically. There is no automatic restart, so you must manually start the search service in order to end a back-off state. Note that there is little reason to start the search service until you have solved the problem that caused the back-off in the first place.</p> |
| | <p>Timeouts:</p> <p>Indicates the number of timeouts detected by the system during the last measurement period.</p> | Number | <p>Ideally, this value should be zero.</p> |
| | <p>Documents filtered:</p> <p>Indicates the rate at which the documents are filtered in the system.</p> | KB/Sec | <p>If this rate is decreasing over time, you should perform some troubleshooting to find out why your server is not filtering documents.</p> <p>Look for memory issues, processor issues, network issues, or site hit frequency rules that slow the gatherer process.</p> |
| | <p>Documents successfully filtered:</p> <p>Indicates the rate at which the documents are filtered successfully in the system.</p> | KB/Sec | |
| | <p>Documents delayed due to site hit frequency rules:</p> <p>Indicates the number of documents that were currently delayed due to site hit frequency rules.</p> | Number | <p>If you have a plethora of rules and this number is steadily increasing over time, consider relaxing or simplifying your site hit frequency rules. A very high number may indicate a conflict in the rules that the gatherer cannot resolve or follow with efficiency.</p> |
| | <p>Document entries currently in memory:</p> <p>Indicates the number of document entries that are currently available in the memory of the system.</p> | Number | |
| | <p>Documents filtered:</p> <p>Indicates the total number of documents filtered in the system during the last measurement period.</p> | Number | |

| | | | |
|--|--|--------|---|
| | <p>Documents successfully filtered:</p> <p>Indicates the total number of documents that are successfully filtered in the system during the last measurement period.</p> | Number | If the value of this measure is less than the value of the <i>Documents filtered</i> measure, use the gatherer logs to figure out the cause for the documents that are attempting to be filtered but are failing. |
|--|--|--------|---|

18.2.4.6 Distributed Cache Service Test

SharePoint uses the Distributed Cache to store data for very fast retrieval across all entities. The Distributed Cache service provides in-memory caching services to several features in SharePoint Server 2013. Some of the features that use the Distributed Cache service include:

- Newsfeeds
- Authentication
- pOneNote client access
- Security Trimming
- Page load performance

Besides services, several caches that exist in Sharepoint 2013 depend upon the Distributed Cache service for their proper functioning.

Any server in the farm running the Distributed Cache service is known as a **cache host**. A **cache cluster** is the group of all cache hosts in a SharePoint Server 2013 farm. A cache host joins a cache cluster when a new application server running the Distributed Cache service is added to the farm. When using a cache cluster, the Distributed Cache spans all application servers and creates one cache in the server farm. The total cache size is the sum of the memory allocated to the Distributed Cache service on each of the cache hosts.

If the distributed cache is not able to service requests efficiently, it is bound to significantly impact the performance of the dependent services/caches. Furthermore, it will add significantly to the processing overheads of Sharepoint, as poor cache usage translates into increased database accesses. If this is to be prevented, administrators should keep a close watch on the distributed cache’s ability to service requests, rapidly detect poor cache usage patterns, and accurately pinpoint the reason for the same – is it because adequate objects are not cached in the distributed cache? If so, why? Is it owing to insufficient cache size? Will allocating more memory to the cache help or should more servers be added to the cache cluster? The **Distributed Cache Service** test helps answer all these questions! This test continuously monitors the requests to the cache, reports the count of requests serviced and rejected by the cache, and thus enables administrators to ascertain how well the cache is utilized. In the event of poor cache usage, close scrutiny of these test results will provide administrators with useful pointers to what is impeding cache usage and whether/not right-sizing the cache will help clear the bottleneck.

| | |
|---------------------------------|--|
| Purpose | Continuously monitors the requests to the cache, reports the count of requests serviced and rejected by the cache, and thus enables administrators to ascertain how well the cache is utilized |
| Target of the test | A Sharepoint Server 2013 |
| Agent deploying the test | An internal agent |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|---|---|-------------------------|---|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results each for the Sharepoint server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Cache data transferred rate:</p> <p>Indicates the number of cached entries transferred per second.</p> | Number | |
| | <p>Cache hit count:</p> <p>Indicates the number of requests serviced by the cache during the last measurement period.</p> | Number | <p>A high value is desired for this measure. A sudden/steady dip in this value indicates that the cache is unable to process requests, thereby increasing direct database accesses.</p> |

| | | | |
|--|---|----------------|---|
| | <p>Cache hit ratio:</p> <p>Indicates the percentage of requests that were serviced by the cache.</p> | <p>Percent</p> | <p>A high value is desired for this measure. A sudden/steady drop in this value is indicative of poor cache usage, which in turn can cause direct database accesses to increase and strain the database.</p> <p>One of the common reasons for a low cache hit ratio is insufficient memory allocation to the cache. In the absence of adequate memory resources, the cache may not be able to hold many frequently-accessed objects within, and may hence not be able to service many requests. Under such circumstances, you may want to consider allocating more memory to the cache. Here are a few recommendations from Microsoft with regard to how to size the distributed cache:</p> <ul style="list-style-type: none"> • The Distributed Cache service actually uses twice the allocated amount of RAM, using the extra for housekeeping. In a small farm with fewer than 10,000 users, Microsoft recommends allocating 1GB of RAM for the Distributed Cache. This can be either a dedicated server or collocated with other SharePoint services, such as the Web Application Service. Beyond this the recommendation is using dedicated servers for the cache. A medium farm with fewer than 100,000 users should look to allocate around 2.5GB for the cache, and a large farm with up to 500,000 users should set aside around 12GB of RAM allocated for the cache. • It is a very strong recommendation that you should not allocate more than 16GB to any one Cache Host. This may cause the Cache Service to timeout during housekeeping operations and become unresponsive for several seconds at a time. If you need a cache size of greater than 16GB, it is better to use multiple servers in a Cache Cluster. You can have up to a maximum of 16 hosts in a Cache Cluster. |
|--|---|----------------|---|

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|--|
| | <p>Cache miss count:</p> <p>Indicates the number of requests that were not serviced by the cache since the last measurement period.</p> | Number | <p>Ideally, the value of this measure should be low. A sudden/steady increase in this value is indicative of poor cache usage, which in turn can cause direct database accesses to increase and strain the database.</p> |
| | <p>Cache read requests rate:</p> <p>Indicates the number of read requests to the cache per second, during the last measurement period.</p> | Number | <p>A high value for these measures is often indicative of heavy load on the distributed cache.</p> <p>In such a situation, for better cache performance, it is recommended that you opt for the dedicated mode of cache deployment. In this mode, all services other than the Distributed Cache service are stopped on the application server that runs the Distributed Cache service, thus ensuring that all critical resources on the server are at the disposal of the distributed cache. This in turn, will help the cache handle the load efficiently!</p> |
| | <p>Cache write requests rate:</p> <p>Indicates the number of write requests to the cache per second, during the last measurement period.</p> | Number | |
| | <p>Total cache read requests:</p> <p>Indicates the total number of read requests received by the cache since the last measurement period.</p> | Number | |
| | <p>Total cache write requests:</p> <p>Indicates the total number of write requests received by the cache since the last measurement period.</p> | Number | |
| | | | |

18.2.5 Sharepoint Search Content Feed Layer

The key components of the Sharepoint content feeding chain are:

- Crawl Database
- Crawl Component
- Content Processing Component
- Index Component

When search queries execute slowly, administrators need to figure out where in the feeding chain the slowdown originated. The tests mapped to this layer run checks on all the aforesaid components, so that administrators can accurately isolate the probable cause of this slowdown.

18.2.5.1 Search Gatherer Threads Test

Search in SharePoint 2013 enables users to find relevant information more quickly and easily than ever before and makes it easy for Search administrators to customize the search experience.

The search architecture consists of the following areas:

- Crawl and content processing
- Index
- Query processing
- Search administration
- Analytics

Figure 19.1 depicts how these components work together to implement the search functionality in Sharepoint 2013.

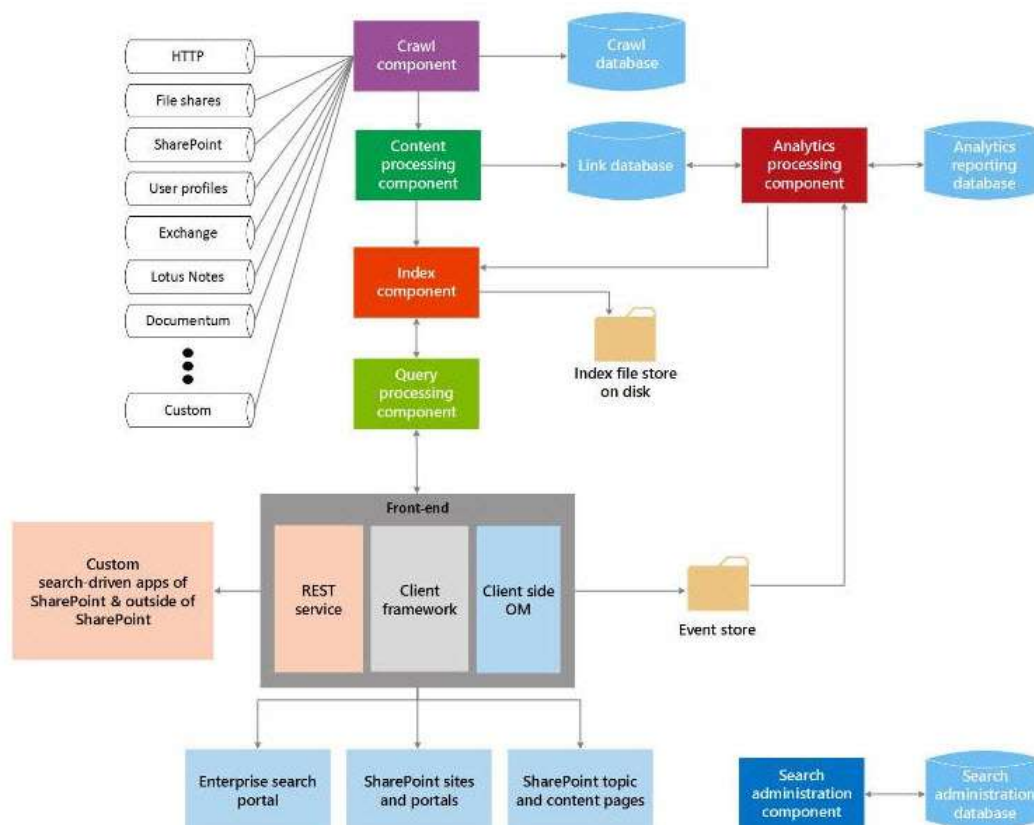


Figure 19.7: How search works in Sharepoint 2013?

From Figure 19.7, it is clear that the crawl component lays the foundation for the search mechanism! The crawl component crawls content sources to collect crawled properties and metadata from crawled items and sends this information to the content processing component. This means that if the crawl component is unable to crawl the content hosts, it could impact the speed of every dependent operation – be it content processing, indexing, query processing etc. – thereby crippling the entire search engine! Hence, for search in Sharepoint 2013 to be quick and

MONITORING MICROSOFT SHAREPOINT

efficient, administrators should primarily keep an eye on the crawl component, swiftly isolate painpoints in crawling, and clear them rapidly. To achieve this, administrators can use the **Search Gatherer Threads** test. This test monitors the crawling process and reveals how well the crawling worker threads are doing their jobs. While at it, the test proactively notifies administrators of a potential slowdown (if any) in crawling and pinpoints what is causing the slowdown – a hungry content host? or improperly configured crawls? .

| | | | |
|---|--|-------------------------|---|
| Purpose | Monitors the crawling process and reveals how well the crawling worker threads are doing their jobs. While at it, the test proactively notifies administrators of a potential slowdown (if any) in crawling and pinpoints what is causing the slowdown – a hungry content host? or improperly configured crawls? | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Threads accessing the network: Indicates the number of threads that are waiting on the content host to return the requested content. | Number | If this number is consistently high then you are either network bound or you are bound by a "hungry" host. If you are <u>not</u> meeting your crawl freshness goals, you can either change your crawl schedules to minimize overlapping crawls or look at the remote repositories you are crawling to optimize them for more throughput. |
| | Filtering threads: Indicates the current number of filtering threads in the system. | Number | If the value of the <i>Threads accessing the network</i> measure is close to that of the <i>Filtering threads</i> measure, it is an indication that a bottleneck exists at the content source/host. When this happens, you may also want to check whether processor usage on the crawl component servers is low. Likewise, look for disk latency issues on the crawl database. If all the above exist, it is a clear indicator that the content host/source is where the bottleneck lies! |

| | | | |
|--|--|--------|--|
| | <p>Idle threads:</p> <p>Indicates the number of threads that are currently waiting for documents.</p> | Number | <p>These threads are not currently doing any work and will eventually be terminated. If you consistently have a more than <i>Max Threads/Hosts</i> idle threads you can schedule an additional crawl. If this number is 0 then you are starved. Do not schedule another crawl in this time period and analyze the durations of your crawls during this time to see if they are meeting your freshness goals. If your goals are not being met you should reduce the number of crawls.</p> |
|--|--|--------|--|

18.2.5.2 Search Gatherer Transactions Test

Crawls, when scheduled to occur too frequently, can significantly impact the processing ability of the content processing component, the level of I/O activity on the crawl database, and ultimately, the search throughput! Likewise, a resource-starved content processing component and/or a crawl database can also considerably slowdown Sharepoint search, as they may not be able to handle the workload generated by the crawler! This is why, when end-users complain of slow searching by Sharepoint, administrators need to be able to quickly figure out where the bottleneck is and how to clear it – should the crawl schedules be changed so that less crawls occur? Or should the processing power of the content processor and crawl database change in tandem with the frequency of crawls? This is where the **Search Gatherer Transactions** test helps!

This test monitors the transactions on the crawl component and reports the count of transactions that are waiting for processing by the content processor and those that have completed processing. In the process, the test turns the spotlight on a potential processing slowdown and accurately pinpoints what is causing it – is it owing to too many crawls? Or is it because the content processor and/or the crawl database are incorrectly sized? Based on the results of this test, administrators can clearly understand what needs to be fine-tuned and how.

| | |
|---|--|
| Purpose | Monitors the transactions on the crawl component and reports the count of transactions that are waiting for processing by the content processor and those that have completed processing. In the process, the test turns the spotlight on a potential processing slowdown and accurately pinpoints what is causing it – is it owing to too many crawls? Or is it because the content processor and/or the crawl database are incorrectly sized? Based on the results of this test, administrators can clearly understand what needs to be fine-tuned and how |
| Target of the test | A Sharepoint Server |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. |
| Outputs of the test | One set of results for the Sharepoint server monitored |

MONITORING MICROSOFT SHAREPOINT

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|-------------------------------|---|------------------|---|
| | <p>Waiting transactions:</p> <p>Indicates the number of transactions that are currently waiting to be processed by the content processing component.</p> | Number | <p>Ideally, this value should be low (less than a few thousand). If so, it implies that content processing is keeping up with content crawling.</p> <p>On the other hand, if the value of this measure is high and/or consistently rising, then it means that the crawl component is pushing more data for processing than what the content processing component can handle. This will slow down content processing and eventually affect Sharepoint search! Under such circumstances, you can do either of the following:</p> <ul style="list-style-type: none"> • Provide more processing power to the content processing component, so that it is able to handle the load imposed by the crawl component. You can also add more content processing components to uniformly distribute the processing load. • Reconfigure the crawl component to run crawls less frequently, so that the crawl component does not overload the content processing component |
| | <p>Transactions in progress:</p> <p>Indicates the number of transactions that are currently being processed by the crawl component.</p> | Number | This is a good indicator of the current load on the crawl component. |
| | <p>Completed transactions:</p> <p>Indicates the number of transactions that are completed</p> | Number | <p>If this value is very high (say, greater than a few hundred), it means that too many transactions are getting completed and are written to the crawl database, causing disk activity on the database to increase. At this juncture, check the crawl database for disk latency. If the disk latency and disk queue length are also high, you can conclude that the crawl database is where the bottleneck is.</p> |

18.2.5.3 Search Submission Test

Like problems in the content acquisition process, snags in the content processing routine can also delay searching. Content processing in Sharepoint is performed by the content processing component (CPP) and the index component. Once crawling is complete, the Content plug-in on the crawl component first routes the content to the **Content Submission Service** (CSS) of the content processing component. An instance of the CSS runs alongside each instance of a content processing component. Once the content plug-in on the crawl component establishes a session with the CSS, the CSS load-balances the incoming content by uniformly distributing the content to the content processing components (CPC). Upon receipt of documents from the CSS, the content processing component processes the documents and then sends them to the indexer for indexing.

If a crawler session is unexpectedly terminated by CSS, then some crawled content may not even reach the CSS, and will hence not be processed or indexed; this will eventually impact the search service! Moreover, if CSS is not able to push its document load to the content processing component fast enough, documents may get timed out from the CSS itself, and will hence be omitted from the search index; this again will result in a poor search experience. Likewise, if the content processing component suffers a slowdown, document processing and indexing will be significantly delayed, which in turn can affect querying. If such problems are to be avoided, administrators should closely monitor the availability and processing ability of the CSS and the CPC, and rapidly isolate bottlenecks. This is where the **Search Submission** test helps. This test periodically checks the sessions to CSS, monitors how quickly the CSS load-balances the content and transmits it to the CPC, and measures the processing capacity of the CPC. When users complain of their search queries being slow, then this test will shed light on the probable cause of the delay – is it owing to sudden/sporadic breaks in the crawler sessions to CSS? Is it because of a load-balancing bottleneck experienced by the CSS? Or is it due to a processing slowdown at the CPC? Based on the findings reported by this test, administrators can initiate the appropriate remedial measures.

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Periodically checks the sessions to CSS, monitors how quickly the CSS load-balances the content and transmits it to the CPC, and measures the processing capacity of the CPC | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|---|
| | <p>Aborted sessions:</p> <p>Indicates the number of sessions that aborted since the start of the component.</p> | Number | <p>Ideally, the value of this measure should be 0. A high value is a cause for concern as it indicates frequent breaks in the crawler sessions on the CSS. Too many broken sessions can seriously impede the transfer of crawled content from the crawler to the CSS, resulting in incomplete transfers! This warrants an investigation into the reason for the frequent session failures.</p> |
| | <p>Active sessions:</p> <p>Indicates the number of crawler sessions that are currently active on the CSS.</p> | Number | <p>This is a good indicator of the current load on the CSS.</p> |
| | <p>Available callbacks:</p> <p>Indicates the current number of callbacks ready for consumption, but not yet consumed by the client.</p> | Number | <p>Once the content processing component processes the content it receives and writes it to the index, it sends out a 'call back' to the content plug-in on the crawler indicating the processing status of that content.</p> <p>A high value for this measure indicates that while the CPC has been able to generate callbacks, many of these callbacks have not yet been consumed by – i.e., have not yet reached – the crawler. This hints at an error in network communication between the crawler and the CPC.</p> |
| | <p>Total callbacks:</p> <p>Indicates the total number of callbacks produced by the submission service since the start of the component.</p> | Number | <p>You may want to compare the value of the <i>Available callbacks</i> measure with that of this measure to understand what fraction of callbacks is still to be consumed by the crawl component.</p> |
| | <p>Client polls:</p> <p>Indicates the total number of client polls since the start of the component.</p> | Number | <p>Each time a client refreshes the session to check for callbacks this measure will be incremented.</p> |
| | <p>Client submits:</p> <p>Indicates the total number of submits performed by clients since the start of the component.</p> | Number | |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|--|--------|--|
| | Skipped documents: Indicates the total number of documents skipped in the submission service before being delivered to the content processing component. | Number | A non-zero value is desired for this measure. A high value is disconcerting as it indicates that too many crawled documents are not reaching the CPC for processing as the CSS disregards them. Further investigation into the reasons is necessitated. |
| | Timed out documents: Indicates the total number of documents that timed out in the submission service. | Number | A low value is desired for this measure. A high value implies that the search index may not include many crawled documents as they have been timed out of the submission queue itself. This in turn may result in ineffective search queries. You may hence want to reset the timeout value for documents in the submission service. |
| | Flows used for feeding: Indicates the current number of flows used for feeding. | Number | The CPC uses Flows and Operators to process the content. Flows define how to process content, queries and results and each flow processes one item at a time. The number of current flows is hence an indicator of the number of documents that are being processed by the CPC. |
| | Pending items: Indicates the current number of items delivered to the content processing component but where no callback has yet been received. | Number | A high value or a consistent rise in the value for this measure could indicate a bottleneck in content processing. |

18.2.5.4 Search Flow Test

Content processing in Sharepoint is performed by the content processing component (CPC) and the index component.

The Content Processing Component (CPC) uses Flows and Operators to process the content (see Figure 19.8). Flows define how to process content, queries and results and each flow processes one item at a time. Flows consist of operators and connections organized as graphs. This is where activities like language detection, word breaking, security descriptors, content enrichment (web service callout), entity and metadata extraction, deep link extraction and many others take place. The flow has branches that handle different operations, like inserts, deletes and partial updates.

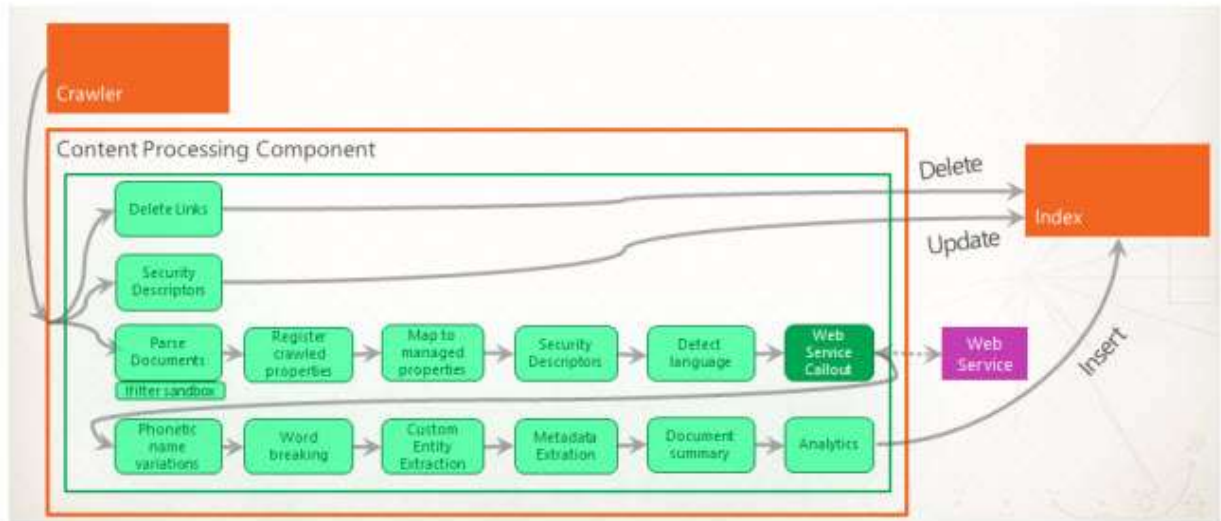


Figure 19.8: Flows and operators in CPC

Once content is processed by the CPC, the index component receives the processed items from the CPC and writes them to the search index. The index component also handles incoming queries, retrieves information from the search index, and sends back the result set to the query processing component.

Whether it is the CPC that fails to process the content rapidly or the index component that writes to the index slowly, what suffers is the end-user’s experience with Sharepoint search! To ensure that Sharepoint delivers to users a fast and flawless searching experience, administrators should not only be able to detect slowdowns before they impact query processing, but also tell where the slowdown originated – is it with the CPC or the index component? The **Search Flows** test answers this question accurately! This test monitors the flows on CPC, keeps track of documents that are in queue waiting to be processed by the flows, and reports how quickly the CPC and the index component process the enqueued contents. While at it, the test points to potential bottlenecks in content processing and accurately isolates the source of the bottleneck – is it the CPC or the index component?

| | | | |
|---|---|-------------------------|-----------------------|
| Purpose | Monitors the flows on CPC, keeps track of documents that are in queue waiting to be processed by the flows, and reports how quickly the CPC and the index component process the enqueued contents. While at it, the test points to potential bottlenecks in content processing and accurately isolates the source of the bottleneck – is it the CPC or the index component? | | |
| Target of the test | A Sharepoint Server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> TEST PERIOD - How often should the test be executed HOST - The host for which the test is to be configured PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the Sharepoint server monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|-----------|--|
| | <p>Total inbound items:</p> <p>Indicates the total number of items placed on input queues.</p> | Number | |
| | <p>Items queued for processing:</p> <p>Indicates the number of items that are currently in queues in front of input operators that are ready for processing.</p> | Number | A high value or a consistent increase in the value of this measure is indicative of bottlenecks in content processing. |
| | <p>Active threads:</p> <p>Indicates the number of threads that are currently active.</p> | Number | |
| | <p>Input queue empty time:</p> <p>Indicates the total time spent by input operators waiting for items.</p> | Millisecs | <p>If this value is low (say, less than a thousand), it indicates that the input queues are rarely ever empty! You may then want to check the processor usage on the CPC component. If this is very high, it is a clear indication that the CPC is stressed and could be the key contributor to the slowdown in content processing.</p> <p>On the other hand, if the value of this measure is high (say, over a thousand) , it indicates that the input queues are empty for long time spells. This implies that the CPC is processing content quickly. In this case, check the disk I/O and latency on the index component. If these parameters are high, it implies that the index component is stressed and is unable to handle the load imposed by the CPC. You can thus conclude that the bottleneck lies with the index component.</p> |

MONITORING MICROSOFT SHAREPOINT

| | | | |
|--|---|--------------|---|
| | <p>Input queue full time:</p> <p>Indicates the total time spent waiting for space to become available on input queues.</p> | Milliseconds | <p>If this value is high (say, over a thousand), it indicates that the CPC is taking a long time to process the contents in the input queues and free up the queues! You may then want to check the processor usage on the CPC component. If this is very high, it is a clear indication that the CPC is stressed and could be the key contributor to the slowdown in content processing.</p> <p>On the other hand, if the value of this measure is low (say, less than a thousand), it indicates that the input queues are getting cleared very quickly. This implies that the CPC is processing content quickly. In this case, check the disk I/O and latency on the index component. If these parameters are high, it implies that the index component is stressed and is unable to handle the load imposed by the CPC. You can thus conclude that the bottleneck lies with the index component.</p> |
|--|---|--------------|---|

Monitoring Microsoft Dynamics AX

Microsoft Dynamics® AX is an integrated, adaptable business management solution that streamlines financial, customer relationship, and supply chain processes. This ERP solution consolidates and standardizes processes, provides visibility across your organization, and simplifies compliance.

Since decision-makers rely on this solution for working efficiently and taking prompt and accurate decisions, slowdowns experienced by the solution and exceptions thrown by the AX portal can greatly impair the productivity and the decision-making ability of the users, and can ultimately affect revenues.

To avert this, the AX Application Object Server (AOS) and the AX portal need to be continuously monitored, and users promptly alerted to processing delays, overloads, and errors.

eG Enterprise provides a dedicated *Microsoft Dynamics AX* monitoring model that proactively detects and promptly alerts users to issues in the performance of the Dynamics AX solution.

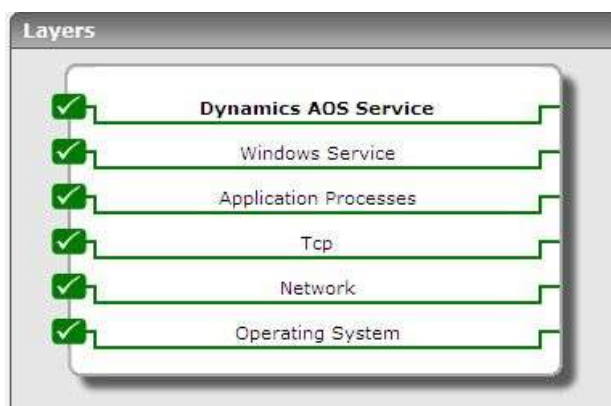


Figure 19.9: The layer model of the Microsoft Dynamics AX solution

Each layer in Figure 19.9 is mapped to a set of tests, which employ agent-based or agentless techniques to extract critical performance statistics from the AX solution. These metrics provide answers to the following key question:

- Is the AX server overloaded with requests?

MONITORING MICROSOFT DYNAMICS AX

- Is the server able to process the requests quickly?
- Has the AX Enterprise portal encountered any .NET business connector exceptions? If so, how many, and of what type?

Since the last 5 layers of Figure 19.9 have been discussed in-depth in the *Monitoring Unix and Windows Servers* document, this chapter will be discussing the top layer alone.

19.1 Dynamics AOS Service

The tests mapped to this layer monitors the load and the processing ability of the Application Object Server (AOS), and also captures exceptions (if any) that are encountered by the AX Enterprise Portal.

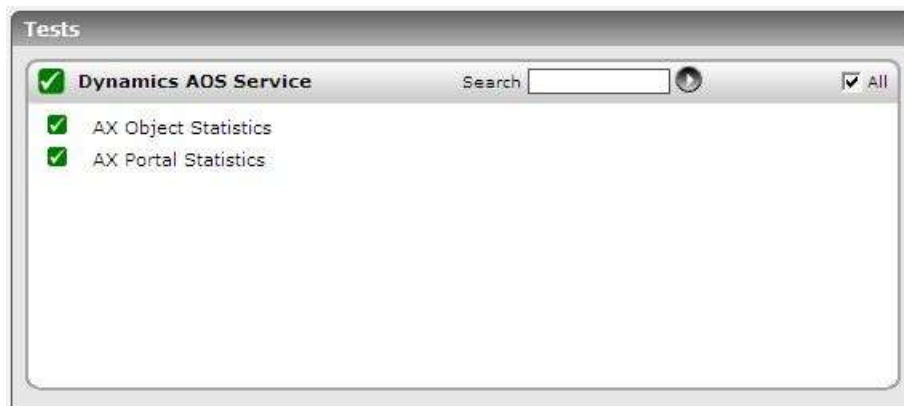


Figure 19.10: The tests mapped to the Dynamics AOS Service

19.1.1 AX Object Statistics Test

This test reports useful statistics with the help of which the session, request, and data load on the Application Object Server (AOS) can be ascertained.

| | |
|---------------------------------|--|
| Purpose | Reports useful statistics with the help of which the session, request, and data load on the Application Object Server (AOS) can be ascertained |
| Target of the test | A Microsoft Dynamics AX server |
| Agent deploying the test | An internal agent |

| | | | |
|---|---|-------------------------|--|
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the server being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | Total sessions to AOS: Indicates the total number of active sessions on the server during the last measurement period. | Number | This is a good indicator of the session load on the server. |
| | Currently active sessions to AOS: Indicates the number of currently active server sessions. | Number | |
| | Client-to-server requests handled: Indicates the number of client-to-server requests during the last measurement period. | Number | This measure is a good indicator of the workload on the server. |
| | Client-to-server processing rate: The number of client-to-server requests processed per second | Reqs/Se | A low rate could indicate a processing bottleneck. |
| | Server-to-client requests processed: Indicates the number of server-to-client requests processed during the last measurement period. | Number | |
| | Data transmitted by server: Indicates the number of bytes sent by the server during the last measurement period. | Number | These measures are good indicators of the data load on the server. |

| | | | |
|--|---|--------|--|
| | <p>Data received by server:</p> <p>Indicates the number of bytes received by the server since the last measurement period.</p> | Number | |
|--|---|--------|--|

19.1.2 AX Portal Statistics Test

This test reports critical statistics related to the .NET Business Connector sessions on the Microsoft Dynamics server.

| | | | |
|---|---|-------------------------|---|
| Purpose | Reports critical statistics related to the .NET Business Connector sessions | | |
| Target of the test | A Microsoft Dynamics AX server | | |
| Agent deploying the test | An internal agent | | |
| Configurable parameters for the test | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. | | |
| Outputs of the test | One set of results for the AX portal being monitored | | |
| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
| | <p>Active .NET business connector sessions to portal:</p> <p>number of currently active .NET Business Connector sessions.</p> | Number | This is a good indicator of the session load on the server. |
| | <p>Web part execution and rendering time:</p> <p>Indicates the time in seconds taken to execute and render a Web Part.</p> | Secs | A high value indicates that Web Part renditions takes too long. |
| | <p>Fatal .NET business connector session exceptions:</p> <p>Indicates the Fatal .NET business connector session exceptions.</p> | Number | For Enterprise Portal, this means that the page was not rendered. A Windows Sharepoint Services error page was displayed to the user. |

| | | | |
|--|---|------------|---|
| | <p>Non-fatal .NET business connector session exceptions:</p> <p>Indicates the number of nonfatal .NET Business Connector session exceptions.</p> | Number | For Enterprise Portal, this means that the page was rendered, but some Web Parts on the page were not rendered. |
| | <p>X++ .NET session exceptions:</p> <p>Indicates the number of X++ .NET session exceptions.</p> | Number | |
| | <p>.NET business connector sessions allocated:</p> <p>Indicates the total number of .NET Business Connector sessions allocated during the last measurement period.</p> | Number | |
| | <p>.NET business connector sessions disposed:</p> <p>Indicates the total number of .NET Business Connector sessions disposed during the last measurement period.</p> | Number | |
| | <p>.NET business connector session allocation rate:</p> <p>Indicates the NET business connector session allocation rate.</p> | Number/Sec | |

Monitoring the Microsoft RDS License Server

A Microsoft RDS License server is a computer on which the TS Licensing role service is installed. A license server stores all RDS CALs (Microsoft RDS server Client Access Licenses) that have been installed for a group of Microsoft RDS servers and tracks the RDS CALs that have been issued. One license server can serve many Microsoft RDS servers simultaneously. As clients connect to a Microsoft RDS server, the Microsoft RDS server determines if the client needs a RDS CAL, requests a RDS CAL from a license server, and then delivers that RDS CAL to the client. In the absence of RDS CALs, users will neither be able to connect to the Microsoft RDS server, nor access any of the applications published on it. To avoid this, you will have to continuously track license usage by the Microsoft RDS clients, proactively detect a potential contention for licenses, and ensure that the Terminal License server has adequate number of licenses to support the current and future load of the Microsoft RDS server.

eG Enterprise provides a *Microsoft RDS License* monitoring model that periodically monitors the usage of the licenses stored on the Microsoft RDS License server and promptly alerts administrators if the license server is about to run out of RDS CALs.

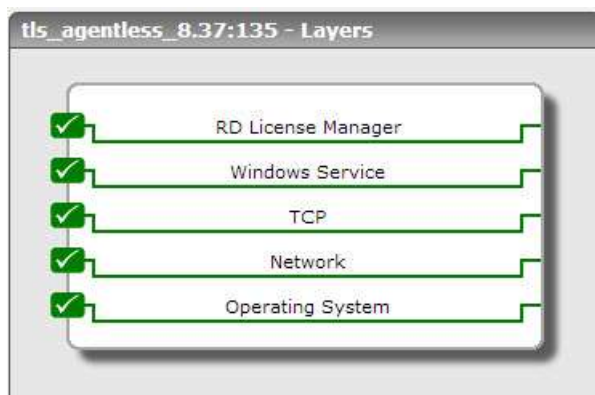


Figure 20.1: Layer model of the Microsoft RDS License server

Each layer of Figure 20.1 is mapped to a variety of tests that capture even the smallest of non-conformances that a Terminal license server experiences. Using the metrics reported, the following performance queries can be accurately answered:

MONITORING THE MICROSOFT RDS LICENSE SERVER

- Is the Microsoft RDS License server available over the network? Is it responding quickly to requests?
- How many RDS CALs are managed by the Microsoft RDS License server per Microsoft RDS server?
- Are too many users connecting to any particular Microsoft RDS server causing excessive usage of TS CALs?
- Will any Microsoft RDS server require additional licenses to be installed? If so, which Microsoft RDS server is it?
- Is any license about to expire?
- Are there any inactive licenses? If so, which ones are they?

Since the four layers at the bottom of Figure 20.1 have already been dealt with in the *Monitoring Windows and Unix Servers* document, this chapter will discuss the first layer only.

20.1 RD License Manager Layer

This layer monitors license usage.



Figure 20.2: The tests mapped to the TS CAL Licenses Utilization test

20.1.1 TS CAL Licenses Utilization Test

Without an RDS CAL, a Microsoft RDS client cannot connect to a Microsoft RDS server and access the applications operating on that server. It is hence imperative that administrators periodically check whether/not the Microsoft RDS License server has enough RDS CALs to support the current and future user load of the Microsoft RDS server. To achieve this, administrators can use the **TS CAL Licenses Utilization** test. For every Microsoft RDS server that is managed by the license server, this test reports the number, type, and usage of RDS CALs installed on the Microsoft RDS License server under a particular *Key pack ID* and purchased under a specific *License program or Purchase method* (this can be, *Unknown, Retail, Built-in, Volume, Concurrent, Temporary, Open*). Optionally, you can also

MONITORING THE MICROSOFT RDS LICENSE SERVER

group license usage by Microsoft RDS server alone (and not by key pack ID and purchase method). Using these statistics, you can rapidly detect probable license shortages and accurately point to the Microsoft RDS server that will potentially run out of licenses.

| | |
|---------------------------------|--|
| Purpose | For every Microsoft RDS server that is managed by the license server, this test reports the number, type, and usage of RDS CALs installed on the Microsoft RDS License server under a particular <i>Key pack ID</i> and purchased under a specific <i>License program or Purchase method</i> (this can be, <i>Unknown, Retail, Built-in, Volume, Concurrent, Temporary, Open</i>) |
| Target of the test | A Microsoft RDS License server |
| Agent deploying the test | An internal/remote agent |

| | | | |
|--|--|--------------------------------|------------------------------|
| <p>Configurable parameters for the test</p> | <ol style="list-style-type: none"> 1. TEST PERIOD - How often should the test be executed 2. HOST - The host for which the test is to be configured 3. PORT – Refers to the port used by the HOST. 4. REPORT TOTAL - By default, this flag is set to Yes. This indicates that by default, the test reports license usage per <i><Microsoft_RDS_server>_<License_program>_<KeypackID></i> combination and also reports the total license usage across all key packs and license programs relevant to a particular Microsoft RDS server. This is why, by default, in addition to descriptors represented by a combination of <i><Microsoft_RDS_server>_<License_program>_<KeypackID></i>, a Total descriptor also appears for this test for every Microsoft RDS server. If you want the test to report metrics per <i><Microsoft_RDS_server>_<License_program>_<KeypackID></i> combination only, then set this flag to No. 5. REPORT ONLY TOTAL - If you want the test to report metrics for the Total descriptor (of every Microsoft RDS server) alone and not for each <i><Microsoft_RDS_server>_<License_program>_<KeypackID></i> combination, set this flag to Yes. In this case, the test will report metrics for the <i>Licenses in use</i> measure alone. By default, this flag is set to No. 6. IGNORE PER USER CALS - Microsoft RDS servers can operate in two licensing modes: Per Device (default factory setting) and Per User. A Per Device CAL gives each client computer or device the right to access a Microsoft RDS server. Using Per User licensing on the other hand, one user can access a Microsoft RDS server from an unlimited number of devices and only one CAL is needed instead of a CAL for each device. If you want this test to ignore the CALs that have been installed in the 'Per User' mode when computing license usage, set this flag to Yes. By default, this flag is set to No, indicating that the test, by default, also considers the CALs installed in the per user mode when reporting license utilization. 7. IGNORE TEMPORARY LICENSES - By default, this flag is set to No. This implies that the test, by default, includes temporary licenses as well in the count of installed and used licenses. To make sure that the test disregards temporary licenses when computing license usage, set this flag to Yes. 8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| <p>Outputs of the test</p> | <p>One set of results for every <i><Microsoft_RDS_server>_<License_program>_<KeypackID></i> combination</p> | | |
| <p>Measurements made by the test</p> | <p>Measurement</p> | <p>Measurement Unit</p> | <p>Interpretation</p> |

| | <p>CAL Type:</p> <p>Indicates the CAL type of the licenses for this Microsoft RDS server, installed under this Key pack ID and license program.</p> | <p>CALs apply to either a "device" (as defined in the license agreement) or a "user". A business is free to choose either mode. In <i>Per-User</i> mode, a CAL is purchased to allow one user to connect to the server software. Any user can connect, but only one user may use a given CAL at any given time. Any number of CALs can be purchased to allow five, five hundred, or any number of users to simultaneously connect to the server. Any number of devices may connect to the server software, but only a set number of users can connect to it at once.</p> <p>Per-device mode operates in much the same way, but limits connections made by devices, rather than users. One CAL enables one device to connect to and use the server software, regardless of how many users are connecting.</p> <p>If the CALs apply to a <i>user</i>, then the value of this measure will be <i>Per User</i>. If the CALs apply to a <i>device</i>, then the value of this measure will be <i>Per Device</i>. In the case of invalid CALs, the value of this measure will be <i>Not Valid</i>. The numeric values that correspond to these measure values have been discussed in the table below:</p> <table border="1" data-bbox="933 1129 1417 1327"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>Per User</td> <td>1</td> </tr> <tr> <td>Per Device</td> <td>0 or 3</td> </tr> <tr> <td>Not Valid</td> <td>2</td> </tr> </tbody> </table> <p>Note:</p> <p>By default, the measure reports the Measure Values listed in the table above to indicate the CAL type. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</p> | Measure Value | Numeric Value | Per User | 1 | Per Device | 0 or 3 | Not Valid | 2 |
|---------------|--|--|---------------|---------------|----------|---|------------|--------|-----------|---|
| Measure Value | Numeric Value | | | | | | | | | |
| Per User | 1 | | | | | | | | | |
| Per Device | 0 or 3 | | | | | | | | | |
| Not Valid | 2 | | | | | | | | | |

| | | | |
|--|---|---------|--|
| | | | <p>This measure will not be available for the Total descriptor.</p> <p>You can use the detailed diagnosis of this measure to know the license program and expiration date of each license.</p> |
| | <p>Total license:</p> <p>Indicates the total number of licenses installed for this Microsoft RDS server under this key pack ID and license program.</p> | Number | <p>This measure will not be available for the Total descriptor.</p> |
| | <p>Available licenses:</p> <p>Indicates the number of licenses under this key pack ID and license program that are still to be used by this Microsoft RDS server.</p> | Number | <p>A high value is desired for this measure.</p> <p>This measure will not be available for the Total descriptor.</p> |
| | <p>Licenses in use:</p> <p>Indicates the number of licenses under this key pack ID and license program that are currently used by this Microsoft RDS server. For the Total descriptor, this measure reports the number of licenses currently used by this Microsoft RDS server across all relevant key pack IDs and license programs.</p> | Number | <p>A low value is desired for this measure. Compare the value of this measure for the Total descriptor across all Microsoft RDS servers to identify which Microsoft RDS server is over-utilizing the CALs.</p> <p>Using the detailed diagnosis of this measure, you can view the complete details of license usage. This includes the License ID of every license installed under a key pack, the license program under which each license was purchased, who it was issued to and when, the expiry date of license and its current status.</p> |
| | <p>License utilization:</p> <p>Indicates the percentage of licenses under this key pack ID and license program that are currently used by this Microsoft RDS server.</p> | Percent | <p>A value close to 100% indicates excessive CAL utilization. This in turn implies that too many users are connecting to the Microsoft RDS server. You may want to install additional licenses to ensure that subsequent users are able to connect to and work with the Microsoft RDS server.</p> <p>This measure will not be available for the Total descriptor.</p> |

The detailed diagnosis of the *CAL type* measure reveals the license program and expiration date of each license. If you have installed multiple licenses using a key pack ID, you can use the detailed diagnosis to know the purchase method and expiry date of every license under that key pack ID.

MONITORING THE MICROSOFT RDS LICENSE SERVER



Figure 20.3: The detailed diagnosis of the CAL type measure

Using the detailed diagnosis of the *Licenses in use* measure, you can view the complete details of license usage. This includes the License ID of every license installed under a key pack, the license program under which each license was purchased, who it was issued to and when, the expiry date of license and its current status. If you notice abnormal license usage on a Terminal license server, you can use the detailed diagnosis to figure out which Microsoft RDS server was issued the maximum number of licenses. You can also identify licenses that are inactive currently, so that such licenses can be revoked and made available for the use of active connections to the Microsoft RDS server.

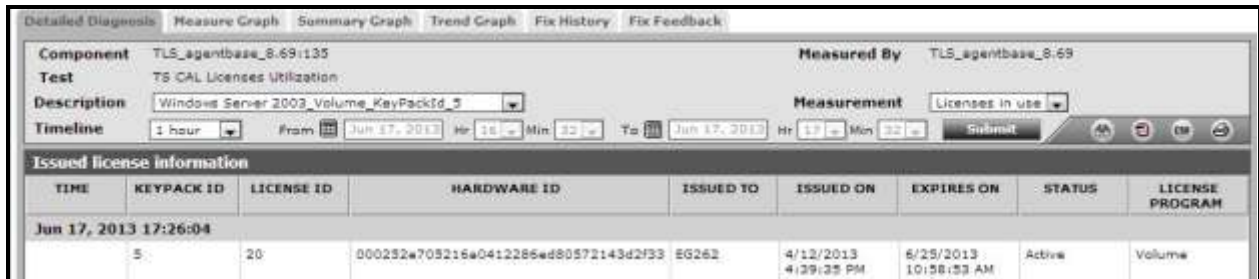


Figure 20.4: The detailed diagnosis of the Licenses in use measure

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 **Microsoft applications**. 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.