

Foglight™ for Sybase 5.7.5.40
User and Reference Guide



© 2017 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend



WARNING: A WARNING icon indicates a potential for property damage, personal injury, or death.



CAUTION: A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.



IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO: An information icon indicates supporting information.

Contents

Using Foglight for Sybase	4
Exploring the Sybase Dashboards	4
Accessing the Sybase Dashboards	5
Sybase MDA Global View Dashboard	5
MDA Home Page Dashboard	8
Top Activity Dashboard	14
Engine Activity Dashboard	35
Disk I/O Dashboard	39
Buffer Pools Dashboard	42
Database Spaces and Trends Dashboard	48
RS Home Page Dashboard	60
Replication Server Activity Dashboard	65
Databases Dashboard	68
Manage Thresholds Dashboard	74
About the Sybase_MDA Agent	74
Configuring Sybase Login Privileges	75
Configuring the Sybase_MDA Agent for Remote Monitoring	76
Sybase_MDA Agent Properties	76
Sybase_MDA Investigations	88
About the Sybase_RS Agent	96
Sybase_RS Agent Properties	96
Sybase_RS Agent Investigations	101
Investigating Health and Availability of the Replication Server	102
Generating SybaseMDA Reports	105
Cache Activity Report	106
Change Tracking Summary Report	107
Engine Activity Report	108
I/O and Network Activity Report	108
Locks Activity Report	109
Top SQL Top Procedure Report	109
Workload Summary Extended Report	109
Reference	112
Views	112
We are more than just a name	115
Our brand, our vision. Together.	115
Contacting Quest	115
Technical support resources	115

Using Foglight for Sybase

The Sybase cartridge monitors the Sybase database activity. It contains three major components:

- *Sybase cartridge dashboards* provide a visual representation of the status of the major components of the Sybase database. They allow you to determine any potential bottlenecks in database performance. For more information, see [Exploring the Sybase Dashboards](#) on page 5.
- *Sybase_MDA Agents* monitor the collection of data by the Sybase Web-enabled relational database management system.

The Sybase_MDA Agent identifies potential resource bottlenecks, investigates current server problems, and tunes the server. The Sybase_MDA Agent also allows you to select a “monitoring level” (called Availability, Performance, Configuration and Trend). Each level is associated with a group of tables and their associated collection. By default, Availability is activated and all others are turned off. A user can optionally select Performance, Configuration or Trend to tune the agent to collect the metrics important to their environment.

Some of the functions that are available with the Sybase_MDA Agent are:

- Ability to collect performance data using the MDA tables.
- The Limit Incoming Rows agent property feature allows all tables that can return multiple rows to have a corresponding entry in this agent property. This table list displays a set of criteria developed to limit the amount of data written to the Foglight database. These values are used strictly to manage the amount of data written to the Foglight database and NOT to configure the rules. Rules are still configured by registry variables.
- Improved command and connection time-out handling.
- Cluster-tolerant support allows you to monitor Sybase cluster nodes. Each cluster node is monitored separately. To enable this behavior, you need to create a separate Foglight Sybase user. The Sybase dashboards show shared resources as well as node-specific private resources for these cluster nodes. Each monitored cluster node requires a separate instance of the Sybase_MDA Agent.

For more information, see [About the Sybase_MDA Agent](#) on page 74.

- *Sybase_RS Agents* provide a consolidated view of all database instances within a Sybase server.

The Sybase_RS Agent monitors the Replication Server on the host it is deployed on. The agent collects internal performance and availability metrics by connecting to the Replication Server directly and from its related RSSD database via a connection to the respective Adaptive Server. The agent monitors Replication Server health, internal threads, partitions (space and status), connections, and exceptions. It also scans the external error log of the Replication Server looking for any default or user defined errors.

In addition, the Sybase_RS Agent also gathers availability information of the Replication Agent (RepAgent) from the primary sites (Adaptive Servers / databases) connected directly to this Replication Server as well as latency information from the Replicate Sites (Adaptive Servers / databases) controlled directly from this Replication Server.

For more information, see [About the Sybase_RS Agent](#) on page 96.

Exploring the Sybase Dashboards

The Foglight for Sybase includes dashboards which aid in the monitoring, analysis, and investigation of Sybase performance.

These dashboards provide the status of the major components of the Sybase agents. They allow you to determine any potential bottlenecks in database performance.

Accessing the Sybase Dashboards


To access the Sybase dashboards:

i | **NOTE:** The user must be a member of a Foglight group which has been assigned the Sybase DBA role in order to view the dashboards.

- 1 On the navigation panel, under **Dashboards**, click one of the following:
 - **SYBA: Sybase_MDA** to display the list of Sybase_MDA Agent views. See [Views](#) on page 112 for details.
 - **SYBA: Sybase_RS** to display the list of Sybase_RS Agent views. See [Views](#) on page 112 for details.
 - **SybaseMDA > Sybase MDA Global View** to display the [Sybase MDA Global View Dashboard](#).
 - **SybaseMDA > Sybase MDA Global View**, from there, click a Sybase_MDA agent instance to display the [MDA Home Page Dashboard](#).
 - **SybaseMDA > Sybase MDA Global View**, from there, click a Sybase_RS agent instance to display the [RS Home Page Dashboard](#).
- 2 On the navigation panel, under **Homes**, click **Databases** to display the [Databases Dashboard](#).

Sybase MDA Global View Dashboard

This Sybase Global View dashboard provides many drilldowns for Sybase agents. An overview of the dashboard navigation is shown below.



The Sybase Global View dashboard is comprised of two parts:

- [Sybase Alarms Selection](#)
- [Sybase Agent Table](#)

From the [Sybase Agent Table](#), you can drill down to [Sybase Agents Home Pages](#).

Figure 1. Sybase Agents Home Pages

Alarms

Health	Agent Name	F	C	W	Host Name	Agent Status	Sybase Version	Up Since
	isrvmrh500_Sybase		1	49	isrvmrh500	Collecting data	ASE 15.0.3 EBF 16374	Jun 16 2009 11:49AM
	devx174_Sybase			12	isrdevx174	Collecting data	ASE 12.5.4 EBF 15438 ESD#8	Jun 15 2009 11:57AM

Agent Table

Sybase Alarms Selection

By default this list shows all of the Sybase agents that are configured for this Foglight Management Server.

To view the Sybase alarms:

- 1 On the navigation panel, under Dashboards, click **Alarms > Alarms**.

The All Systems Alarms & Changes view appears.

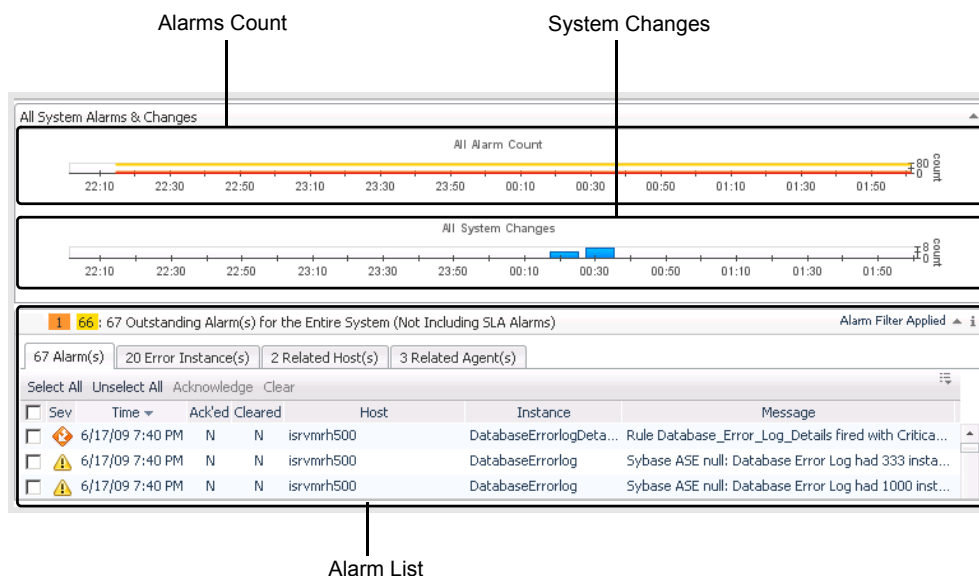
- 2 Click **All Alarms**.

The Alarms view appears with three embedded views:

- The top graph shows the Alarm Count plotted over time.
- The second graph shows the number of System Changes plotted over time.
- The table shows the Alarm List.

See the *Foglight Online Help, Monitoring System-Wide Alarms* for details on the alarm information.

Figure 2. Alarms View



- 3 Drill down on the alarms shown in the table. You can acknowledge or clear any alarm.

Sybase Agent Table

To view details of the Sybase MDA Agent table, clicking on most fields allows you to drill down for more detailed information.

The Sybase Agent Name table columns are described as follows:

Table 1. Sybase Agent Table

Name	Description
Health	<p>The health of the Sybase Agent Name, based upon the highest severity alarm.</p> <p>You can click on an Instance Health icon. A TopologyObject aggregateState Summary popup appears, providing further drill downs to:</p> <ul style="list-style-type: none">• Health of All Alarm Sources• Health of Current Object• All Outstanding Alarms
Agent Name	<p>The name of the Sybase database agent.</p> <p>You can drill down by clicking on a Sybase agent name:</p> <ul style="list-style-type: none">• Clicking a Sybase_MDA agent displays the MDA Home Page Dashboard.• Clicking a Sybase_RS agent displays the RS Home Page Dashboard.
F	<p>Number of fatal alarms for this Sybase agent. The dwell shows the number of most recent Fatal alarms raised by this agent.</p> <p>You can drill down by clicking on the alarm number to view the Alarms list. This allows you to browse through the alarms.</p>
C	<p>Number of critical alarms for this Sybase agent. The dwell shows the number of most recent Critical alarms raised by this agent.</p> <p>You can drill down by clicking on the alarm number to view the Alarms list. This allows you to browse through the alarms.</p>
W	<p>Number of warning alarms for this Sybase agent. The dwell shows the number of most recent Warning alarms raised by this agent.</p> <p>You can drill down by clicking on the alarm number to view the Alarms list. This allows you to browse through the alarms.</p>
Host Name	<p>Name of the computer which is hosting the Sybase database.</p> <p>You can drill down by clicking on a host name and then selecting:</p> <ul style="list-style-type: none">• Host Monitor—a detailed dashboard for the selected host appears.• Host Browser—a dashboard appears listing the agents running on this host and any alarms for this host.

Table 1. Sybase Agent Table

Name	Description
Agent Status	<p>The operational status of the monitoring agent.</p> <p>When the agent instance is running, the following status message appears:</p> <ul style="list-style-type: none">Collecting Data <p>When the agent instance is running but not collecting data, one of the following status messages appear:</p> <ul style="list-style-type: none">StartingStoppedStoppingUnknown <p>If you have administrator role permissions, click on the status and then select a command to:</p> <ul style="list-style-type: none">Activate the agent.Deactivate the agent.Start Data Collection by the agent.Stop Data Collection by the agent.Edit Properties of the agent.
Sybase Version	<p>The Sybase version being monitored.</p> <p>You can expand the version detail by clicking on any version number.</p>
Up Since	The date and time that the Sybase agent was last restarted.

Sybase Agents Home Pages

Clicking an agent name on the [Sybase MDA Global View Dashboard](#) or [Databases Dashboard](#) shows a dashboard containing more information about the selected agent instance. The dashboard that appears in the display area depends on the selected agent type:

- Clicking a Sybase_MDA agent instance displays the [MDA Home Page Dashboard](#).
- Clicking a Sybase_RS agent instance displays the [RS Home Page Dashboard](#).

MDA Home Page Dashboard

The MDA Home Page dashboard provides performance details for a specific database. The following sub-sections describe the dashboard element metrics:

- [MDA Instance Metrics](#)
- [MDA Alarms](#)
- [MDA Processes Metrics](#)
- [MDA Engines Metrics](#)
- [MDA Shared Memory Metrics](#)
- [MDA Host Metrics](#)
- [MDA Devices Metrics](#)
- [MDA Input/Output Metrics](#)

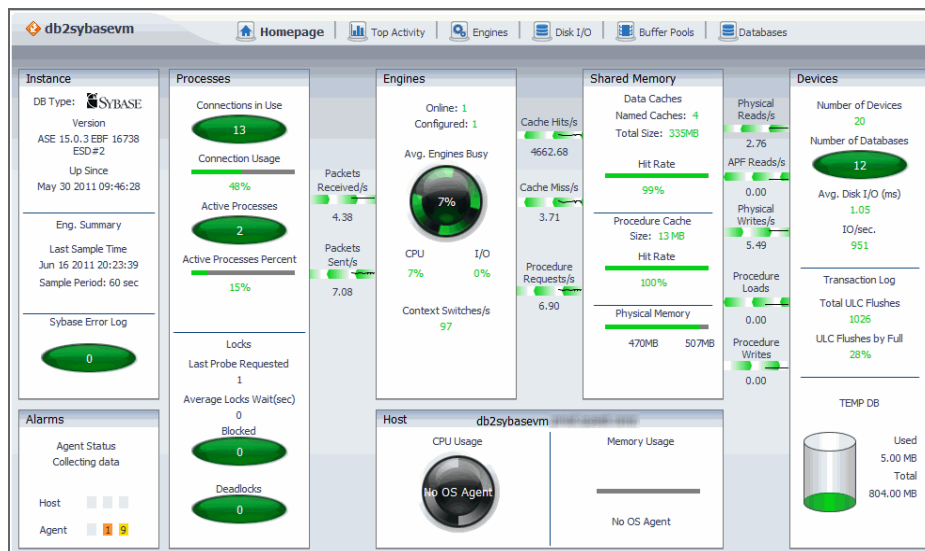
Drilldown dashboards may be accessed through some of the dashboard element metrics or directly through the dashboard toolbar, where:

Table 2. Drilldown dashboards

Name	Description
Homepage	This link opens the MDA Home Page Dashboard .
Top Activity	This link opens the Top Activity Dashboard .
Engines	This link opens the Engine Activity Dashboard .
Disk I/O	This link opens the Disk I/O Dashboard .
Buffer Pools	This link opens the Buffer Pools Dashboard .

If no data is being collected by the Sybase_MDA agent, the “Not Collecting data” error message appears beside the Sybase agent name. In this case, the dashboard shows the data collected from the last sample time.

Figure 3. Dashboard data



Click a point on the plot line of any metric to zoom in to the corresponding graph. Click a data point on the plot line of the graph to view the corresponding top SQL (displayed on the Top SQL tab of the Top Activity dashboard). For more information, see [Top SQL Tab](#) on page 15.

MDA Instance Metrics

Table 3. MDA Instance Metrics

Name	Description
DB Type	Type of database.
Version	The Sybase version being monitored. You can expand the version detail by clicking on or hovering on any version number.
Up Since	The date and time that the Sybase agent was last restarted.
Eng. Summary	
Last Sample Time	The last date (month/day/year) and time (hour:minute, AM/PM) that data was sampled.
Sample Period	The frequency (in seconds) of data sampling.

Table 3. MDA Instance Metrics

Name	Description
Sybase Error Log	<p>The number of Sybase errors.</p> <p>This field displays color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the device symbol. A popup appears showing a detailed list of the errors.</p>

MDA Alarms

Table 4. MDA Alarms

Name	Description
Agent Status	<p>When the agent instance is running, the following status message appears:</p> <ul style="list-style-type: none"> Collecting Data <p>When the agent instance is running but not collecting data, one of the following status messages appear:</p> <ul style="list-style-type: none"> Starting Stopped Stopping Running but not collecting data (initiated by clicking Unknown
Host	<p>The number of fatal, critical, and warning alarms for this host.</p> <p>You can drill down to the detailed alarm list.</p>
Agent	<p>The number of fatal, critical, and warning alarms for this database agent.</p> <p>You can drill down to the detailed alarm list.</p>

MDA Processes Metrics

Table 5. MDA Processes Metrics

Name	Description
Connections in Use	<p>The number of client application connections currently in use.</p> <p>You can drill down by clicking on the device symbol. A popup appears showing the number of current connections in use plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> Top Activity Drilldown—a Top Activity Dashboard appears. Edit Threshold—a Manage Thresholds Dashboard appears.
Connection Usage	<p>The percentage of total available connections currently in use.</p> <p>You can drill down by clicking on the number. A popup appears showing the percentage of connections that are available, plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> Top Activity Drilldown—a Top Activity Dashboard appears. Edit Threshold—a Manage Thresholds Dashboard appears.
Active Processes	<p>The number of active client application connections.</p> <p>You can drill down by clicking on the device symbol. A popup appears showing the active connections plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> Top Activity Drilldown—a Top Activity Dashboard appears. Edit Threshold—a Manage Thresholds Dashboard appears.

Table 5. MDA Processes Metrics

Name	Description
Active Processes Percent	<p>The percentage of total available processes currently in use.</p> <p>This field displays color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the number. A popup appears showing the percentage of processes that are active, plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> • Top Activity Drilldown—a Top Activity Dashboard appears. • Edit Threshold—a Manage Thresholds Dashboard appears.
Locks	
Last Probe Requested	<p>The number of the last probe requested.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Average Locks Wait (sec)	<p>The average amount of time spent waiting for a lock to release.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Blocked	<p>The number of blocked processes detected.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Deadlocks	<p>The number of deadlocks detected.</p> <p>This field displays color-coded real-time behavior for severity.</p>

MDA Engines Metrics

Table 6. MDA Engines Metrics

Name	Description
Online	<p>The number of database engines currently online.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Configured	<p>The number of database engines currently configured.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Avg. Engines Busy	<p>The average number of database engines that are busy.</p> <p>This field displays color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the icon. A popup appears showing the percentage of engine total busy and engine I/O busy, plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> • Engines Drilldown—a Engine Activity Dashboard appears. • Edit Threshold—a Manage Thresholds Dashboard appears.
CPU	<p>The percentage of CPU time all engines spent not idle.</p> <p>This field displays color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the number. A popup appears showing the percentage of engine total busy and engine I/O busy, plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> • Engines Drilldown—a Engine Activity Dashboard appears. • Edit Threshold—a Manage Thresholds Dashboard appears.

Table 6. MDA Engines Metrics

Name	Description
I/O	<p>The percentage of input/output resources currently being used by the database engines. This field displays color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the number. A popup appears showing the percentage of engine total busy and engine I/O busy, plotted over time. It provides further drill downs to:</p> <ul style="list-style-type: none"> • Engines Drilldown—a Engine Activity Dashboard appears. • Edit Threshold—a Manage Thresholds Dashboard appears.
Context Switches /s	<p>The number of context switches for the engines in this database, per second.</p> <p>This field displays color-coded real-time behavior for severity.</p>

MDA Shared Memory Metrics

Table 7. MDA Shared Memory Metrics

Name	Description
Data Caches	
Named Caches	<p>The number of named caches.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Total Size	<p>The total size of the data cache in megabytes.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Hit Rate	<p>The percentage of times a requested page occurred in the named cache.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Procedure Cache	
Size	The total size of the procedure cache in megabytes.
Hit Rate	<p>The percentage of times a requested page occurred in the procedure cache.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Physical Memory	
Physical Memory	<p>The size of the used and total available physical memory in megabytes. This is shown in the percentage bar and as numbers:</p> <ul style="list-style-type: none"> • The size of the used physical memory is shown on the left. • The size of the total available physical memory is shown on the right. <p>This field displays color-coded real-time behavior for severity.</p>

MDA Host Metrics

Table 8. MDA Host Metrics

Name	Description
CPU Usage	The amount of CPU resources being consumed by the database host.
Memory Usage	The amount of memory being consumed by the database host.

MDA Devices Metrics

Table 9. MDA Devices Metrics

Name	Description
Number of Devices	<p>The number of devices in this database.</p> <p>This field displays color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the number. A popup appears showing the Disk I/O Summary chart. It provides further drill downs. Click:</p> <ul style="list-style-type: none">• Show Me the “Disk I/O” Drilldown—a Disk I/O Dashboard appears.
Number of Databases	<p>The number of databases in this instance.</p> <p>This field displays color-coded real-time behavior for severity.</p>
Avg. Disk I/O (ms)	The average disk input/output time, in milliseconds.
IO/sec	The number of inputs/outputs per second.
Transaction Log	
Total ULC Flushes	The total number of User Log Cache flushes.
ULC Flushes by Full TEMP DB	The percentage of User Log Cache flushes resulting from the cache becoming full.
Used	The amount of memory (in megabytes) in the temp database currently in use.
Total	The total amount of memory (in megabytes) in the temp database that is available for use.

MDA Input/Output Metrics

Table 10. MDA Input/Output Metrics

Name	Description
Packets Received/s	<p>The number of message packets received per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the Packet Received chart.</p>
Packets Sent/s	<p>The number of message packets sent per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the Packet Sent chart.</p>
Cache Hits/s	<p>The number of cache hits per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the Cache Hits chart.</p>
Cache Miss/s	<p>The number of cache misses per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the Cache Misses chart.</p>
Procedure Requests/s	<p>The number of procedure requests per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the Procedure Cache chart.</p>
Physical Reads/s	<p>The number of physical reads per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the Physical I/O Reads chart.</p>
APF Reads/s	<p>The number of APF (Asynchronous Prefetch Activity) reads per second.</p> <p>You can drill down by clicking on the number. A popup appears showing the APF Reads chart.</p>

Table 10. MDA Input/Output Metrics

Name	Description
Physical Writes/s	The number of physical writes per second. You can drill down by clicking on the number. A popup appears showing the Disk I/O chart.
Procedure Loads	The number of procedure loads. You can drill down by clicking on the number. A popup appears showing the Procedure Loads chart.
Procedure Writes	The number of procedure writes. You can drill down by clicking on the number. A popup appears showing the Procedure Writes chart.

Top Activity Dashboard

The Top Activity dashboard provides performance details for SQL operations. This dashboard contains the following tabs:

- [Top Hash Tab](#)
- [Top SQL Tab](#)
- [Top Procedures Tab](#)
- [Top Users Tab](#)
- [Current Top Sessions Tab](#)
- [SQL Details Tab](#)
- [Blocked Tree Tab](#)
- [System Wait Events Tab](#)
- [Spinlocks Tab](#)
- [Execution Time Tab](#)

Top Hash Tab

The Top Hash information is provided in a table. The Top Hash table has the same columns and drilldowns as the [Top SQL Table](#).

Adaptive Server uses the statement cache to store the text of cached SQL statements. This tab displays the SQL statements executed from the statement cache along with their bind variables, if they exist. Information that appears on this tab depends on a set of the following Sybase configuration parameters:

```
sp_configure "enable stmt cache monitoring" 1
sp_configure 'statement cache size'
sp_configure 'enable literal autotparam'
```


i | **TIP:** If `enable literal autotparam` is disabled, all SQL statements appear separately. Using this parameter has a number of advantages. For complete information, see your Sybase documentation.

For complete information about the required parameters, see “Adaptive Server Enterprise configuration parameters” in the *Foglight for Sybase Release Notes*.

Selecting an option button for an individual SQL text line shows the [SQL Execution Statistics Dialog Box](#).. This dialog box shows the SQL details about a selected line of SQL text.

i | **IMPORTANT:** This dashboard is only supported with Sybase Adaptive Server Enterprise version 15.02 and higher.

Figure 4. SQL Execution Statistics Dialog Box.



Top Hash Filters

The Top Hash filters are the same as the [Top SQL Filters](#).

Server Statistics Pane

The Server Statistics pane on the Top Hash tab is the same as the [Server Statistics Pane](#) on the Top SQL tab.

For more information on the Server Statistics pane, click the ? toward the top left of the pane.

Top SQL Tab

The Top SQL information is provided in a table, which is sorted by the Avg Duration (ms) column.

Figure 5. Top SQL Tab

SQL Text	SQL Executions	Avg Duration (ms)	Max Duration (ms)	Min Duration (ms)	Avg CPU (ms)	Avg Phys. Reads	Avg Logical Reads	Avg Wait Time (ms)
SELECT EngineNumber, ErrorMessage, ErrorNumber, Severity, Time, FamilyID, KPID, SPID, 0 ...	10	61	63	60	36	0	241	0
SET cis_rpc_handling off exec SYB_BACKUP...sp_who	9	60	70	50	18	0	19	0
EXEC sybtempprocs.sp_fgl_dbspace 'vmrh502_tdb_1'	10	55	63	50	55	0	17	0
EXEC sybtempprocs.sp_fgl_dbspace 'master'	10	42	59	26	42	0	39	0
SELECT name FROM master..sysdatabases	9	31	33	30	31	0	38	0
EXEC sybtempprocs.sp_fgl_dbspace 'sybtempdb'	10	28	40	23	28	0	26	0
EXEC sybtempprocs.sp_fgl_dbspace 'model'	10	26	40	20	26	0	36	0
EXEC sybtempprocs.sp_fgl_dbspace 'tempdb'	10	26	33	23	26	0	27	0
EXEC sybtempprocs.sp_fgl_dbspace 'sybtempprocs'	10	24	33	20	24	0	6	0
EXEC sybtempprocs.sp_fgl_dbspace 'vmrh502_tdb_2'	10	19	23	16	19	0	0	0

Top SQL Table

Table 11. Top SQL Table

Name	Description
Option button	Shows additional information about the selected statement. For more information, see SQL Execution Statistics Dialog Box . on page 15.
SQL Text	The SQL (Structured Query Language) statement.
SQL Executions	The number of SQL executions for this SQL statement.
Avg. Duration (ms)	The average amount of time the SQL statement took to run in milliseconds.
Max. Duration (ms)	The maximum amount of time the SQL statement took to run in milliseconds.
Min. Duration (ms)	The minimum amount of time the SQL statement took to run in milliseconds.
Avg. CPU (ms)	The average CPU time spent executing the SQL statement, in milliseconds.
Avg. Phys. Reads	Average physical reads: the average number of buffers read from disk.
Avg. Logical Reads	Average logical reads: the average number of buffers read from cache.

Table 11. Top SQL Table

Name	Description
Avg. Packets Sent	Average packets sent: the average number network packets sent by Adaptive Server.
Avg. Wait Time (ms)	Average wait time in milliseconds: the average number of milliseconds a task has waited during execution of the SQL statement.

Top SQL Filters

This filter allows you to thoroughly investigate a large number of SQL statements and is especially helpful when a long retention policy is configured. The filter applies to the Top SQL, Top Procedures and Top Hash values.

The following rules apply to the search methods:

- The search is *not* case sensitive (for example, "abc", "Abc" and ABC" are considered to be the same string).
- The search does not support wildcards, but does not require you to enter the entire word.
- To exclude certain strings from search, use the dash (-) sign
For example, type `-sa` to exclude sa user
- To search using multiple parameters, separate each parameter with a comma
For example, to find all abc and def users, type "abc,def"
- To zoom-in on a specific point in time, rather than the entire time range:
 - Click the requested point on the plotted line (for example: 15:40).
 - Data is then displayed for this point, and for points that are within a user-defined amount of time before and after it, as specified in seconds on the Time before and after (seconds) field.

In the example above, if the specified time is 15:40 and the value entered in the field is 300, the search would include the time range between 15:35 and 15:45.

The search process is carried out using the following steps:


- The SQL statements are filtered using one or more fields, and are sorted based on the criterion selected in the "Sort by" list.
For example, only the SQL statements with the login name "QA" and the application name "DBISQL" will be retrieved.
- Because even after filtering, the number of rows retrieved can be extremely large (50,000 or higher), the value entered in the "Max Rows After Filter" field of the Default values section to the right of this box limits the number of rows to be aggregated (by default: the top 5,000 rows).
- The filtered and sorted SQL statements are aggregated.

The combination of a filter and the "Max rows after filter" field significantly reduces the dashboard response time, by limiting the highly time-consuming aggregation process.

Sybase_MDA Agent Max Response Time Graph View

Purpose — The Max Response Time graph plots the SQL queries maximum response time and the SQL procedures maximum response time, for the database, over the specified time range.

Figure 6. Sybase_MDA Agent Max Response Time Graph View



Description of the View

Table 12. Max Response Time View


Data Displayed	SQLs Max Response Time. The maximum amount of time (in ms) it took the database to respond to a SQL statement at a given time, plotted over the specified time range.
	Proc's Max Response Time. The maximum amount of time (in ms) it took the database to respond to a procedure SQL statement at a given time, plotted over the specified time range.
Where to go next	Drill down on: <ul style="list-style-type: none">• SQLs Max Response Time plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.• Proc's Max Response Time plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Sybase_MDA Agent Engine Total Busy Graph View

Purpose

The Engine Total Busy graph plots the percent of the engine (total) that was busy at a given time, the percent of the engine I/O that was busy at a given time, and the percent of the engine CPU that was busy at a given time, for the database, over the specified time range.

Figure 7. Sybase_MDA Agent Engine Total Busy Graph View



Description of the View

Table 13. Sybase_MDA Agent Engine Total Busy Graph View Description


Data Displayed	Avg. Engine Total Busy. The percent of the engine (total) that was busy at a given time, plotted over the specified time range.
	Avg. Engine I/O Busy. The percent of the engine I/O that was busy at a given time, plotted over the specified time range.
	Avg. Engine CPU Busy. The percent of the engine CPU that was busy at a given time, plotted over the specified time range.
Where to go next	Drill down on:
	• Avg. Engine Total Busy plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.
	• Avg. Engine I/O Busy plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.
	• Avg. Engine CPU Busy plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Sybase_MDA Agent Cache Hits/s Graph View

Purpose

The Cache Hits/s graph plots the number of hits per second, the number of misses per second, and the number of dirty buffers per second, for the database, over the specified time range.

Figure 8. Sybase_MDA Agent Cache Hits/s Graph View



Description of the View

Table 14. Sybase_MDA Agent Cache Hits/s Graph View Description

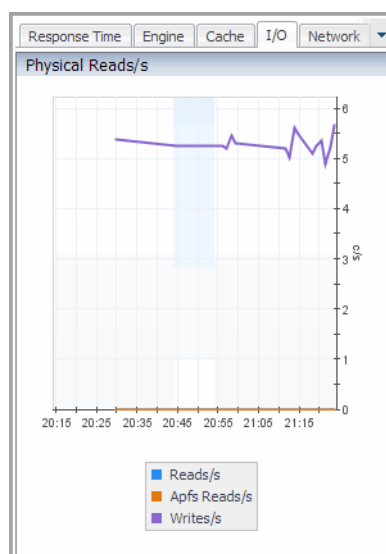
Data Displayed	Hits/s. The number of hits per second, plotted over the specified time range.
	Miss/s. The number of misses per second, plotted over the specified time range.
	Dirty Buffer/s. The number of dirty buffers per second, plotted over the specified time range.
Where to go next	Drill down on:
	<ul style="list-style-type: none">• Hits/s plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.
	<ul style="list-style-type: none">• Miss/s plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.
	<ul style="list-style-type: none">• Dirty Buffer/s plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Sybase_MDA Agent I/O Physical Reads Graph View

Purpose

The I/O Physical Reads graph plots the number of physical reads per second, the number of APFs per second, and the number of physical writes per second, for the database, over the specified time range.

Figure 9. Sybase_MDA Agent I/O Physical Reads Graph View



Description of the View

Table 15. Sybase_MDA Agent I/O Physical Reads Graph View Description


Data Displayed	Reads/s. The number of physical reads per second, plotted over the specified time range.
	Apfs Reads/s. The number of APFs per second, plotted over the specified time range.
	Writes/s. The number of physical writes per second, plotted over the specified time range.
Where to go next	Drill down on: <ul style="list-style-type: none"> • Reads/s plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23. • Apfs Reads/s plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23. • Writes/s plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Sybase_MDA Agent Network Graph View

Purpose

The Network graph plots the packet received rate and the packet sent rate, for the database, over the specified time range.

Figure 10. Sybase_MDA Agent Network Graph View



Description of the View

Table 16. Sybase_MDA Agent Network Graph View Description


Data Displayed	Packet Received Rate. The packet received rate, plotted over the specified time range.
	Packet Sent Rate. The packet sent rate, plotted over the specified time range.
Where to go next	Drill down on:
	<ul style="list-style-type: none">• Packet Received Rate plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.• Packet Sent Rate plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Sybase_MDA Agent Waits Graph View

Purpose

The Waits graph plots the total waits for the database over the specified time range.

Figure 11. Sybase_MDA Agent Waits Graph View



Description of the View

Table 17. Sybase_MDA Agent Waits Graph View Description


Data Displayed	Total Waits. The total waits, plotted over the specified time range.
Where to go next	Drill down on: <ul style="list-style-type: none">Total Waits plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Sybase_MDA Agent Active Users Graph View

Purpose

The Active Users graph plots the total number of active users at a given time, for the database, over the specified time range.

Figure 12. Sybase_MDA Agent Active Users Graph View



Description of the View

Table 18. Sybase_MDA Agent Active Users Graph View Description

Data Displayed	Active Processes. The total number of active users at a given time, plotted over the specified time range.
Where to go next	Drill down on: <ul style="list-style-type: none">• Active Processes plot line. Click a data point on the plot line to filter the Top Hash, Top SQL, or Top Procedures table according to a defined and corresponding time range. For more information, see Zooming In on a Data Point Using the Filter Pane on page 23.

Zooming In on a Data Point Using the Filter Pane

If you click a data point on any plot line in a Filter pane graph to zoom in on that data point, the **Zoom-in on a data point** check box is selected and the exact time at which the data point was recorded is displayed in the **Selected data metric point** box.

You can adjust the number of seconds for which you want data displayed before and after the selected data metric point by adjusting the value in the **Metric point time margins (sec)** box.

To reset the Filter pane to the default values at any time, click **Reset**.

For more information on the Filter pane, click the **?** at the top right of the pane.

SQL Execution Statistics Dialog Box

This dialog box shows additional information about a selected SQL statement. This dialog appears when you select an option button on the **Top Hash**, **Top SQL**, or the **Top Procedures** tab.

Figure 13. SQL Execution Statistics Dialog Box

SQL Execution Statistics

Top Procedure(Period)- sample rate: 60 sec

sybsystemprocs..sp_fgl_dbseg_space;line-156

SQL Details Statistics

Line Number	SQL Executions	Avg Duration (ms)	Max Duration (ms)	Min Duration (ms)	Avg CPU (ms)	Avg Phys. Reads	Avg Logical Reads	Avg Wait Time (ms)	Avg Packets Sent	Avg Packets Recvd	Avg. Rows	Max Rows	Min Rows
156	11	28	33	3	28	0	0	0	0	0	1	1	1

Separate SQL Statement statistics

End Time	DB Name	Application	Line No	SPID	Login	Duration (ms)	CPU Time (ms)	Phys. Reads	Logical Reads	Wait Time (ms)	Packets Sent	Packets Received	Rows Affected	Error Status	Plan ID	Compile Date
Mar 7 17:17:34:483	sybsystemprocs	FoglightMDA	line-156	393	fogvm	30	30	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:16:34:876	sybsystemprocs	FoglightMDA	line-156	393	fogvm	16	16	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:15:34:360	sybsystemprocs	FoglightMDA	line-156	393	fogvm	30	30	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:14:34:610	sybsystemprocs	FoglightMDA	line-156	393	fogvm	33	33	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:13:33:906	sybsystemprocs	FoglightMDA	line-156	393	fogvm	30	30	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:12:33:843	sybsystemprocs	FoglightMDA	line-156	393	fogvm	33	33	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:11:33:123	sybsystemprocs	FoglightMDA	line-156	393	fogvm	30	30	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:10:32:656	sybsystemprocs	FoglightMDA	line-156	393	fogvm	30	30	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:09:32:180	sybsystemprocs	FoglightMDA	line-156	393	fogvm	30	30	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34
Mar 7 17:08:32:336	sybsystemprocs	FoglightMDA	line-156	393	fogvm	33	33	0	0	0	0	0	1	0	29,780	Mar 3 2013 20:54:34

Table 19. SQL Execution Statistics Dialog Box Description

Name	Description
SQL Text	The SQL (Structured Query Language) statement.
SQL Details Statistics	
SQL Executions	The number of SQL executions for this SQL statement.
Avg Duration (ms)	The average amount of time the SQL statement took to run in milliseconds.
Max Duration (ms)	The maximum amount of time the SQL statement took to run in milliseconds.
Min Duration (ms)	The minimum amount of time the SQL statement took to run in milliseconds.
Avg CPU (ms)	The average CPU time spent executing the SQL statement, in milliseconds.
Avg Phys. Reads	Average physical reads: the average number of buffers read from disk.
Avg Logical Reads	Average logical reads: the average number of buffers read from cache.
Avg Wait Time (ms)	Average wait time in milliseconds: the average number of milliseconds a task has waited during the execution of the SQL statement.
Avg Packets Sent	Average packets sent: the average number network packets sent by Adaptive Server.
Avg Packets Recvd	Average packets received: the average number network packets received by Adaptive Server.
Avg. Rows	The average number of rows affected by the SQL statement.
Max Rows	The maximum number of rows affected by the SQL statement.
Min Rows	The minimum number of rows affected by the SQL statement.
Separate SQL Statement statistics	
End Time	<p>The time that the SQL statement finished running, in the format:</p> <p><month> <day> <hour> <minute> <second> <millisecond>.</p> <p>Click the end time to view a popup that provides more information about the SQL text, including execution details.</p>

Table 19. SQL Execution Statistics Dialog Box Description

Name	Description
DB Name	The name of the database that this SQL statement used during its execution.
Application	The name of the application associated with this SQL statement.
SPID	The ID of the session process that executed the SQL statement.
Login	The login name of the user associated with the session process that executed the SQL statement.
Duration (ms)	The time it took to execute the SQL statement, in milliseconds.
CPU Time (ms)	The time it took the CPU to execute the SQL statement, in milliseconds.
Phys. Reads	The number of buffers read from disk.
Logical Reads	The number of buffers read from cache.
Wait Time (ms)	The number of milliseconds a task has waited during execution of the SQL statement.
Packets Sent	The number of network packets sent by the Adaptive Server.
Packets Received	The number of packets received by the Adaptive Server.
Rows Affected	The average number of rows affected by the SQL statement.
Error Status	Indicates if any errors are encountered during the statement execution.
Plan ID ^a	The unique identifier for the query plan for the object in the procedure cache.
Compile Date ^a	The date that the procedure was compiled.

a. This column is only available on the **Top Activity > Top Procedures** tab.

Server Statistics Pane

The Server Statistics pane uses graphs and tables to present performance-related statistics and their impact, with focus on the top-consuming SQL queries and SQL procedures, for the entire ASE Server during the specified time range.


On any graph, hover over a point on a plot line for precise details about that point.

For more information on the Server Statistics pane, click the **?** toward the top left of the pane.

Top Procedures Tab

The top procedures information is provided in a table, which is sorted by the Avg Duration (ms) column. The Top Procedures table has the same columns as the [Top SQL Table](#). In addition, the **SQL Text** column provides a navigation tree for each top procedure SQL statement, which contains the procedure name data. Expanding the root procedure node lists the line numbers containing the SQL statements in the procedure for which the execution statistics are available.

Figure 14. Expanded root procedure node list



Click an individual line to display the stored procedure text with its execution statistics in the [Procedure Activity Drilldown](#). This behavior is enabled using the call back functionality. To use this feature, you need to configure a set of call back parameters. For more information, see [Setting Callback Connection Properties](#) on page 78.

Selecting an option button for an individual line shows the [SQL Execution Statistics Dialog Box](#). This dialog box shows the SQL details about a selected procedure.

The procedures shown in this table are determined by the values set for the following Sybase_MDA agent properties:

- Top SQL - Max Rows Returned
- Top SQL - Minimum Threshold Value
- Top SQL - Sort Criteria
- Top SQL - Show Top Procedure Lines

For more information, see [Setting Data Retrieval Properties](#) on page 81.

Figure 15. SQL details

SQL Text	Period	Last Sample	Avg. Duration (ms)	Max Duration (ms)	Min Duration (ms)	Avg CPU (ms)	Avg Phys. Reads	Avg Logical Reads	Avg Packets Sent	Avg Wait Time (ms)
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			121	123	120	90	0	127	0	31
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			1	120	120	120	0	520	0	0
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			60	91	100	90	70	185	0	21
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			63	80	109	46	71	307	0	7
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			4	26	26	19	0	45	0	6
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			63	11	30	10	9	162	0	1
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			11	28	33	3	28	0	0	0
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			2	15	30	10	15	180	0	0
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			10	65	90	60	39	206	0	25
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			5	65	86	60	53	223	0	12
SELECT SPID, CPUTime, PhysicalReads, LogicalReads, PhysicalWrites, MemUsageKB			10	60	63	60	56	208	0	24

Top Procedures Filters

The Top Procedures filters are the same as the [Top SQL Filters](#).

Procedure Activity Drilldown

The Procedure activity drilldown shows a stored procedure's text and the execution statistics for each executed SQL statement.

Figure 16. Procedure Activity Drilldown

Procedure Text	SQL Executions	Avg Duration (ms)	Max Duration (ms)	Min Duration (ms)	Avg CPU (ms)	Avg Phys. Reads	Avg Logical Reads	Avg Wait Time	Avg Packets Sent	Avg Packets Recvd	Avg Rows
1 : 2 : CREATE procedure sp_fgl_topapplications 3 : @criteria integer, 4 : @threshold integer, 5 : @topRows integer = 50 6 : as 7 : -- 8 : -- 9 : -- @criteria values: 10 : -- 1 CPUTime(seconds) 11 : -- 2 Physical Reads 12 : -- 3 Logical Reads 13 : -- 4 Physical Writes 14 : -- 5 Memory Usage(KB) 15 : -- 16 : 17 : SELECT SPID, CPUTime, PhysicalReads, LogicalReads, 18 : PhysicalWrites, MemUsageKB 19 : into #tmpMonProcessActivity 20 : from master..monProcessActivity 21 : 22 : SELECT SPID, Application 23 : into #utmpMonProcess 24 : from master..monProcess 25 : 26 : SELECT sum(b.CPUTime) AS CPUTime, 27 : sum(b.PhysicalReads) as PhysicalReads, sum(b.LogicalReads) ... 28 : sum(b.PhysicalWrites) as PhysicalWrites, 29 : sum(b.MemUsageKB) as MemUsageKB, 30 : case 31 : when a.Application like 'Omni%' then 'OmniServer' 32 : else a.Application	1	200	200	200	0	0	120	200	0	0	0

Navigate through the text and statistics using **PREV**, **FULL TEXT**, and **NEXT**.

Table 20. Procedure Activity Drilldown Description

Name	Description
Procedure Text	The SQL (Structured Query Language) statement.
SQL Executions	The number of SQL executions for this SQL statement.

Table 20. Procedure Activity Drilldown Description

Name	Description
Avg Duration (ms)	The average amount of time the SQL statement took to run, in milliseconds.
Max Duration (ms)	The maximum amount of time the SQL statement took to run, in milliseconds.
Min Duration (ms)	The minimum amount of time the SQL statement took to run, in milliseconds.
Avg CPU (ms)	The average CPU time spent executing the SQL statement, in milliseconds.
Avg Phys. Reads	Average physical reads: the average number of buffers read from disk.
Avg Logical Reads	Average logical reads: the average number of buffers read from cache.
Avg Wait Time (ms)	Average wait time in milliseconds: the average number of milliseconds a task has waited during execution of the SQL statement.
Avg Packets Sent	Average packets sent: the average number of network packets sent by Adaptive Server.
Avg Packets Recvd	Average packets received: the average number of network packets received by Adaptive Server.
Avg Rows	The average number of rows affected by the SQL statement.

Server Statistics Pane

The Server Statistics pane on the Top Procedures tab is the same as the [Server Statistics Pane](#) on the Top SQL tab.

For more information on the Server Statistics pane, click the ? toward the top left of the pane.

Top Users Tab

This tab covers the same users listed in the [Current Top Sessions Tab](#).

The difference are:

- The Current Top Session tab shows data aggregated since user login.
- The Top Users tab shows data aggregated for each user according to the time range. This allows you to view users that were displayed in the Current Top Session tab, but may not currently be displayed; for example, the user has logged out.

The users shown in this table are determined by the values set for the following Sybase_MDA agent properties:

- Top Users - Max Rows Returned
- Top Users - Minimum Threshold Value
- Top Users - Sort Criteria

For more information, see [Setting Data Retrieval Properties](#) on page 81.

Top user information is provided in two tables:

- [Top Users Table](#)
- [SQL Text Table](#)

Figure 17. Top User Information

Top Hash	Top SQL	Top Procedures	SQL Details	Top Users	Current Top Sessions	Blocked Tree	System Wait Events
Top Users							
Login Name	CPU Time (ms)	Memory Usage (kB)	Logical Reads	Physical Reads	Physical Writes		
fog5582	484,800	26	85,813,389	1,367	841		

SQL Text	SQL Executions	Avg Duration (ms)	Max Duration (ms)	Min Duration (ms)	Avg CPU (ms)	Avg Phys. Reads	Avg Logical Reads	Avg Packets Sent	Avg Wait Time (ms)
EXEC sybsystemprocs..sp_fgl_dbseg_space 'master'	6	15	25	13	15	0	38	0	0
sybsystemprocs..sp_fgl_topsql; line-45	6	17	20	16	17	0	503	0	0
EXEC sybsystemprocs..sp_fgl_dbseg_space 'model'	1	16	16	16	16	0	20	0	0
EXEC sybsystemprocs..sp_fgl_dbseg_space 'sales'	6	15	19	13	15	0	8	1	0
EXEC sybsystemprocs..sp_fgl_dbseg_space 'sybsystemdb'	2	13	13	13	13	0	41	0	0
EXEC sybsystemprocs..sp_fgl_dbseg_space 'sybsystemprocs'	3	13	13	13	13	0	20	0	0
EXEC sybsystemprocs..sp_fgl_dbseg_space 'tempdb'	3	13	13	13	13	0	6	0	0
sybsystemprocs..sp_fgl_datacached_objects; line-32	6	12	13	10	12	0	61	1	0

Top Users Table

Table 21. Top User Table Description

Name	Description
Login Name	The login name of the top user.
CPU Time (ms)	The CPU time (in milliseconds) used by the process.
Memory Usage (kB)	The amount of memory (in kilobytes) allocated to the process.
Logical Reads	The number of buffers read from cache.
Physical Reads	The number of buffers read from disk.
Physical Writes	The number of buffers written to disk.

SQL Text Table

This table has the same columns as the [Top SQL Table](#). The only difference is that the SQL Text table provides data for selected user-executed SQL statements.

Current Top Sessions Tab

Current top session information is provided in three tables:

- [Top Sessions Table](#)
- [Top SQL of Session with SPID Table](#)
- [Wait Events of SPID xx for Period Table](#)

You can filter the Current Top Sessions according to the Wait Classes, by clicking **Choose Resource**. Then select a resource from the list of the Wait Classes for events that occurred in the time period. This allows you to view only those processes for which that specific Wait Class occurred. The possible wait classes are:

- Workload
- waiting for a disk write to complete
- waiting for a disk read to complete
- waiting to be scheduled
- waiting for memory or a buffer
- waiting to output to the network

- waiting for internal system event

Top Sessions Table

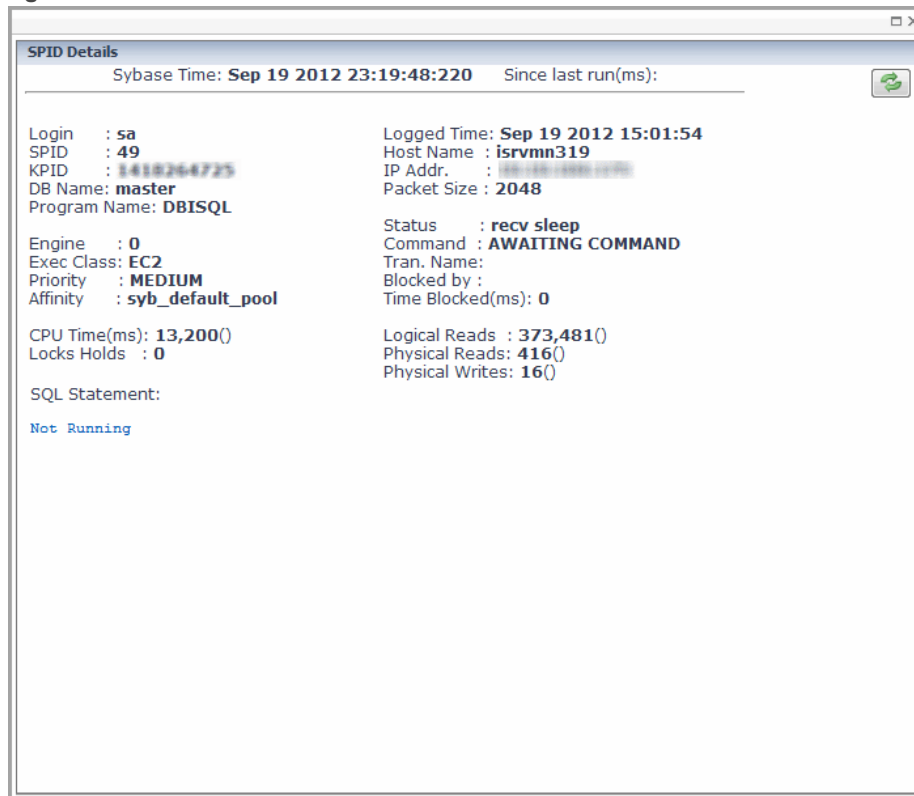
The Top Sessions table lists all of the top sessions.

The filters appearing just above the table allows you to look for specific entry based on some known criteria. You can filter on:

- System or user processes
- Time range, by selecting **Since Logon**, **Last Sample**, or **Last 10 min**

Click the Login Name for a session to view a popup that provides real-time monitoring details about the associated process.

Figure 18. SPID Details



These details include the time (in milliseconds) that has elapsed since the process was last run. Values for other statistics since the process was last run are provided in brackets.

Select an option button on the Top Sessions table to specify the SPID (session process identifier) number.

Table 22. Top Sessions Table Descriptions

Name	Description
Option button	Shows additional information about the selected SPID in the Top SQL of Session with SPID Table and Wait Events of SPID xx for Period Table .
Login Name	The login name used to open the Sybase session.
SPID	The Session Process Identifier (SPID) number.
Block by	If there is a block, this column shows the blocking process ID in red.
Database Name	The name of the database for which execution statistics are being monitored.
Application	The name of the application for which execution statistics are being monitored.

Table 22. Top Sessions Table Descriptions

Name	Description
CPU Time since Logon (ms)	The CPU time (in milliseconds) used by the session since the user logon.
Wait Time (ms) Since Logon	The wait time (in milliseconds) since the logon.
Logical Reads Since Logon	The number of buffers read from cache since the logon.
Physical Reads Since Logon	The number of buffers read from disk since the logon.
Physical Writes Since Logon	The number of buffers written to disk since the logon.
Current SQL	The SQL statement currently being executed by this session.
Command	The command code for the last SQL statement executed by this session.
Table Pages Accessed Since Logon	The number of table pages accessed since the logon.
Index Pages Accessed Since Logon	The number of index pages accessed since the logon.
Locks Currently Held	The number of locks the process currently holds.
TempDB pages Currently Used	The number of temporary database pages the process is currently using.
Commits Since Logon	The number of database commits since the logon.
Rollbacks Since Logon	The number of database rollbacks since the logon.
Engine	The number of the engine on which the process is running.

Top SQL of Session with SPID Table

The Top SQL of Session with SPID (session process identifier) table has the same columns as the Top SQL tab, [Top SQL Table](#).

This table provides detail for the top SQL session selected in the [Top Sessions Table](#).

Wait Events of SPID xx for Period Table

This table provides detail for SPID xx for the time period. The SPID number is selected in the [Top Sessions Table](#).

Table 23. Wait Events of SPID Description

Name	Description
Event ID	The event identifier number.
Event Description	The description of the Sybase wait event.
Wait Time (ms)	The wait time for this event in milliseconds.

SQL Details Tab

The SQL statement details are provided in a table.

You can display the SQL details for specified SQL text using the Filter on SQL Text box and clicking **Apply**. You can view the full list by clicking **Show All**.

Figure 19. SQL Details Tab

EndTime	Database Name	SQL Text	SPID	Execution Duration (ms)	CPU Time (ms)	Wait Time (ms)	Physical Reads	Logical Reads	Pages Modified	Memory Used (KB)	Packets Sent	Packets Received
Mar 14 2011 9:45PM	sybsystemprocs	sybsystemprocs..sp_fgl_topsql;line-45	86	16	16	0	0	506	0	22	0	0
Mar 14 2011 9:45PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'master'	86	13	13	0	0	0	0	12	0	0
Mar 14 2011 9:45PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'sales'	86	13	13	0	0	0	0	12	1	0
Mar 14 2011 9:45PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'tempdb'	86	13	13	0	0	0	0	12	0	0
Mar 14 2011 9:45PM	sybsystemprocs	sybsystemprocs..sp_fgl_datacached_objects;line-32	86	13	13	0	0	61	0	26	1	0
Mar 14 2011 9:44PM	sybsystemprocs	sybsystemprocs..sp_fgl_topsql;line-45	86	16	16	0	0	504	0	22	0	0
Mar 14 2011 9:44PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'sales'	86	16	16	0	0	20	0	204	1	0
Mar 14 2011 9:44PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'sybsystemdb'	86	13	13	0	0	51	0	204	0	0
Mar 14 2011 9:44PM	sybsystemprocs	sybsystemprocs..sp_fgl_datacached_objects;line-32	86	13	13	0	0	61	0	26	1	0
Mar 14 2011 9:39PM	sybsystemprocs	sybsystemprocs..sp_fgl_topsql;line-45	86	16	16	0	0	508	0	22	0	0
Mar 14 2011 9:39PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'sales'	86	16	16	0	0	0	0	12	1	0
Mar 14 2011 9:39PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'master'	86	13	13	0	0	0	0	12	0	0
Mar 14 2011 9:39PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'tempdb'	86	13	13	0	0	0	0	12	0	0
Mar 14 2011 9:39PM	sybsystemprocs	sybsystemprocs..sp_fgl_datacached_objects;line-32	86	13	13	0	0	61	0	26	1	0
Mar 14 2011 9:29PM	sybsystemprocs	exec sybsystemprocs..sp_fgl_dbseg_space 'sales'	86	19	19	0	0	51	0	204	1	0
Mar 14 2011 9:29PM	sybsystemprocs	sybsystemprocs..sp_fgl_topsql;line-45	86	16	16	0	0	496	0	22	0	0

SQL Details Table

Table 24. SQL Details Table Description

Name	Description
End Time	The time that the SQL statement finished running, in the format: <ul style="list-style-type: none"> Month, Day, Year Hour, Minute, am/pm.
Database	The name of the database that this SQL statement used during its execution.
SQL Text	The SQL statement that was run.
SPID	The ID of the session process that executed the SQL statement.
Execution Duration (ms)	The time it took to execute the SQL statement, in milliseconds.
CPU Time (ms)	The time it took the CPU to execute the SQL statement, in milliseconds.
Wait Time (ms)	The number of milliseconds a task has waited during execution of the SQL statement.
Physical Reads	The number of buffers read from disk.
Logical Reads	The number of buffers read from cache.
Pages Modified	The number of pages modified by the SQL statement.
Memory Used (KB)	The amount of memory (in kilobytes) allocated to the SQL script.
Packets Sent	The number of network packets sent by the Adaptive Server.
Packets Received	The number of packets received by the Adaptive Server.

Blocked Tree Tab

The Blocked tree tab shows the:

- Current blocked process situation.
- Blocked processes that occurred in the Time Range.
- Currently blocked processes, marked by lock.

The blocked tree information is provided in the [Blocked Tree Table](#). You can select the time for the metrics displayed:

- Show Last Sample; shows only the current blocks.
- Show All Time Range.

Figure 20. Blocked Tree Tab

The screenshot shows the 'Blocked Tree' tab in the ISRVMS318 application. The title bar indicates the time range: 'Blocked Tree (Oct 14 2009 13:24:14 - Oct 14 2009 14:24:14 - 60 minutes)'. Below the title bar, there's a 'Show Last Sample' button. The main table displays the following data:

SPID	kpid	Login Name	Database Name	Application	Locked Table	SQL Text	Time Blocked (hh:mm:ss)	Status	Lock Ended Time
130		sa	sales	SQL_Advantage					
149	306446830	sa	sales	MultiCon_1	leo1	begin tran update leo1 set c = "b" where i < 50000 ;	00:01:53	lock sleep	
150	306512367	sa	sales	MultiCon_3	leo1	begin tran update leo1 set c = "b" where i < 50000 ;	00:02:00	lock sleep	
148	306381293	sa	sales	MultiCon_2	leo1	begin tran update leo1 set c = "b" where i < 50000 ;	00:01:56	lock sleep	
21		sa	sales	SQL_Advantage					
119	304546257	sa	sales	MultiCon_2	leo1	begin tran update leo1 set c = "b" where i < 50000 ;	00:01:19	lock sleep	2009/10/14 14:16:41
120	304611794	sa	sales	MultiCon_3	leo1	begin tran update leo1 set c = "b" where i < 50000 ;	00:01:24	lock sleep	2009/10/14 14:16:41
118	304480720	sa	sales	MultiCon_1	leo1	begin tran update leo1 set c = "b" where i < 50000 ;	00:01:51	lock sleep	2009/10/14 14:17:11

Blocked Tree Table

Table 25. Blocked Tree Table Description


Name	Description
SPID	The session process identifier. A lock icon appears when there is an active current block.
kpid	The kernel process identifier.
Login Name	The login name of the blocked user.
Database Name	The name of the database being used by the current process.
Application	The application name of the current process.
Locked Table	The name of the table holding the locked object that causes the block.
SQL Text	The SQL (Structured Query Language) statement.
Time Blocked (hh:mm:ss)	The amount of time (in hours:minutes:seconds) the process has been blocked.
Status	The status of the blocked process.
Last Sample Time	The time at which the last sample block occurred, and was subsequently cleared.

System Wait Events Tab

The System Wait Events tab information is provided in a chart and two tables:

- [System Wait Events Chart](#)
- [System Wait Classes Table](#)
- [Wait Events Details Table](#)

Figure 21. System Wait Events Tab



System Wait Events Chart

Table 26. System Wait Events Chart Description

Name	Description
System Wait Classes	<div>The following wait classes, plotted over time:<ul style="list-style-type: none">• Process is running• waiting to be scheduled• disk read• disk write• log semaphore• lock• memory or a buffer• network (input)• network (output)• system event• another thread</div>

System Wait Classes Table

You can choose to display the system wait classes for either:

- Last Sample
- Period

The columns are common for both the Show Last Sample and Show Time Range tables.

You can view detailed information for a wait class by clicking a row in the System Wait Classes table. The detailed information appears in the [Wait Events Details Table](#), located at the bottom-right of the screen at the next data refresh.

Table 27. System Wait Events Table Description

Name	Description
Class ID	A sequential number for the class identifier.
Class Description	The description of the class: <ul style="list-style-type: none"> • Process is running • waiting to be scheduled • waiting for a disk read to complete • waiting for a disk write to complete • waiting to acquire the log semaphore
% of Total Wait Time	The percentage total wait time for this class.
Number of Waits	The number of waits for this class.
Wait Time (sec)	The wait time for this class in seconds.

Wait Events Details Table

The content of this table is determined by the selection made in the [System Wait Classes Table](#).

Table 28. Wait Events Details Table Description

Name	Description
Class ID	A sequential number for the class identifier.
Event ID	The event identifier number.
Description	The description of the wait event; for example: <ul style="list-style-type: none"> • waiting to run queue after sleep • waiting to run queue after yield
Number of Waits	The number of waits for this event.
Wait Time (sec)	The wait time for this event in seconds.

Spinlocks Tab

ASE uses spinlocks as synchronization mechanisms in SMP environment in order to protect shared resources. Use this dashboard to find areas where spinlock are a contention and tune their configuration. Data is available for Sybase version ASE 15.7 ESD#2 and above.

To monitor spinlocks, enable the 'Enable spinlock monitoring' Sybase configuration parameters using the following command:

```
sp_configure "enable spinlock monitoring", 1
```

The dashboard shows the following type of objects in which ASE applies the spinlock mechanism:

- Object Manager
- Object
- Index
- Index Hash
- Partition Hash
- Lock Address Manager
- Lock Manager
- Procedure Cache
- Procedure Cache Manager

For information about Data Caches Spinlock Contention, refer to the [Buffer Pools Dashboard](#) on page 42.

Description

Table 29. Spinlocks Tab Description

Name	Description
Name	Name of spinlock
Contention	Spinlock contention as percentage of waits in grabs.
Grabs	Number of grabs of this spinlock
Spins	Number of spins of this spinlock
Waits	Number of waits of this spinlock
Spins per wait	Number of spins per wait.

Execution Time Tab

This dashboard shows the execution time of each of the following operations performed in an instance:

- Compilation
- DeviceIO
- NetworkIO
- Execution
- Sorting.

Use this dashboard to understand where does the instance spends most of its time and to analyze trends in its behavior. Data is available for Sybase version ASE 15.7 SP100 and above.

Table 30. Execution Time Tab Description

Name	Description
Module	Name of operation category.
Execution Time	Execution time, in microseconds, of each operation performed.
Execution Count	Total number of this operation type that took place.

Engine Activity Dashboard

The Engine Activity dashboard provides performance details for Engine operations. This dashboard contains the following:

- [Engine Details Tab](#)
- [Engines Processes Tab](#)
- [Engines Host Tab](#)
- [Engines Compare Tab](#)


Engine Details Tab

Engine details information is provided in a table and a chart:

- [Engine Details Table](#)

- Engine Summary Charts

Figure 22. Engine Details Tab



Engine Details Table

You can filter the Engine Details table to include:

- Engines
- Engine Groups / Thread Pools
- Engines and Groups / Thread Pools

IMPORTANT: Thread pool information is only available from Sybase ASE 15.7 ESD2 and later.

Table 31. Engine Details Table Description

Name	Description
Engine	
No.	The engine number.
Current Status	The current status of the engine: online or offline. Hovering over this value provides more information.
<div> <div>Engine 0</div> <div>Start Time: Mar 13 2011 10:39AM</div> <div>Stop Time :</div> </div>	
Last Sample	
Eng. Busy %	The percentage busy for this engine.
CPU Busy %	The percentage CPU time all engines spent not idle.
I/O Busy %	The percentage of input/output resources currently being used by the database engine.
Cont. Switches %	The percentage of the total number of context switches for the engine divided by those completed for all engines.
Context Switches rate (c/s)	The count of context switches for the engine occurring per second.
Disk IO Checks rate	The number of times the engine checked for asynchronous disk I/O.
Checks Returning I/O %	The percentage of disk I/O checks with I/O pending.

Table 31. Engine Details Table Description

Name	Description
Async IO Completed	The number of times the engine completed I/O.
Processes Affinitied	The number of processes affined with the engine.
Time Range Period	
Total Busy	The percentage busy for all of the database engines.
Avg Busy %	The average percentage busy for all of the database engines.
Max Busy %	The maximum percentage busy for all of the database engines.
Avg I/O Busy %	The average percentage of input/output resources currently being used by all of the database engines.
Avg CPU Busy %	The average percentage CPU time all database engines spent not idle.
Cont. Switches %	The percentage of the total number of context switches for one engine divided by those completed for all engines.

Engine Summary Charts

Name	Description
Engine Summary	The following three values, plotted over time: <ul style="list-style-type: none">• Average Engine Total Busy• Average Engine I/O Busy• Average Engine CPU Busy
Engines Busy	Plots how busy an engine was over time. Select from the list of available engines.
Engine Groups / Thread Pools Busy ^a	Plots an average of how busy an engine group or thread pool was over time. Select from the list of available engine groups or thread pools.

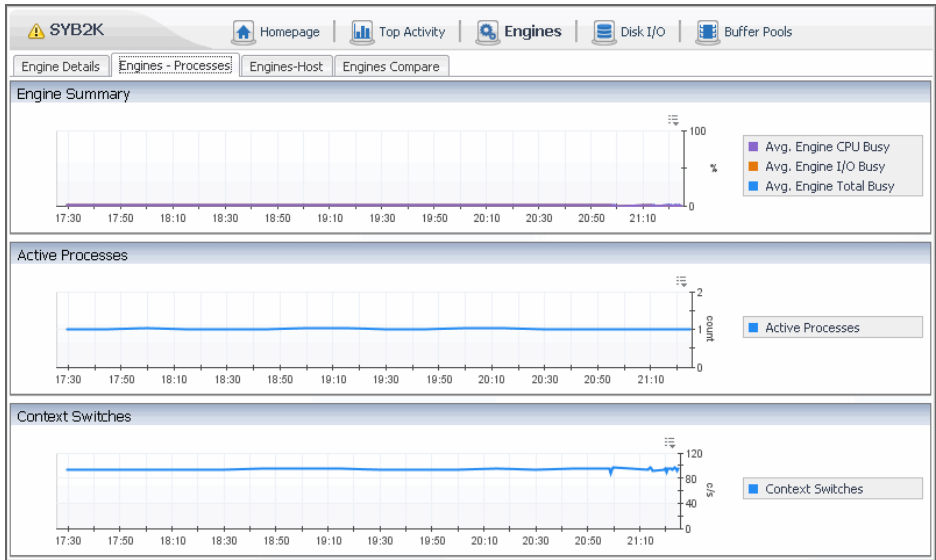
a. For thread pools, this information is only available from Sybase ASE 15.7 ESD2 and later.

To view a plot of the maximum value over the course of the time period, select **Show Maximum**.

Engines Processes Tab

The engine processes information is provided in three charts.

Figure 23. Engines Processes Tab



Engine Processes Charts

Table 32. Engine Processes Chart Description

Name	Description
Engine Summary	The following values, plotted over time: <ul style="list-style-type: none">• Average Engine Total Busy• Average Engine I/O Busy• Average Engine CPU Busy
Active Processes	The number of engine processes running, plotted over time.
Context Switches	The number of engine context switches, plotted over time.

Engines Host Tab

The engines host information is provided in two charts.

Engines Host Charts

Table 33. Engine Host Charts Descriptions

Name	Description
Engines Host	The following values, plotted over time, for each engine: <ul style="list-style-type: none">• Sybase Engines. The percentage of engine utilization.• Host CPUs. The percentage of utilized host CPU resources being used by the Sybase engines.
All Engines Activity	The percentage of activity for each engine.

Engines Compare Tab

The engines compare information is provided in three charts.

Engines Compare Charts

Table 34. Engines Compare Charts Description

Name	Description
Engines - time range	The following values, plotted over time: <ul style="list-style-type: none">• Average Engine Total Busy• Average Engine I/O Busy• Average Engine CPU Busy
Engines - month, day, year	This displays the plotted values for the previous day (that is, yesterday). The following values, plotted over time: <ul style="list-style-type: none">• Average Engine Total Busy• Average Engine I/O Busy• Average Engine CPU Busy
Engines - month, day, year	This displays the plotted values from today minus seven days (that is, a week ago) for a 24 hour period. The following values, plotted over time: <ul style="list-style-type: none">• Average Engine Total Busy• Average Engine I/O Busy• Average Engine CPU Busy

Disk I/O Dashboard

The Disk I/O dashboard provides performance details for read/write operations. This dashboard contains the following:

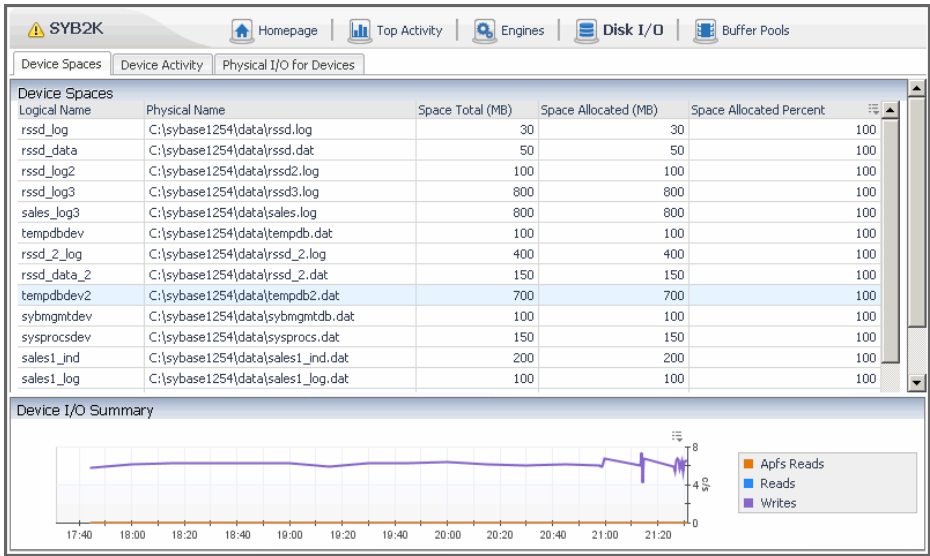
- [Device Spaces Tab](#)
- [Device Activity Tab](#)
- [Physical I/O for Devices Tab](#)

Device Spaces Tab

Devices spaces information is provided in a table and a chart:

- [Device Spaces Table](#)
- [Device I/O Summary Chart](#)

Figure 24. Device Spaces Tab



Device Spaces Table

Table 35. Device Spaces Table Description

Name	Description
Logical Name	The logical name of the device.
Physical Name	The physical name of the device with directory path.
Space Total (MB)	The total memory space available to this device, in megabytes.
Space Allocated (MB)	The total memory space allocated to this device, in megabytes.
Space Allocated Percent	The percentage of device space allocated in the device space available.

Device I/O Summary Chart

Table 36. Device I/O Summary Chart Description

Name	Description
Device I/O Summary	The following three values, plotted over time: <ul style="list-style-type: none">ReadsAPFs (Asynchronous Prefetch Activity) readsWrites

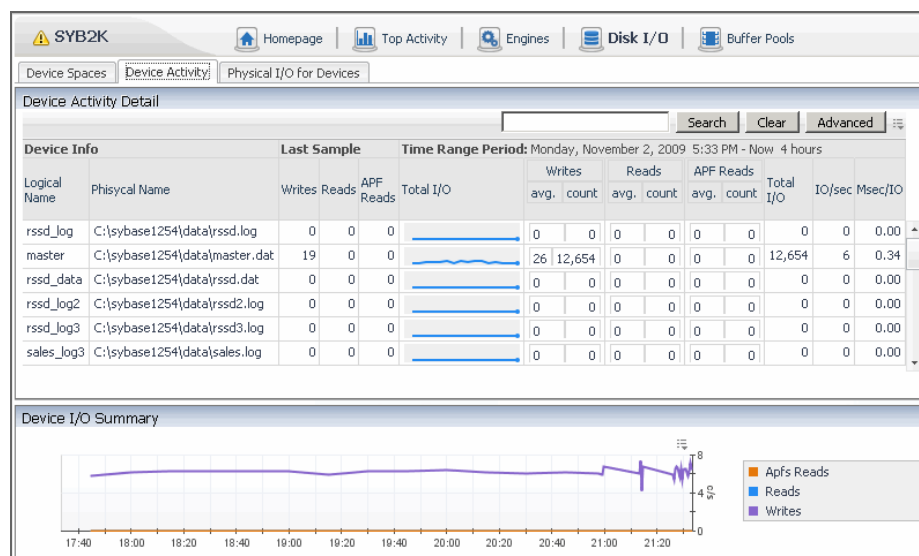
Device Activity Tab

Device activity information is provided in a table and a chart:

- [Device Activity Detail Table](#)
- [Device I/O Summary Chart](#)

Drill down by clicking on any device. A Sybase Disk I/O Detail chart appears.

Figure 25. Sybase Disk I/O Details chart



Device Activity Detail Table

Table 37. Device Activity Detail Table Description

Name	Description
Device Info	
Logical Name	The logical name of the device.
Physical Name	The physical name of the device with directory path.
Last Sample	
Writes	The number of writes for the selected time range.
Reads	The number of reads for the selected time range.
APF Reads	The number of APF (Asynchronous Prefetch Activity) reads for the selected time range.
Time Range Period:	
Total I/O	<p>The total number inputs and outputs for the selected time range.</p> <p>You can drill down by clicking on any field in this column. A popup appears showing the following three values, plotted over time:</p> <ul style="list-style-type: none"> Reads for this device. APFs (Asynchronous Prefetch Activity) reads for this device. Writes for this device.
Writes avg	The average number of writes for the selected time range.
Writes count	The number of writes for the selected time range.
Reads avg.	The average number of reads for the selected time range.
Reads count	The number of reads for the selected time range.
APF Reads avg.	The number of APF (Asynchronous Prefetch Activity) reads for the selected time range.
APF Reads count	The average number of APF (Asynchronous Prefetch Activity) reads for the selected time range.
IO/sec	The number of inputs and outputs per second.
Msec/IO	The number of milliseconds per input/output.


Device I/O Summary Chart

Table 38. Device I/O Summary Chart Description

Name	Description
Device I/O Summary	The following three values, plotted over time for all of the monitored devices: <ul style="list-style-type: none"> Reads APFs (Asynchronous Prefetch Activity) reads Writes

Physical I/O for Devices Tab

Figure 26. Physical I/O for Devices Tab



Physical I/O for Devices Charts

Table 39. Physical I/O for Devices Charts Description

Name	Description
Physical Writes (All Devices)	The number of physical writes for each device, plotted over time.
Physical Reads (All Devices)	The number of physical reads for each device, plotted over time.

Buffer Pools Dashboard

The Buffer Pools dashboard provides performance details for cache operations. This dashboard contains the following:

- [Data Cache Activity Tab](#)
- [Buffer Pools Activity Tab](#)
- [Procedure Cache Activity Tab](#)
- [Data Caches Spinlocks Tab](#)

Data Cache Activity Tab

Data Cache can be divided to Buffer Pools with different I/O sizes.

The cache objects shown in this table are determined by the values set for the following Sybase_MDA agent properties:

- Cached Objects - Max Rows Returned
- Cached Objects - Minimum Threshold Value
- Cached Objects - Sort Criteria

For more information, see [Setting Data Retrieval Properties](#) on page 81.


Data cache information is provided in a table and charts:

- [Data Cache Activity Table](#)

This table includes two [Data Cache Activity Buttons](#), which provide a simple way to find peaks times.

- [Data Cache Activity Charts](#)

Figure 27. Data Cache Activity Charts



Data Cache Activity Buttons

Table 40. Data Cache Activity Buttons Description

Name	Description
Cache Hit Rate Periods	A popup table lists: <ul style="list-style-type: none">• Start Time• End Time• Data Cache Summary dcache_hit_rate (%): minimum, average, and maximum. The minimum value is plotted in the Data Cache Activity Charts, Data Cache Hit Rate chart.
Cache Hits - Misses	A popup table lists: <ul style="list-style-type: none">• Start Time• End Time• Data Cache Summary Cache Hits/s (c/s): minimum, average, and maximum.• Data Cache Summary Cache Misses/s (c/s): minimum, average, and maximum.

Data Cache Activity Table

You can drill down by clicking any table field. A popup appears providing further drill downs to:

- **Cache Detail Graph**— Sybase Data Each Detail Graph popup appears showing Hit Rate Average (in percent) and Hit Rate Minimum (in percent) plotted over time.
- **Cached Objects Table**—a [Data Cache Objects Table Popup](#) appears.

Table 41. Data Cache Activity Table Description

Name	Description
Cache Name	The name of the cache.
Last Sample	
Hit Rate %	The percentage of times a requested procedure was found and therefore not read into the named cache.
Cache Hits (buffers) per sec	The number of times per second a requested page was found and therefore not read into the named cache.
Cache Hits (buffers) count	The number of buffers retrieved (that is, the number of Logical Reads) from the cache.
Cache Misses (buffers) per sec	The number of times per second a requested page had to be read into the named cache because it was not found.
Cache Misses (buffers) count	The number of buffers read (that is, the number of Physical Reads) into the cache from disk.
“Dirty” Buffer Retrievals	The number of dirty buffer retrievals.
Time Range Period:	
Hit Rate	<p>The number of times per second a requested page was found and therefore not read into the named cache for the selected time range.</p> <p>You can drill down by clicking on any field in this column. A popup appears showing the following three values, plotted over time:</p> <ul style="list-style-type: none">• Reads for this device.• APFs (Asynchronous Prefetch Activity) reads for this device.• Writes for this device.
Avg. Hit Rate %	The average percentage of times a requested page was found and therefore not read into the named cache for the selected time range.

Table 41. Data Cache Activity Table Description

Name	Description
Avg. Cache Hits (buffers) per sec	The average number of cache hits per second for the selected time range.
Avg. Cache Hits (buffers) count	The average number of cache hits for the selected time range.
Avg. Cache Misses (buffers) per sec	The average number of cache misses per second for the selected time range.
Avg. Cache Misses (buffers) count	The average number of cache misses for the selected time range.

Data Cache Objects Table Popup

Table 42. Data Cache Objects Table Popup Description

Name	Description
Cache Name	The name of the cache.
Name	The name of the data cache object.
Index ID	The identifier of the index. If this value is: <ul style="list-style-type: none"> zero, then the table is cached greater than zero, then the index is cached.
DB Name	The database name.
Cached (MB)	The amount of cached memory in megabytes.
Total Size (MB)	The total size of the cached object in the database in megabytes. NOTE: This field only shows data for Sybase version 15 and later.
Percent Cached	The percentage of the total object's size that is cached memory. NOTE: This field only shows data for Sybase version 15 and later.
Process Accessing	The number of processes currently accessing the object. NOTE: This field only shows data for Sybase version 15 and later.

Data Cache Activity Charts

Table 43. Data Cache Activity Charts Description


Name	Description
Data Cache Summary	The following three values are the total for all of the monitored caches, plotted over time: <ul style="list-style-type: none"> Number of hits per second. Number of misses per second. Number of dirty buffer retrievals per second.
Data Cache Hit Rate	The following two values for the monitored caches, plotted over time: <ul style="list-style-type: none"> Cache Hit Average. Cache Hit Minimum.

Buffer Pools Activity Tab

Buffer pools information is provided in a table and charts:

- [Buffer Pools Spaces Table](#)
- [Buffer Pools Charts](#)

Figure 28. Buffer Pools Activity Tab



Buffer Pools Spaces Table

Table 44. Buffer Pools Spaces Table Description

Name	Description
Cache Name/Pool	The name of the cache being monitored.
Pool Size (MB)	The number of megabytes that have been allocated for the pool.
Last Sample	
Physical Reads incl. APF (buffers)	The number of buffers, including APF (Asynchronous Prefetch Activity) that have been read from disk into the pool.
Physical Reads (Pages)	The number of buffer pages that have been read from disk into the pool.
Pages Touched	The number of pages that are currently being used within the pool.
Percent Used (%)	Data space used as a percentage of total data space available.
Buffers Grabbed Dirty	The number of dirty buffer retrievals.
Time Range Period (Average)	
Physical Reads incl. APF (buffers)	The number of buffers, including APF (Asynchronous Prefetch Activity) that have been read from disk into the pool for this time period.
Buffers Grabbed Dirty	The number of dirty buffer retrievals for this time period.
Pages to LRU	The number of buffers that were fetched and replaced in the least recently used portion of the pool.
Pages to MRU	The number of buffers that were fetched and replaced in the most recently used portion of the pool.

Buffer Pools Charts

Table 45. Buffer Pools Charts Description

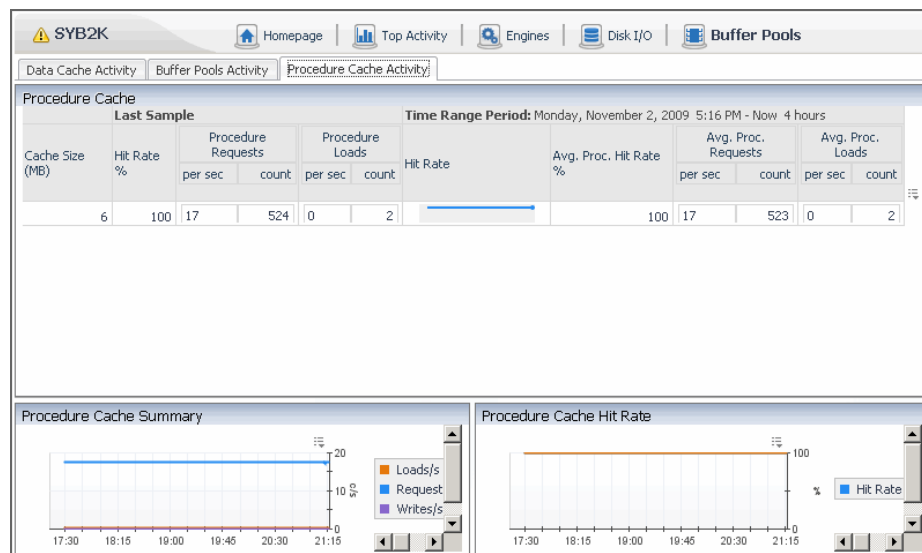
Name	Description
Data Cache Summary	The following three values, plotted over time: <ul style="list-style-type: none"> • Dirty Buffers per second • Hits per second • Miss per second
Data Cache Hit Rate	The Hit Rate Average plotted over time:

Procedure Cache Activity Tab

Procedure cache activity information is provided in a table and charts:

- [Procedures Cache Activity Table](#)
- [Procedure Cache Activity Charts](#)

Figure 29. Procedure Cache Activity Tab



Procedures Cache Activity Table

Table 46. Procedures Cache Activity Table Description

Name	Description
Cache Size (MB)	The size of the cache in megabytes.
Last Sample	
Hit Rate %	The percentage of times a requested page was found and therefore not read into the named cache.
Procedure Requests per sec	The number of procedure requests per second resulting in reads from disk.
Procedure Requests count	The number of procedure requests resulting in reads from disk.
Procedure Loads per sec	The number of procedures per second loaded into cache.
Procedure Loads count	The number of procedures loaded into cache.

Table 46. Procedures Cache Activity Table Description

Name	Description
Time Range Period:	
Hit Rate	<p>The number of times per second a requested page was found and therefore not read into the named cache for the selected time range.</p> <p>You can drill down by clicking on any field in this column. A popup appears showing the following three values, plotted over time:</p> <ul style="list-style-type: none"> • Reads for this device. • APFs (Asynchronous Prefetch Activity) reads for this device. • Writes for this device.
Avg. Proc. Hit Rate %	The average percentage of procedure requests resulting in reads from disk.
Avg. Proc. Requests per sec	The average number of procedure requests per second resulting in reads from disk.
Avg. Proc. Requests count	The average number of procedure requests resulting in reads from disk.
Avg. Proc. Loads per sec	The average number of procedures per second loaded into cache.
Avg. Proc. Loads count	The average number of procedures loaded into cache.

Procedure Cache Activity Charts

Table 47. Procedure Cache Activity Charts Description

Name	Description
Procedure Cache Summary	<p>The following three values are the total for all of the monitored procedure caches, plotted over time:</p> <ul style="list-style-type: none"> • Loads per second. • Requests per second. • Writes per second.
Procedure Cache Hit Rate	The Hit Rate Average plotted over time.

Data Caches Spinlocks Tab

This dashboard shows the Spinlocks that occur for each Data Cache. Use this dashboard for detecting data cache that may suffer a contention and monitor it over time.

The table columns are the same as in Spinlocks table in [Spinlocks Tab](#).

Database Spaces and Trends Dashboard

The Database Spaces and Trends dashboard shows details about databases and table spaces and the overall trends. This dashboard contains the following tabs:

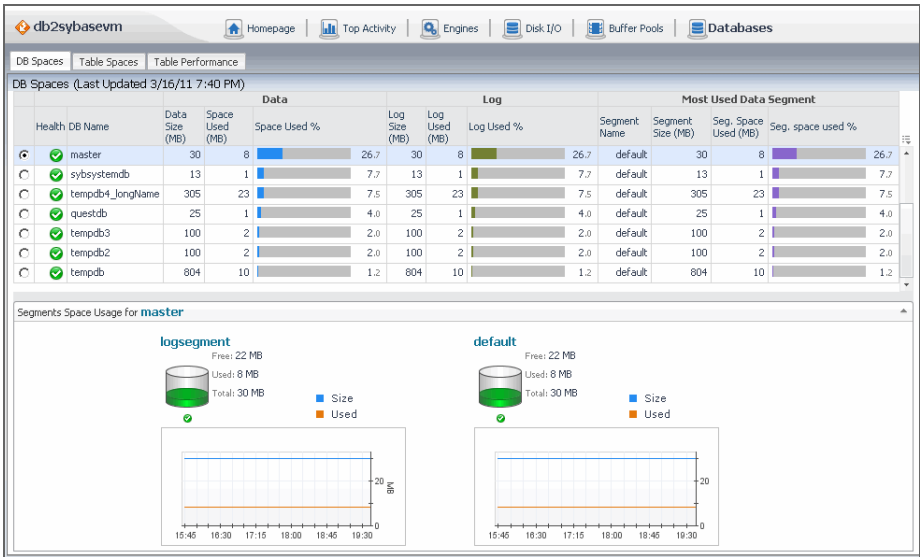
- [DB Spaces Tab](#)
- [Table Spaces Tab](#)
- [Table Performance Tab](#)
- [Top I/O Active Tables](#)
- [Top Lock Activity Tables](#)

DB Spaces Tab

Device activity information is provided in a table and a chart:

- [DB Spaces Table](#)
- [Segment Space Usage Composite View](#)

Figure 30. DB Spaces Tab



DB Spaces Table

Table 48. DB Spaces Tab Description

Name	Description
Option button	Shows additional information about the selected database in the Segment Space Usage Composite View .
Data	
Health	The state of the database health, indicated by the alarm of the highest severity generated against the database. Hovering over this column results in a message box showing the alarm message. <div><div>sales - Warning</div><div>"data222" Segment Space Percent reached 85 percent.</div></div>
Data Size (MB)	The total amount of space in MB that is available to the database.
Space Used (MB)	The amount of space used by the database.
Space Used %	The percentage of the database space currently in use, shown as a number and a graphic indicator

Table 48. DB Spaces Tab Description

Name	Description
Log	
Log Size (MB)	The total amount of space in MB that is available to the database log.
Log Used (MB)	The amount of space used by the database log.
Log Used %	The percentage of the database log currently in use, shown as a number and a graphic indicator.
Most Used Data Segment	
Segment Name	The name of the database segment that is most used/
Segment Size (MB)	The total amount of space in MB that is available to the database segment.
Seg. Space Used (MB)	The amount of space used by the database segment.
Seg. space used %	The percentage of the database segment currently in use, shown as a number and a graphic indicator.

Segment Space Usage Composite View

This composite view contains information about space usage for each segment of the database selected in the [DB Spaces Table](#).

Table 49. Segment Space Usage Composite View Description

Name	Description
Current segment usage	
Free	The amount of the segment space that is currently not in use.
Used	The amount of the segment space that is currently in use.
Total	The total amount of space that is available to the database segment.
Segment usage over time (chart)	
Size	The total amount of space that is available to the database segment, shown over time.
Used	The amount of the segment space that is currently in use, shown over time.

Table Spaces Tab

The Table Spaces information is provided in a table.

To enable the data collection for a database and to display this data on the Table Spaces tab, the database must be added to the Table Space Monitoring list in the **Table Space Management** group of the Sybase_MDA Agent properties. For more information, see [Setting Table Space Management Properties](#) on page 80.

Figure 31. Table Spaces Tab

db2sysbasevnm

Homepage

Table Activity

Engines

Disk I/O

Buffer Pools

Databases

DB SpacesTable SpacesTable Performance

Filter byTable Name:DB Name:AllShow:AllTablesIndexesClear

Tables and Partitions (sample rate: 1 min)

Last Updated 5/8/11 6:50 PM

Table Identification							Statistics					
Table Name	DB Name	Owner	Index Name	Ind ID	Locks Scheme/Partition type	Table Size (MB)	Data/Index Size (MB)	Incremental Delta Size (MB)	Rows Inserted (1.5 Hours)	Rows Updated (1.5 Hours)	Rows Deleted (1.5 Hours)	Rows
mymsgsg_HASH_PARTITIONS...	sales	dbo	NCIND222_mysmsg...	3	datarows	923.0	43.2		6,000			6,203,192
NCIND333_mysmsgsg_HASH_P...					hash		15.1		6,000			
NCIND333_mysmsgsg_HASH_P...					hash		14.9					
NCIND333_mysmsgsg_HASH_P...					hash		3.4					
NCIND333_mysmsgsg_HASH_P...					hash		9.7					
mymsgsg_HASH_PARTITIONS...	sales	dbo		0	datarows	923.0	729.6		6,000	5,000		6,203,192
p1					hash		183.5		6,000	5,000		528,106
p4					hash		261.8					2,522,108
p3					hash		214.8					1,632,783
p2					hash		69.6					570,595
mymsgsg_HASH_PARTITIONS...	sales	dbo	mymsgsg_HASH_PA...	2	datarows	923.0	150.3		5,000	5,000		6,203,192
mymsgsg_HASH_PARTITIONS...					hash		38.8		5,000	5,000		
mymsgsg_HASH_PARTITIONS...					hash		28.3					
mymsgsg_HASH_PARTITIONS...					hash		36.0					
mymsgsg_HASH_PARTITIONS...					hash		47.2					
leo1	sales	dbo		0	datapages	339.0	171.2					7,932,795
leo1	sales	dbo	leo1_c	3	datapages	339.0	50.2					7,932,795
leo1	sales	dbo	leo1_j	2	datapages	339.0	117.6					7,932,795

Table Spaces Filters

The filters appearing just above the table allows you to look for specific entry based on some known criteria. You can filter on table and/or database name, and choose between showing the table and/or indices.

Table Spaces Table

Table 50. Table Space Table Description

Name	Description
Last Updated	
Table Identification	
Table Name	For table entries, this column contains the table name. If the table has one or more partition, the table name can be expanded to show the names of table partitions in this column. For index entries, this column contains the name of the table associated with the index. When the row shows information about a table, clicking this column shows additional information in the Table Usage Drilldown . For index entries, clicking this column shows the Index Usage Drilldown . Clicking a partition the Partition-Table Usage Drilldown .
DB Name	The name of the database to which the table belongs.
Owner	The table owner.
Index Name	The name of the table index.
Ind ID	For index entries, this column contains the ID of the table index. For table entries, this column shows zero '0'. For partitions, this column is blank.
Locks Schema/Partition type	For table entries, this column indicates its locking schema: allpages, datapages, or datarows. For partition entries, this column indicates its type: range, hash, list, or roundrobin.

Table 50. Table Space Table Description

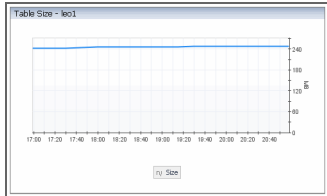
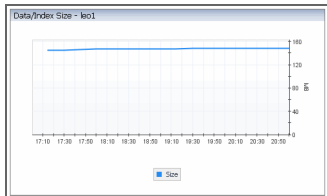

Name	Description
Table Size (MB)	<p>The table size in MB, including the amount of space occupied by the table data and index.</p> <p>Clicking or hovering over this column shows a chart indicating the table size over time.</p> 
Statistics	
Data/Index Size (MB)	<p>For index entries, this column shows the amount of space occupied by the table index.</p> <p>For table entries, this column shows the amount of space occupied by the table data.</p> <p>Clicking or hovering over this column shows a chart indicating the displayed data or index size over time.</p> 
Incremental Delta Size (MB) (1.5 Hours)	The difference in size of the table data and index over the past one and a half hour.
Rows Inserted (1.5 Hours)	The number of rows inserted over the last one and a half hour.
Rows Updated (1.5 Hours)	The number of rows updated over the past one and a half hour.
Rows Deleted (1.5 Hours)	The number of rows deleted over the past one and a half hour.
Rows	<p>For index entries, this column shows zero '0'.</p> <p>For table entries, this column shows the number of rows in the table. Clicking or hovering over this column shows a chart indicating the displayed number of table rows over time.</p> 

Table Usage Drilldown

The Table Usage drilldown displays a dialog box with detailed information about the selected table.

Figure 32. Table Usage Drilldown details



Table 51. Table Usage Drilldown Descriptio

Name	Description
Table	The table name, its locking schema, and the number of partitions.
Data row count	The number of data rows.
Table Size (MB)	The total table size in MB.
Data Size (MB)	The amount of table space occupied by the data.
Data row size	The data row size.
Forwarded row count	The number of rows forwarded to other pages.
Deleted row count	The number of deleted rows.
Derived statistics	<ul style="list-style-type: none"> Data page cluster ratio: The data page cluster ratio, used to measure the packing and sequencing of pages on extents. Space utilization: The number of pages compared to the expected minimum number of pages, calculated using the average row size and the number of data rows. Large I/O efficiency: The number of useful pages brought in by a large I/O. <p>NOTE: This is the same information that can be obtained with the Sybase <code>optdiag</code> utility. For more information about this command, see your Sybase documentation.</p>
Charts	
Table Rows Total for Today	The count of table rows, shown over the course of the current day.
Table Size for Today	The table size in MB, shown over the course of the current day.
Data Modification Each 1.5 Hours for Today	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current day.
Table Rows Total for Week Days	The number of table rows, shown over the course of the current week.
Table Size for Week Days	The table size in MB, shown over the course of the current week.
Data Modification Each 1.5 Hours for Week Days	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current week.

Index Usage Drilldown

The Index Usage drilldown displays a dialog box with detailed information about the selected index.

Figure 33. Index Usage Drilldown



Table 52. Index Usage Drilldown Description

Name	Description
Table	The name of the table associated with the index, the table's locking schema, and the number of partitions.
Index	The index name.
Index type	The index type.
Table row count	The number of table rows.
Index Size (MB)	The index size in MB.
Leaf row size	The leaf row size.
Derived statistics	<ul style="list-style-type: none"> Data page cluster ratio: The data page cluster ratio, used to measure the packing and sequencing of pages on extents. Index page cluster ratio: The data page cluster ratio, used to measure the packing and sequencing of index leaf pages on extents. Space utilization: The number of pages compared to the expected minimum number of pages, calculated using the average row size and the number of data rows. Large I/O efficiency: The number of useful pages brought in by a large I/O. <p>NOTE: This is the same information that can be obtained with the Sybase <code>optdiag</code> utility. For more information about this command, see your Sybase documentation.</p>

Charts

Table Rows Total for Today	The count of table rows, shown over the course of the current day.
Index Size for Today	The index size in MB, shown over the course of the current day.
Data Modification Each 1.5 Hours for Today	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current day.
Table Rows Total for Week Days	The number of table rows, shown over the course of the current week.
Index Size for Week Days	The index size in MB, shown over the course of the current week.
Data Modification Each 1.5 Hours for Week Days	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current week.

Partition-Table Usage Drilldown

The Partition-Table Usage drilldown displays a dialog box with detailed information about the selected partition.

Figure 34. Partition-Table Usage Drilldown




Table 53. Partition-Table Usage Drilldown Description

Name	Description
Partition Name	The partition name.
Table	The name of the table to which the selected partition belongs, the table's locking schema, the number of partitions it contains and their type.
Partition row count	The number of rows in the partition.
Table Size (MB)	The table size in MB.
Partition Size (MB)	The partition size in MB.
Data row size	The data row size.
Forwarded row count	The number of rows forwarded to other pages.
Deleted row count	The number of deleted rows.
Derived statistics	<ul style="list-style-type: none"> Data page cluster ratio: The data page cluster ratio, used to measure the packing and sequencing of pages on extents. Space utilization: The number of pages compared to the expected minimum number of pages, calculated using the average row size and the number of data rows. Large I/O efficiency: The number of useful pages brought in by a large I/O. <p>NOTE: This is the same information that can be obtained with the Sybase <code>optdiag</code> utility. For more information about this command, see your Sybase documentation.</p>

Charts

Partition Rows Total for Today	The count of partition rows, shown over the course of the current day.
Partition Size for Today	The partition size in MB, shown over the course of the current day.
Partition Size Delta/Hour for Today	The change in partition size in MB, shown over the course of the current day.
Partition Rows Total for Week Days	The number of partition rows, shown over the course of the current week.
Partition Size for Week Days	The partition size in MB, shown over the course of the current week.
Size Delta/Hour for Week Days	The change in partition size in MB, shown over the course of the current week.

Table Performance Tab

The Table Performance information is provided in a table.

To enable the data collection for a database and to display this data on the Table Performance tab, the database must be added to the Table Space Monitoring list in the **Table Space Management** group of the Sybase_MDA Agent properties. For more information, see [Setting Table Space Management Properties](#) on page 80.

Figure 35. Table Performance Tab

Table Identification						Performance Statistics						
Table Name	Owner	DB Name	Index Name	Ind ID	Table Size (MB)	Logical Reads	APF Reads	Physical Reads	Physical Writes	Rows Inserted	Rows Updated	Rows Deleted
mymsgs_HASH_PARTITIONS...	dbo	sales	NCIND222_mymsgs...	3	902.2	405,003	0	193	3,803	100,600	0	0
NCIND033_mymsgs_HASH_P...						406,245	0	193	3,803	100,600	0	0
NCIND033_mymsgs_HASH_P...						0	0	0	0	0	0	0
NCIND033_mymsgs_HASH_P...						0	0	0	0	0	0	0
NCIND033_mymsgs_HASH_P...						0	0	0	0	0	0	0
mymsgs_HASH_PARTITIONS...	dbo	sales		0	902.2	444,001	617	507	5,830	100,600	0	0
p3						3,035	0	0	0	0	0	0
p2						3,035	0	0	0	0	0	0
mymsgs_HASH_PARTITIONS...	dbo	sales	mymsgs_HASH_PA...	2	902.2	635,808	0	11,878	29,304	0	0	0
mymsgs_HASH_PARTITIONS...						637,049	0	11,878	29,304	0	0	0
mymsgs_HASH_PARTITIONS...						0	0	0	0	0	0	0
mymsgs_HASH_PARTITIONS...						0	0	0	0	0	0	0
mymsgs_HASH_PARTITIONS...						0	0	0	0	0	0	0

Table Performance Filters

The filters appearing just above the table allow you to look for specific entry based on some known criteria. You can filter on table and/or database name, choose between showing the table and/or indices, show delta values or rates, and indicate if you want to display the data collected in the last collection period, most recent sample, or in the last ten minutes.

Table Performance Table

Table 54. Table Performance Table Description

Name	Description
Table Identification	
Table Name	For table entries, this column contains the table name. If the table has one or more partition, the table node can be expanded to show the names of table partitions in this column. For index entries, this column contains the name of the table associated with the index. When the row shows information about a table, clicking this column shows additional information in the Table Performance Drilldown . For index entries, clicking this column shows the Index Usage Drilldown . Clicking a partition the Partition-Table Usage Drilldown .
Owner	The table owner.
DB Name	The name of the database to which the table belongs.
Index Name	The name of the table index.
Ind ID	For index entries, this column contains the ID of the table index. For table entries, this column shows zero '0'. For partitions, this column is blank.
Table Size (MB)	The table size in MB.
Performance Statistics	
Logical Reads	The number of buffers read from cache.
APF Reads Delta	The change in the APF reads for the selected period.
Physical Reads Delta	The change in the number of buffers read from disk for the selected period.

Table 54. Table Performance Table Description

Name	Description
Physical Writes	The number of buffers written to disk.
Rows Inserted	The number of inserted rows.
Rows Updated	The number of updated rows.
Rows Deleted	The number of deleted rows.

Table Performance Drilldown

The Table Performance drilldown displays a dialog box with detailed information about the selected table.

Figure 36. Table Performance Drilldown



Table 55. Table Performance Drilldown Description

Name	Description
Table	The table name, its locking schema, and the number of partitions.
Data row count	The number of data rows.
Table Size (MB)	The total table size in MB.
Data Size (MB)	The amount of table space occupied by the data.
Data row size	The data row size.
Forwarded row count	The number of rows forwarded to other pages.
Deleted row count	The number of deleted rows.
Derived statistics	<ul style="list-style-type: none"> Data page cluster ratio: The data page cluster ratio, used to measure the packing and sequencing of pages on extents. Space utilization: The number of pages compared to the expected minimum number of pages, calculated using the average row size and the number of data rows. Large I/O efficiency: The number of useful pages brought in by a large I/O. <p>NOTE: This is the same information that can be obtained with the Sybase <code>optdiag</code> utility. For more information about this command, see your Sybase documentation.</p>

Charts

Logical Reads for Today (value)	The amount of data read from cache, shown over the course of the current day.
Physical I/O for Today (value)	The numbers of physical reads, writes, and APF reads, shown over the course of the current day.

Table 55. Table Performance Drilldown Description

Name	Description
Data Modification for Today (value)	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current day.
Logical Reads for Week Days (value)	The amount of data read from cache, shown over the course of the current week.
Physical I/O for Week Days (value)	The numbers of physical reads, writes, and APF reads, shown over the course of the current week.
Data Modification for Week Days (value)	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current week.
Table Rows Total for Today	The count of table rows, shown over the course of the current day.
Table Size for Today	The table size in MB, shown over the course of the current day.
Data Modification Each 1.5 Hours for Today	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current day.
Table Rows Total for Week Days	The number of table rows, shown over the course of the current week.
Table Size for Week Days	The table size in MB, shown over the course of the current week.
Data Modification Each 1.5 Hours for Week Days	Data modification, including the amounts of inserted, updated, and deleted data, shown over the course of the current week.

Index Usage Drilldown

This table displays the same layout and data as the [Index Usage Drilldown](#).

Partition-Table Usage Drilldown

This table displays the same layout and data as the [Partition-Table Usage Drilldown](#).

Top I/O Active Tables

Use this dashboard to examine activity of top tables. Tables are displayed based on the sort criteria defined in agent ASP data retrieval. The default order is "Logical Reads". Other criteria possible are:

- Lock Waits
- Logical Reads
- Physical Reads
- Row Count
- Rows Inserted

Top I/O Active Tables Description

Table 56. Top I/O Active Tables Description

Name	Description
Table Name	Name of the table.
DB Name	The name of the database to which the table belongs.
Index Name	The name of the table index.
Ind ID	For index entries, this column contains the ID of the table index. For table entries, this column shows zero '0'. For partitions, this column is blank.
Statistics	
Logical Reads	The number of buffers read from cache.
Physical Reads	The number of buffers read from disk.
APF Reads	The number of APF reads.
Physical Writes	The number of buffers written to disk.
Used Count	Number of times the object was used in a plan.
Last Used Date	Last date the index was used in a plan.

Top Lock Activity Tables

Use this tab to examine lock activity of top tables. Tables are displayed based on the sort criteria defined in agent ASP data retrieval. The default order is Logical Reads. Other criteria possible are

- Lock Waits
- Logical Reads
- Physical Reads
- Row Count
- Rows Inserted

Top Lock Activity Table Description

Table 57. Top Lock Activity Table Description

Name	Description
Table Name	Name of the table.
DB Name	The name of the database to which the table belongs.
Index Name	The name of the table index.
Ind ID	For index entries, this column contains the ID of the table index. For table entries, this column shows zero '0'. For partitions, this column is blank.
Row Count	The number of rows in the table.
Lock Scheme	Indicates the locking schema: allpages, datapages, or datarows.
Statistics	
Lock Requests	Number of physical locks requested.
Lock Waits	Number of times a task waited for a lock.
Lock Cont Pct.	The percentage of Lock Waits in Requests.
Rows Inserted	The number of rows inserted.
Rows Deleted	The number of rows deleted.
Rows Updated	The number of rows updated.

RS Home Page Dashboard


The Replication Server Home Page provides performance details and indicates problems occurring on a specific Replication Server. The following sub-sections describe the dashboard element metrics:

- [RS Instance Metrics](#)
- [RS Alarms](#)
- [RSSD Metrics](#)
- [RS Host Metrics](#)
- [Replication Server Metrics](#)

Drilldown dashboards may be accessed through some of the dashboard element metrics. Clicking **Rep Server Activity** opens the [Replication Server Activity Dashboard](#).

If no data is being collected by the Sybase_RS agent, the “Not Collecting data” error message appears beside the Sybase agent name. In this case, the dashboard shows the data collected from the last sample time.

Figure 37. RS Home Page Dashboard



RS Instance Metrics

Table 58. RS Instance Metrics

Name	Description
company name	The name of the company that provides the Replication Server tool.
Version	The Sybase replication server version being monitored.
Last Sample Time	The last date (month/day/year) and time (hour:minute, AM/PM) that data was sampled.
Sample Period	The frequency (in seconds) of data sampling.
Host	The name of the host computer for the replication server.
Logs	
Agent Log	<p>The location of the log listing the agent errors.</p> <p>You can drill down to the detailed log by clicking this link. The Agent Error Log appears, showing:</p> <ul style="list-style-type: none"> • Message Time • Severity • Message
Error Log	<p>The location of the log listing the replication server errors.</p> <p>You can drill down to the detailed log by clicking this link. The Replication Server Error Log appears, showing:</p> <ul style="list-style-type: none"> • Message Time • Level • Error Number • Message

RS Alarms

Table 59. RS Alarms

Name	Description
Agent Status	<p>When the agent instance is running, the following status message appears:</p> <ul style="list-style-type: none">Collecting Data <p>When the agent instance is running but not collecting data, one of the following status messages appear:</p> <ul style="list-style-type: none">StartingStoppedStoppingRunning but not collecting dataUnknown
Host	<p>The number of fatal, critical, and warning alarms for this host.</p> <p>You can drill down to the detailed alarm list.</p>
Agent	<p>The number of fatal, critical, and warning alarms for this database agent.</p> <p>You can drill down to the detailed alarm list.</p>

RSSD Metrics

Table 60. RSSD Metrics

Name	Description
Server	<p>The name of the Replication Server System Database (RSSD) server.</p> <p>You can drill down by clicking on the server name. The MDA Home Page Dashboard appears.</p>
DBName	<p>The name of the database.</p>
Sybase_MDA agent status	<p>The status of the Sybase_MDA agent:</p> <ul style="list-style-type: none">Available—the Sybase_MDA Agent and the Replication Server are running on the same Foglight Management Server. This also indicates that the Sybase_MDA Agent is in the Collecting Data state.No Agent—no Sybase_MDA Agent is running. <p>You can drill down by clicking on this status field. The MDA Home Page Dashboard appears for the appropriate server.</p>
Data Size	<p>The amount of data memory (in megabytes) in the RSSD:</p> <ul style="list-style-type: none">Available memoryTotal memory <p>This field displays color-coded real-time behavior for severity.</p>
Log Size	<p>The amount of log memory (in megabytes) in the RSSD:</p> <ul style="list-style-type: none">Available memoryTotal memory <p>This field displays color-coded real-time behavior for severity.</p>

RS Host Metrics

From the fields in this pane you can drill down by clicking on the device symbol. A popup appears showing Health, Alarms, CPU, Memory, Disk, and Network metrics.

From the CPU, Memory, Disk, or Network metrics you can:

- Hover to see the metrics plotted over time
- Click on the metric to view the detailed host dashboard.

You can drill down further by clicking:

- **Host Monitor.** The Host Monitor Dashboard appears.
- **Host Browser.** The Host Browser Dashboard appears.
- **Host Information.** A popup appears showing the host attributes for the selected database instance.

For detailed information on the host dashboards, refer to the *Managing Operating Systems User Guide*.

Table 61. RS Host Metrics

Name	Description
CPU Usage	The amount of CPU resources being consumed by the database host. This field displays color-coded real-time behavior for severity.
Memory Usage	The amount of memory being consumed by the database host. This field displays color-coded real-time behavior for severity.

Replication Server Metrics

There are multiple sections in the Replication Server pane:

- Each of the Primary Server sections shows [Primary Server - Replicated Server Status](#) information for a primary server.
- The bottom section shows [Partitions](#) information.

The pane title line indicates the overall status of the replication server, which is either:

- HEALTHY—all threads are executing as expected.
- SUSPECT—a thread is down and the Replication Server expected it to be running. Or, a thread is in a Connecting state. The Connecting state means that either the server to which Replicated Server is connecting is unavailable and a problem exists, or the Replicated Server connects successfully in a moment and the suspect status is transitory.

Primary Server - Replicated Server Status

The top sections show sets of related primary and replicated server information:

- [Primary Server Information](#)
- [Replicated Servers Table](#)

Click the **Pop Up Screen** button to see the full list of replicated server databases.

Primary Server Information

Table 62. Primary Server Information

Name	Description
Primary Server	<p>The primary server is shown as combination of the server name and the database name. Accordingly, the Data and Log sizes are indicated for the appropriate database inside the appropriate Sybase server.</p> <p>You can drill down by clicking on this field. A popup appears showing the Data Size and the Log Size.</p> <p>The amount of data memory and log memory (in megabytes) is shown for the primary server:</p> <ul style="list-style-type: none">• Available memory• Total memory <p>You can drill down further by clicking on the Data Size or Log Size device image:</p> <ul style="list-style-type: none">• The Data Space Used is shown plotted over time.• The Log Space Used is shown plotted over time. <p>The Data Size and Log Size fields displays color-coded real-time behavior for severity.</p>
Sybase_MDA agent status	<p>The status of the Sybase_MDA agent:</p> <ul style="list-style-type: none">• Available—the Sybase_MDA Agent and the Replication Server are running on the same Foglight Management Server. This also indicates that the Sybase_MDA Agent is in the Collecting Data state.• No Agent—no Sybase_MDA Agent is running. <p>You can drill down by clicking on this status field. The MDA Home Page Dashboard appears for the appropriate server.</p>
Threads Status	<p>The number of threads that are either:</p> <ul style="list-style-type: none">• Down—the thread has not started or has terminated.• Suspended—the thread has been suspended by the user. <p>These fields display color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on any Thread Status. A table appears showing the thread status details.</p>
Inbound Q Size	<p>The amount of memory being used for the inbound queue, in megabytes.</p> <p>These fields display color-coded real-time behavior for severity.</p> <p>You can drill down by clicking on the text or number. A popup appears showing the Inbound queue history, plotted over time.</p>

Replicated Servers Table

Table 63. Replicated Servers Table Description

Name	Description
Replicated Server DB	<p>The replicated server database name.</p> <p>You can drill down to view the Data Size and Log Size.</p> <p>You can drill down further to view a chart of the Data Space Used or Log Space Used plotted over time.</p>
Agent Status	<p>The status of the Sybase_MDA agent:</p> <ul style="list-style-type: none">• Available—the Sybase_MDA Agent and the Replication Server are running on the same Foglight Management Server. This also indicates that the Sybase_MDA Agent is in the Collecting Data state.• No Agent—no Sybase_MDA Agent is running. <p>You can drill down by clicking on this status field. The MDA Home Page Dashboard appears for the appropriate server.</p>

Table 63. Replicated Servers Table Description

Name	Description
Latency (sec)	The difference between the time the transaction was committed at the destination (Replicate) database, and the Time at the origin (Primary Site) when the transaction was committed, in seconds. You can drill down to view the Latency plotted over time.
SQM	The SQM (Stable Queue Manager) status. You can drill down by clicking on this field. A popup appears showing the Blocks Written Into the Queue per second, plotted over time.
Q. Size (MB)	The amount of data remaining in the queue for processing, in megabytes. You can drill down to view the Queue Size plotted over time.
DSI	The DSI (Data Server Interface) thread status. You can drill down by clicking on this field. A popup appears showing the Blocks Read from the Queue per second, plotted over time.

Partitions

Table 64. Partitions

Name	Description
Partitions	The total number of partitions. This field displays color-coded real-time behavior for severity. You can drill down by clicking on the device symbol. A popup appears showing the used space summary, plotted over time.
Three top partition names	The amount of data memory (in megabytes) in the partition: <ul style="list-style-type: none">• Available• Total This field displays color-coded real-time behavior for severity.

Replication Server Activity Dashboard

The Replication Server dashboard contains several views that show the performance metrics of the CPU of the machine hosting the Replication Server.

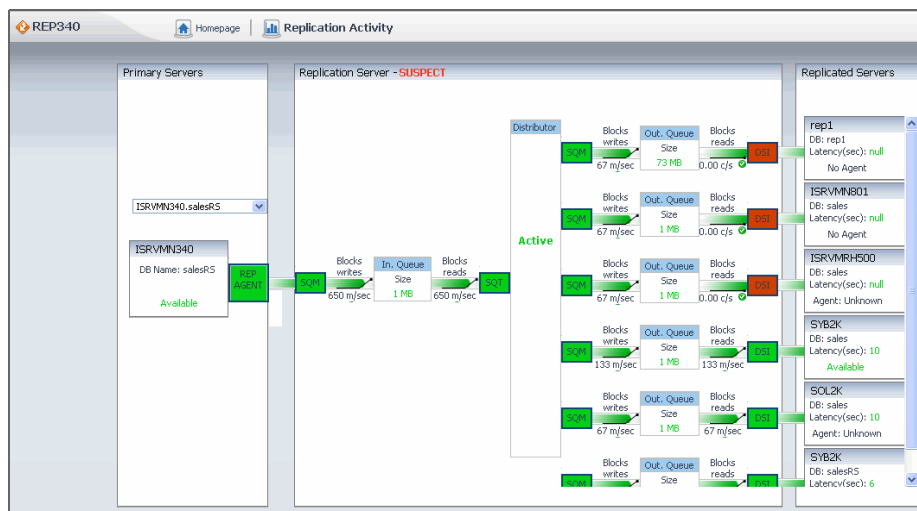
To display this dashboard, on the [Sybase MDA Global View Dashboard](#), drill down on a Sybase_RS agent, and on the [RS Home Page Dashboard](#) that appears, click **Rep Server Activity**.

The dashboard contains the following sections:

- [Primary Servers Metrics](#)
- [Replication Server Metrics](#)
- [Replicated Servers Metrics](#)

Drilldown dashboards may be accessed through some of the dashboard element metrics.

Figure 38. Replication Server Activity Dashboard



Primary Servers Metrics

Table 65. Primary Servers Metrics

Name	Description
Primary Server box	Lists all primary servers related to the replicated server. The primary server is shown as combination of the server name and the database name. Click the box to select a Primary Server and show its metrics in the Replication Server dashboard.
Primary Server view	Shows the database name and the status of the Sybase_MDA Agent. The status of the Sybase_MDA agent: <ul style="list-style-type: none"> Available—the Sybase_MDA Agent and the Replication Server are running on the same Foglight Management Server. This also indicates that the Sybase_MDA Agent is in the Collecting Data state. No Agent—no Sybase_MDA Agent is running. You can drill down by clicking on this status field. The MDA Home Page Dashboard appears for the appropriate server.
Rep. Agent	The RepAgent status. The Sybase RepAgent is used to send the changes from the primary database to the Replication Server.

Replication Server Metrics

This pane shows the data flow between a primary server and replicated servers through a replication server. The primary component of a replication server is the Distributor, which processes changes from the primary server and sends it on to replicated servers. The Distributor uses inbound and outbound queues to receive and send data and the related components, including the Stable Queue Transaction (SQT), SQM (Stable Queue Manager), and Data Server Interface (DSI). The following table lists the components that appear on the Replication Server Metrics section in the order of the data flow, appearing on the Replication Server pane.

The pane title line indicates the overall status of the replication server, which is either:

- HEALTHY—all threads are executing as expected.
- SUSPECT—a thread is down and the Replication Server expected it to be running. Or, a thread is in a Connecting state. The Connecting state means that either the server to which Replicated Server is connecting is unavailable and a problem exists, or the Replicated Server connects successfully in a moment and the suspect status is transitory.

Table 66. Replication Table Metrics

Name	Description
SQM	The status of the SQM used to receive information from the Primary Server, color-coded. You can drill down by clicking on this field, to show the SQM status.
Block writes (c/s)	The rate at which blocks are written to the inbound queue, given as count per second. You can drill down by clicking this component. A popup appears, showing the inbound queue history, plotted over time.
In. Queue Size (MB)	The current size of the inbound queue.
Block reads (c/s)	The rate at which data blocks are read from the inbound queue to the SQM associated with the Distributor, given as count per second.
SQT	The status of the SQM associated with the Distributor that receives information from the inbound queue, color-coded. You can drill down by clicking on this field, to show the SQM status.
Distributor	The status of the Distributor component.
SQM	The status of the SQM associated with the Distributor that writes information to the outbound queue, color-coded. You can drill down by clicking on this field, to show the SQM status.
Block writes (c/s)	The rate at which blocks are written to the outbound queue. You can drill down by clicking this component. A popup appears, showing the inbound queue history, plotted over time.
Out. Queue Size	The current size of the outbound queue.
Block reads (c/s)	The rate at which data blocks are read from the outbound queue to the DSI thread that sends the information to the related replicated server.
DSI	The status of the DSI thread that is associated with a replicated server.

Replicated Servers Metrics

The Replicated Servers pane shows the available replicated servers. For each server, it displays the following information:

- The replicated server name
- **DB:** The replicated server database name.
- **Latency (sec):** The difference between the time the transaction was committed at the destination (Replicate) database, and the time at the origin (Primary Site) when the transaction was committed, in seconds.
- **Agent status:** Shows the database name and the status of the Sybase_MDA Agent. The status of the Sybase_MDA agent:
 - Available—the Sybase_MDA Agent and the Replication Server are running on the same Foglight Management Server. This also indicates that the Sybase_MDA Agent is in the Collecting Data state.
 - No Agent—no Sybase_MDA Agent is running.


You can drill down by clicking on this status field. The [MDA Home Page Dashboard](#) appears for the appropriate server.

Databases Dashboard

The Databases dashboard displays summary information about all databases that may exist in your monitored environment. This includes your Sybase databases along with SQL Server, Oracle, and DB2 databases.

At the top of this dashboard you see the total counts of all alarms generated by all monitored database types, including Sybase. The alarm counts are grouped by the severity level: **Fatal**, **Critical**, **Warning**, **Normal**, and **Unknown**. Use this as a starting point to quickly estimate potential bottlenecks within your infrastructure.


Figure 39. Database Dashboard



NOTE: The flows associated with the **Discover more databases** and **Administer Agents** buttons, appearing in the top-right are only available when you are monitoring SQL Server and Oracle databases and have Foglight for SQL Server and Foglight for Oracle installed. For more information, see the *Managing Oracle Database Systems User Guide* and *Managing SQL Server Database Systems User Guide*.

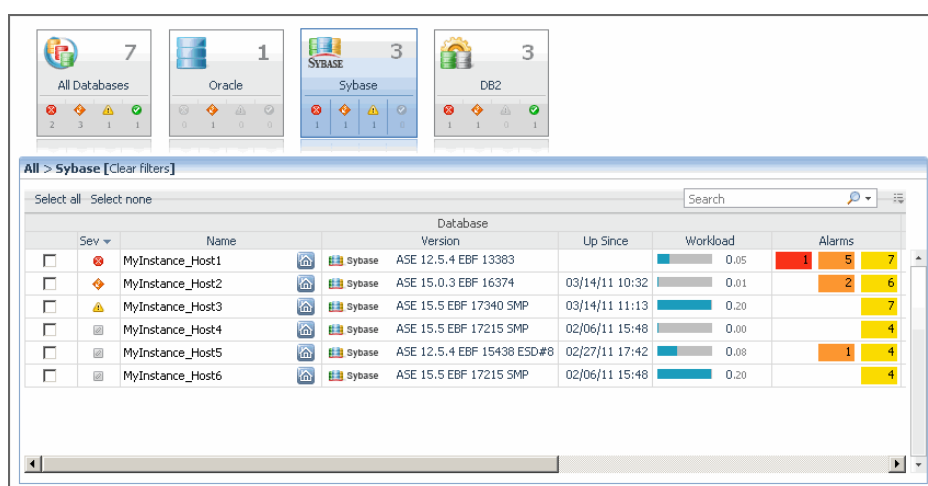
Just below the alarm counts, a collection of tiles displays a high-level overview of your environment. Each tile represents a database type (**All**, **SQL Server**, **Oracle**, **Sybase**, **DB2**) and shows how many database instances exist in your monitored environment, along with the count of objects of that type in each of the alarm states (Normal, Warning, Critical, Fatal).

Figure 40. Database Tiles



The table of database instances appearing just below the tile collection reflects the tile selection. For example, clicking the **Sybase** tile shows only the Sybase databases in the list.

Figure 41. Database Instances




For details about the information appearing in this list, see [Viewing the list of database instances](#) on page 69.

You can use other types of attributes to filter this list. For example, clicking a severity icon in a tile shows only the database instances of that database type that are in that alarm severity state. For more information about filtering this list, see [Filtering the list of database instances](#) on page 71.

In many cases it may be useful to create one or more groups of databases that you want to monitor. You can use database groups to further filter your selection using the **Databases** view appearing on the navigation panel.

Figure 42. Database Groups



For additional information about database groups, see [Managing database groups](#) on page 71.

Clicking the **All Databases** tile or **Clear filters** in the table heading shows a list of all monitored database instances.



Clicking any value in the database list, except the Home icon  in the **Name** column, displays a Sybase cue card at the bottom of the display area, showing additional information about the selected database instance. For details about the information appearing in this view, see [Exploring the state of selected database instances](#) on page 73.

Figure 43. Additional Database Instance Information



Clicking the Home icon  in the **Name** column links to the [MDA Home Page Dashboard](#).

Viewing the list of database instances


The table of database instances can display all available database instances, or a filtered list. For more information about possible filters, see [Filtering the list of database instances](#) on page 71.

Table 67. Database Instance Description


Name	Description
Database	
Sev	The alarm with the highest severity raised against the database instance.
Name	The name of the Sybase_MDA Agent that monitors this database instance.
Version	The type and version number of the monitored database instance.
Up Since	The date and time when the database instance was last restarted.
Workload	The workload graph for the database instance.

Table 67. Database Instance Description

Name	Description
	The number of Warning, Critical, and Fatal alarms for the database instance. The alarms are displayed by their severity levels, with the aggregated number for each severity. Clicking this column shows the Outstanding Alarms popup.



Alarms The popup shows the most recent alarms generated for the specific database instance. To view additional information about an alarm, click its entry in the **Sev** column. The **Alarm Created** dialog box appears, showing details about the selected alarm.



For complete information about alarms in Foglight, see the *Foglight User Help*.

System Utilization	
Host	The name of the computer on which the database instance is running.
CPU (%)	The overall OS CPU usage.
Memory (%)	The amount of memory consumed by all OS processes (including the database instance process). This includes both RAM resident memory and swapped memory.
Disk (% Busy)	The percentage of time the busiest device spent serving system-wide I/O requests. This metric serves as a measure for the system I/O load.


Table 67. Database Instance Description

Name	Description
Agent	The operational status of the monitoring agent.
	When the agent instance is running, hovering over the icon in the State column displays the status message <code>Collecting Data</code> .
State	When the agent instance is running but not collecting data, hovering over the State icon displays one of the following status messages: <ul style="list-style-type: none"> Starting Stopped Stopping Unknown
SC	Indicates whether StealthCollect is installed and configured for this agent. NOTE: StealthCollect connection is available only when you are monitoring SQL Server and Oracle databases, and have Foglight for SQL Server and Foglight for Oracle installed and running.

Filtering the list of database instances

By default, the list of database instances displays all of the currently monitored database instances. If your monitored environment includes a larger number of different database instances, it might be useful to filter this list and focus on individual groups of database instances that share some common attributes.

To filter the list of database instances:

- Choose one of the following ways to filter the list of database instances.
 - To list only the database instances of a selected database type, click the tile identifying that type. For example, to list only the Sybase database instances, click the **Sybase** tile.
 - To list only the database instances of a selected database type that are in a specific alarm state, click the alarm state icon on the database type tile. For example, to list only the Sybase database instances that are in Warning state, click  on the **Sybase** tile.
 - To list only the database instances that belong to a specific database group, on the navigation panel, in the **Databases** view, expand the **Database Groups** node and select a group. For more information about database groups, see [Managing database groups](#) on page 71.

The list of databases refreshes, showing only the entries that match the specified filter.

Managing database groups

In large monitored environments that include a high number of database instances it might be useful to create separate groups of database instances that share some common attributes, and to monitor the individual groups separately. Using the **Databases** view on the navigation panel, it is possible to group database instances either by identifying specific database instances, or services that contain specific database objects, and adding them to the group. This view allows you to create, edit, and remove database groups, thereby serving as a filter for the list of database instances.

When grouping database instances by services, you first need to associate the database instances with services and then identify the services that contain one or more objects used by the monitored databases. For more information about services in Foglight, see the *Foglight User Help*.

To filter the list of databases by selecting services:

- i | IMPORTANT:** Ensure that you already have Sybase services defined in Foglight, and that they include Sybase components.


- 1 To create a Sybase service, using the Service Builder, create a new category with a DB tier and add the applicable Sybase objects to it. For more information about Foglight services, see the *Foglight User Help*.
- 2 On the navigation panel, under **Databases**, click the box on the right of **Group by** and choose **Services**.

The **Databases** view refreshes, showing groups of monitored resources created using the Service Builder dashboard, and containing the Foglight for Sybase database topology. The **SLC** column, also appearing in this view, indicates the current availability of each service over a given period of time.

- 3 Select the services that are associated with the database instances that you want to see listed.

The display area refreshes, showing the database instances whose objects are monitored with the selected services.

Figure 44. Selected Services



To create a database sub-group:

- 1 On the navigation panel, under **Databases**, select **Database Groups** and click **+**. The **Add Sub Group** dialog box appears.
- 2 In the **Name** box, type the name of the sub-group. Optionally, provide the sub-group's description in the **Description** box.
- 3 Select one or more database instances in the **Available** list and click **>** to move them to the **Selected** list.

Alternatively, click **>>** to select all of the databases instances.

- i | TIP:** To remove database instances from **Selected**, use **<** and **<<**, as required.

- 4 Click **Ok**.

The sub-group name appears in the list and the database instances you selected in [Step 3](#) appear in the list of databases.


To manage database groups:

- To remove a database group, select it on the navigation panel and click **×**. When prompted, click **Remove**.
- To edit the list of database instances associated with a database group, select it on the navigation panel and click **✎**. Make your changes, as required.

Exploring the state of selected database instances

Clicking a database instance in the databases table opens a cue card, which gives you more details about the selected database.

Figure 45. Cue Card details



The cue card displays the following information.

Table 68. Cue Card Description




Data	Description
General	
DB type	Specifies the type of the monitored database instance.
Sampling interval	The length of time between data samples.
Last sample time	The time when the last data sample occurred.
Alarms	
	The number of Fatal alarms associated with the selected database instance.
	The number of Critical alarms associated with the selected database instance.
	The number of Warning alarms associated with the selected database instance.
Configuration	
Online engines	The number of engines associated with the selected database instance that are currently online.
Physical memory	The amount of physical memory available to the selected database instance.
Performance	
Connection in use	The number of connections handled by the selected database instance.
Active processes	The number of processes associated with the selected database instance that are currently taking place.
Engines vs. Host Utilization	
Avg. Engine Utilization	The percentage of average engine utilization.
Host CPU Load	The percentage of the host's CPU usage.
Host Memory Utilization	The amount of memory consumed by all OS processes (including the database instance process). This includes both RAM resident memory and swapped memory.
Storage	
Number of devices	The number of devices associated with this database instance.
Number of databases	The number of databases running on the host.
Total	The total amount of the space allocated to the devices.
Usage	The percent of the space allocated to the devices that is currently being used.
Total Wait Events	

Table 68. Cue Card Description

Data	Description
Graph	The wait time for this event in seconds plotted over time.
Home Page	Links to the MDA Home Page Dashboard .

Manage Thresholds Dashboard

The Manage Thresholds dashboard provides access to the list of thresholds. A topology filter is available to search for a specific set of thresholds. For more information about this dashboard, see the *Administration and Configuration Help*.

Figure 46. Manage Thresholds Dashboard

Database Web View > Home Page (Last-Sample) > Manage Thresholds			Mar 3, 2009 7:58:55 PM IST
Showing 1 - 8 of 8 thresholds as of Mar 03, 2009 15:48:15			
Filter by Topology Type	<input type="text"/>	Clear Filters	Refresh
Metric	Topology Type	Summary	
baselineAvailability	ServiceLevelPolicy	Fatal: 0.0(c), Critical: AvailabilityFatal(v), Warning: Availabil	
cpu_busy_pct	Sybase_MDA_EngineSummary	Normal: 0.0(c), Normal: EngineBusyCpuMinWarning(v), Wz	
dcache_hit_rate	Sybase_MDA_DataCacheSumma	Normal: 110.0(c), Warning: CacheEfficiencyHitRateWarnin	
mda_deadlocks_delta	Sybase_MDA_Locks	Critical: 1.0(c), Critical: 9999.0(c)	
mda_io_reads_rate	Sybase_MDA_DiskIOSummary	Normal: 2000.0(c), Warning: 4000.0(c), Critical: 8000.0(c),	
num_active	Sybase_MDA_ProcessSummary	Normal: 10.0(c), Warning: 20.0(c), Critical: 35.0(c), Critical:	
num_con_pct	Sybase_MDA_ResourcesInUse	Normal: -1.0(c), Normal: ConnectionsUsedWarning(v), War	
pcache_hit_rate	Sybase_MDA_ProcedureCache	Normal: 110.0(c), Normal: PcacheHitRate(v), Warning: Pca	

Manage Thresholds Table

Table 69. Manage Thresholds Table Description

Name	Description
Metric	The metric threshold name. You can edit a threshold by clicking on this name.
Topology Type	The topology type to which this threshold is assigned.
Summary	The summary of the threshold properties.

About the Sybase_MDA Agent

The Sybase_MDA Agent monitors the performance of the Sybase Adaptive Server Enterprise 12.5.1 and above. A unique characteristic is the ability for a user to select a "monitoring level" (called Availability Performance and Diagnostic). Each level is associated with a group of tables and their associated collection. By default, Availability and Performance tables are activated and all others are turned off. You can optionally select or clear Configuration, Performance or Space Trending to tune the agent to collect the metrics important to your environment.

In addition, all tables that can return multiple rows have a corresponding entry in the agent properties to limit the incoming data (see the Data Retrieval properties). The Limit Incoming Rows property displays a set of criteria developed to limit the amount of data written to the Foglight database. These values are used strictly to manage the amount of data written to the Foglight database and NOT to configure the Rules. Rules are still configured by registry variables.

Configuring Sybase Login Privileges

The Sybase_MDA agent collects performance metrics using the Sybase MDA tables. This requires the Sybase login to have only the role *mon_role* assigned to it, as described above. The Sybase login does not need to have the role *sa_role* assigned to it. From Cartridge for Sybase version 5.5.8.39, the agent installs the Foglight for Sybase stored procedures during the initial activation. Therefore, if you are upgrading the Cartridge for Sybase, no action is required regarding the stored procedures. If you are installing the agent for the first time, you must perform the following pre-requisite procedure.

To create a login account without 'sa' privileges:

- 1 Ensure that all of your MDA configuration parameters are set properly.

i | TIP: The file *setup_mon_tables.sql* contains all of the required parameters with sample values.

- 2 Create a Foglight login account. For example, fog123:

```
sp_addlogin 'fog123' , 'foglight' , sybsystemprocs
go
grant role mon_role to fog123
go
```

- 3 Run the *rapsAdmin.sql* script as the *sa* user.

The scripts included with the Foglight for Sybase can be found in the directory *<FglAM_Home>/agents/SybaseCartridge/<version>/config/Sybase/scripts*.

```
isql -Usa -P<password> -S<serverName> -i rapsAdmin.sql -o
rapsAdmin.out
```

The above command creates two helper stored procedures: *sp_fgl_addgrant* and *sp_fgl_adduser*. In addition, the script drops previously installed stored procedures and previously created temporary tables.

- 4 Run the *sp_fgl_adduser* procedure as the *sa* user. For example:

```
sp_fgl_adduser 'fog123' , 'foglightGroup'
go
```

The above command creates a group named *foglightGroup* and a user named *fog123* in each database, and in the *sybsystemprocs* database, it grants the permission for creating procedures to that user account.

i | IMPORTANT: The group and user names used here (*foglightGroup* and *fog123*) are just examples, you can substitute them with more appropriate values.

To update a non-sa user after the cartridge upgrade:

- 1 Complete [Step 3](#) and [Step 4](#) described in [To create a login account without 'sa' privileges](#).

The above procedure results in the new user being created (*fog123* in this example) in the group *foglightGroup* in each database, including the model database. Having that user and group in the model database causes a creation of a temporary database using these credentials when the Sybase server is restarted.

Using a secondary database to store collected data

Foglight stores data collected by monitoring an ASE server in the default *tempdb* database. In case the default database fills up, it is possible to store the collected metrics in another database. The cartridge includes a set of scripts that enable this feature:

- *rapsinstall2_questdb.sql*
- *rapsinstall2-1253_questdb.sql*
- *rapsinstall2_cluster_questdb.sql*

The following code block shows how to create a temporary database (*questdb_temp*), and bind a Foglight user (*fogqdb*) to that database.


```
disk init name = 'questdb_temp',  
physname = 'C:\sybase\data\questdb_temp.dat', size = '50M',  
    dsync = false  
create temporary database questdb on questdb_temp = 50  
sp_tempdb bind, 'LG', 'fogqdb', 'DB', 'questdb'
```

Configuring the Sybase_MDA Agent for Remote Monitoring

The Sybase_MDA Agent can be configured for remote monitoring. Remote monitoring does not require a Sybase_MDA Agent instance to run on the monitoring host. The agent instance can run on another machine, provided that you point it to the monitored machine using Sybase_MDA Agent properties.

To configure the Sybase_MDA Agent for remote monitoring:

- 1 Install and run the Foglight Agent Manager on the monitored host. You can use the Agent Manager embedded on the Foglight Management Server.
- 2 On the computer running the Agent Manager, install Sybase Open Client.
- 3 On the Sybase Open Client, use the `dsedit` tool to add an entry representing the monitored Sybase Server in the interface file.
- 4 Define a Sybase user for the Sybase cartridge.
- 5 Configure the Sybase_MDA Agent's **Sybase Home** property to point to the Open Client home directory: `%Sybase%` (Windows) or `$Sybase` (Unix).

 **TIP:** The Sybase Home agent property can be found in the Connection Details group of agent properties. For more information, see [Setting Connection Details](#) on page 77.
- 6 Configure the Sybase_MDA Agent's **Server Name** property to point to the host name or IP address of the monitored server, as defined in `dsedit`.

Sybase_MDA Agent Properties

When an agent connects to the Foglight Management Server, it is provided with sets of properties that it uses to configure its correct running state.

Each agent is provided with a combination of two types of properties: agent properties and shareable properties.

Default versions of these properties are installed with the Foglight for Sybase. However, you can edit the default shareable and agent properties, configure agent properties that apply only to a specific agent instance, and create edited clones of shareable properties that are used by a subset of the agents of a certain type.

For more information about working with agent properties, see the *Foglight Administration and Configuration Guide*.

The Sybase_MDA Agent is shipped with default properties that can be modified to suit your system requirements.

Agent properties described in this section comprise:

- [Setting Connection Details](#) on page 77
- [Setting Callback Connection Properties](#) on page 78
- [Setting Data Management Properties](#) on page 78

- [Setting Table Space Management Properties](#) on page 80
- [Setting Data Retrieval Properties](#) on page 81
- [Setting Benchmark Properties](#) on page 85
- [Setting Configuration Properties](#) on page 86

To modify agent properties:

- 1 Ensure that the navigation panel on the left is open.
To open the navigation panel, click the right-facing arrow on the left ►.
- 2 Open the dashboard that lets you navigate to the agent properties by completing one of the following steps:
 - On the navigation panel, under Dashboards, click **Administration > Agents > Agent Properties**.
In the Agent Properties dashboard, in the Namespace > Type pane, select **Legacy > Sybase_MDA**.
 - On the navigation panel, under Dashboards, click **Administration > Agents > Agent Status**.
In the Agent Status dashboard, select the instance of the Sybase_MDA agent whose properties you want to modify and click **Edit Properties**.

A list of agent properties appears in the **Sybase_MDA** pane.

Figure 47. Agent Properties

Connection Details	
Login Name	foglight
Login Password	*****
Sybase Home (\$SYBASE)	c:\sybase
Server Name (\$DSQUERY)	SYB2K
Backup Server Alias	SYB_BACKUP
Maximum Connection Retries	5
Connection Timeout (seconds)	60

Data Management	
Purge Data After	SYBM_Purging_List [Edit] [Clone] [Delete]
Availability	<input checked="" type="radio"/> True <input type="radio"/> False
Availability Poll Cycle (seconds)	60
Availability Monitoring	AvailMon_Dflt [Edit] [Clone] [Delete]
Performance	<input checked="" type="radio"/> True <input type="radio"/> False

The position of the Properties pane depends on the dashboard you used to access agent properties. If you used the Agent Properties dashboard, the Properties pane appears to the right of the Namespace > Type pane in the display area. If you used the Agent Status dashboard, the Properties pane appears across the display area.

Setting Connection Details

Use the Connection Details agent set of properties in the Sybase_MDA Agent to designate the type of connection properties to use.

To set up the Connection properties:

- 1 Enter the Sybase login in the **Login Name** box.
 - i** **NOTE:** The above user account must be configured to include the permissions for installing and executing stored procedures, as well as read permissions on system tables (including master and user databases). The required stored procedures must be installed manually by an account with appropriate privileges before the Sybase_MDA Agent can start monitoring the Sybase Adaptive Server. This can either be the Sybase *sa* account, or another account that has the appropriate permissions, configured using the *rapsinstall2.sql* script. For complete instructions, see “<Link>Configuring Sybase Login Privileges” on page 75.
- 2 Enter the password associated with the Login Name to access the Sybase database, in the **Login Password** box.
- 3 Enter the full path to the Sybase directory in the **Sybase Home (\$SYBASE)** box. Ensure that this path contains the Sybase ‘interfaces’ file.
 - i** **NOTE:** The directory path should look similar to */apps/Sybase*. Notice that the path does not contain a trailing slash.
- 4 Enter the name of the ASE server to be monitored in the **Server Name (\$DSQUERY)** box.
- 5 Enter the alternate name of the backup server, if one has been defined, in the **Backup Server Alias** box.

To retrieve a list of available servers, run the stored procedure *sp_helpserver*.

The default is SYB_BACKUP.
- 6 Enter the number of times you want Foglight to attempt to connect to Sybase in the **Maximum Connection Retries** box.

This is the maximum number of retries the Sybase_MDA Agent attempts when it encounters an error with the Adaptive Server before the agent stops.
- 7 Enter the number of seconds allowed for an individual connection attempt in the **Connection Timeout (seconds)** box.
- 8 Click **Save**.
- 9 Go to [Setting Callback Connection Properties](#) to enable the appearance of stored procedure details on the [Top Procedures Tab](#).

Setting Callback Connection Properties

Use the callback connection properties in the Sybase_MDA Agent to enable the appearance of stored procedure details, such as the stored procedure text and its execution statistics, on the [Top Procedures Tab](#).

To set up the callback connection properties:

- 1 Enter the name of the Sybase server host machine in the **Sybase Server Host** box.
- 2 Enter the port number used by the Sybase server host machine in the **Sybase Server Port** box.
 - i** **IMPORTANT:** Ensure that the configured port is open between the Foglight Management Server and the Sybase Server.
- 3 Go to [Setting Data Management Properties](#) to set sample frequency agent properties.


Setting Data Management Properties

Use the Data Management set of agent properties to set the sample frequency.

To set the data management properties:

- 1 Ignore the **Purge Data After** box. This field is not functional.

- 2 Select the **Availability Data** option, **True** or **False**, to turn the data filtering on or off. When this field is set to true, the agent monitors the Availability group.
The default value is true.
- 3 Enter the collection rate in the **Availability Poll Cycle (seconds)** box if you want to change the collection rate for the group.
The collection rate, or sample frequency, is the number of seconds between the end of a collection period and the start of the next.
The default is 60 seconds.
- 4 Click the list from the **Availability Monitoring** list that you want to update.
This agent property list allows the activation or deactivation of individual tables for collection.
The default secondary list name is *AvailMon_Dflt*.
- 5 Click **Edit**.
A secondary property list appears.
- 6 Add an entry to the list by clicking **Add Row**.
Fill in the dialog box fields.
- 7 Click **Save Changes**.
- 8 Click the list from the **Performance Monitoring** list that you want to update.
This agent property list includes performance monitoring properties that enable or disable both the collection of performance data and the length of the polling interval per-table.
The default secondary list name is *PerfMon_Dflt*.
- 9 Click **Edit**.
A secondary property list appears, listing all performance tables.
- 10 Click **Add Row** to add an entry to the list.
 - a In the **Table Name** column, type the table name.
 - b Select or clear the check box in the **Collect** column to enable or disable the data collection for this table. By default, this check box is selected.
 - c In the **Poll Cycle (sec)** column, type the length of the poll interval for the table in seconds. This interval represents the number of seconds between the end of one collection period and the start of the next.

 **IMPORTANT:** The *StatementCache-Future-Use* and *TableActivity-Future-Use* tables are for future use only. Do not make any changes to these values.
 - d Click **Save Changes**.
- 11 Select the **Configuration** option, **True** or **False**, to activate the table collection for the Configuration group.
The default value is false; collection deactivated.
- 12 Enter the collection rate in the **Configuration Poll Cycle (hours)** box if you want to change the collection rate for the group.
The collection rate, or sample frequency, is the number of hours between the end of a collection period and the start of the next.
The default is 24 hours.
- 13 Click the list from the **Configuration Monitoring** list that you want to update.
This agent property list allows the activation or deactivation of individual tables for collection.
The default secondary list name is *ConfMon_Dflt*.

- 14 Click **Edit**.

A secondary property list appears.

- 15 Add an entry to the list by clicking **Add Row**.

Fill in the dialog box fields.

- 16 Click **Save Changes**.

- 17 Select the **Space Trending** option, **True** or **False**, to activate the table collection for the Space Trending group.

- 18 Enter the collection rate in the **Trending Poll Cycle (hours)** box if you want to change the collection rate for the group.

The collection rate, or sample frequency, is the number of hours between the end of a collection period and the start of the next.

The default is 24 hours.

- 19 Click the list from the **Space Trending List** that you want to update.

This agent property list allows the activation or deactivation of individual tables for collection.

The default secondary list name is *SPTR_Dflt*.

- 20 Click **Edit**.

A secondary property list appears.

- 21 Add an entry to the list by clicking **Add Row**.

Fill in the dialog box fields.

- 22 Click **Save Changes**.

- 23 Select the **Top SQL** option, **True** or **False**, to activate the table collection for the Top SQL group.

- 24 Enter the collection rate in the **Top SQL Poll Cycle (seconds)** box if you want to change the collection rate for the group.

The collection rate, or sample frequency, is the number of hours between the end of a collection period and the start of the next.

The default is 300 seconds.

- 25 Select the list from the **Top SQL Monitoring** list that you want to update.

This agent property list allows the activation or deactivation of individual tables for collection.

The default secondary list name is *TopSQL_Dflt*.

- 26 Click **Edit**.

A secondary property list appears.

- 27 Add an entry to the list by clicking **Add Row**.

Fill in the dialog box fields.

- 28 Click **Save Changes**.

- 29 Click **Save**.

- 30 Go to [Setting Table Space Management Properties](#) to set the table space properties.

Setting Table Space Management Properties

Use the Table Space Management set of agent properties to set the sample frequency for table spaces, enable the data collection for a database, the minimum table size to retrieve, the number of tables to retrieve, and the total number of tables, indices, and partitions for per agent instance.

To set the table space management properties:

- 1 Select the **Table Space** option, **True** or **False**, to turn the collection of table space data on or off. When this field is set to **True**, the agent collects the table space metrics. The default value is **True**.
- 2 To change the sample frequency for the collection of table space data, edit the collection rate in the **Table Space Poll Cycle (minutes)** box.

The collection rate, or sample frequency, is the number of minutes between the end of a collection period and the start of the next. The default is five minutes.
- 3 Click the list from the **Table Space Monitoring** list that you want to update.

This list property enables the data collection for a database. When you add a database to this list, the database metrics appear on the [Table Spaces Tab](#) and [Table Performance Tab](#) of the [Database Spaces and Trends Dashboard](#) (see page 48).

The default list is *TableSpace_Dflt*.
- 4 Click **Edit**.

A dialog box appears, showing a list of secondary properties.
- 5 Add an entry to the list by clicking **Add Row**.

Specify the following values in the row that appears:
 - **Database Name**: The name of the database.
 - **Threshold MB**: The minimum table size to retrieve. The default is 1 MB.
 - **Number Of Top Tables**: The number of tables to monitor. The default is 10.
- 6 Click **Save Changes** and close the dialog box.
- 7 Edit the total number of tables, indices, and partitions in the **Total Number Of Tables, Indexes and Partitions** box.
- 8 Click **Save**.
- 9 Go to [Setting Data Retrieval Properties](#) to set sample frequency agent properties.

Setting Data Retrieval Properties

Use the Data Retrieval set of agent properties to set the limit for the amount of data collected and written to the Foglight database.

The Top Procedures set of agent properties affect the information that appears on the Top Procedures dashboard. For more information, see [Top Procedures Tab](#) on page 25.

The Top Users set of agent properties affect the information that appears on the Top Users dashboard. For more information, see [Top Users Tab](#) on page 27.

The Cached Objects set of agent properties affect the information that appears on the Data Cache Activity dashboard. For more information, see [Data Cache Activity Tab](#) on page 43.

To set the data retrieval properties:

- 1 Click the list from the **Limit Incoming Rows List** that you want to update.

This property displays a set of criteria developed to limit the amount of data written to the Foglight database. For more information, see [Limit Incoming Rows Criteria](#) on page 83.

The default name is *LIR_List_Dflt*. This list can be renamed to better describe the changes made.

These values are used strictly to manage the amount of data written to the Foglight database and NOT to configure the Rules. Rules are still configured by registry variables.
- 2 Click **Edit**.

A secondary property list appears.

- 3 Add an entry to the list by clicking **Add Row**.

Fill in the dialog box fields.

CAUTION: Only the criteria in the Value column can be changed. DO NOT change the Foglight Table Group and Limit Criteria values.

- 4 Click **Save Changes**.

- 5 Enter the maximum number of top applications to monitor in the **Top Applications - Max Rows Returned** box.

This is the maximum number of applications written to the Foglight database.

The default value is 10 applications. This value specifies the maximum size of the list.

- 6 Enter the minimum threshold value (the minimum value for the selected criteria used for an application to be potentially included on the list based on the maximum number of entries for that table), in the **Top Applications - Minimum Threshold Value** box.

The default value is 10 units of the selected sorting criteria.

- 7 Select a sorting criteria from the **Top Applications - Sort Criteria** list. For more information, see [Minimum Threshold Value Sorting Criteria](#) on page 84.

An application is included on the list if:

- The measurement for the specified criteria is over the threshold value.
AND
- There is room on the list based on the **Top Applications - Max Rows Returned** value specified in Step 5.

- 8 Enter the maximum number of top SQL statements to monitor in the **Top SQL -Max Rows Returned** box.

The number entered is the maximum number of SQL statements written to the Foglight database.

The default value is 10 SQL statements. This value specifies the maximum size of the list. Collecting the SQL statements inside stored procedures typically results in a higher number of SQL statements. If you intend to use this feature, increase this value, as required.

- 9 Enter the minimum threshold value (the minimum value for the selected criteria used for an SQL statement to be potentially included on the list based on the maximum number of entries for that table), in the **Top SQL - Minimum Threshold Value** box.

The default value is 10 units of the selected sorting criteria.

- 10 Select a sorting criteria from the **Top SQL - Sort Criteria** list. For more information, see [Minimum Threshold Value Sorting Criteria](#) on page 84.

An SQL statement is included on the list if:

- The measurement for the specified criteria is over the threshold value.
AND
- There is room on the list based on **Top SQL - Max Rows Returned** value specified in Step 8.

- 11 Select **Top SQL - Include Stored Procedure Lines - YES** to display the list of stored procedures. Select **NO** if the stored procedures list is not to be displayed. Enabling this property allows you to collect the SQL statements inside stored procedures. The number of collected SQL statements depends on the values set by the **Top SQL - Max Rows Returned** property (see [Step 8](#)). It also requires a certain configuration of ASE parameters: the parameter `statement pipe max messages` should be greater than `sql text pipe max messages`. For more information about these parameters, refer to your Sybase documentation.

- 12 Enter the maximum number of top users to monitor in the **Top Users - Max Rows Returned** box.

This is the maximum number of users written to the Foglight database.

The default value is 10 users. This value specifies the maximum size of the list.database.

- 13 Enter the minimum threshold value (the minimum value for the selected criteria used for a user to be potentially included on the list based on the maximum number of entries for that table), in the **Top Users - Minimum Threshold Value** box.

The default value is 10 units of the selected sorting criteria.

- 14 Select a sorting criteria from the **Top Users - Sort Criteria** list. For more information, see [Minimum Threshold Value Sorting Criteria](#) on page 84.

A user is included on the list if:

- The measurement for the specified criteria is over the threshold value
AND
- There is room on the list based on the **Top Users - Max Rows Returned** value specified in step 10.

- 15 Enter the maximum number of cached objects to monitor in the **Cached Objects - Max Rows Returned** box.

The default value is 10 cached objects.

- 16 Enter the minimum threshold value (the minimum value for the selected criteria used for a cached object to be potentially included on the list based on the maximum number of entries for that table), in the **Cached Objects - Minimum Threshold Value** box.

The default value is 1000 units of the selected sorting criteria.

- 17 Select a sorting criteria from the **Cached Objects - Sort Criteria** list. For more information, see [Minimum Threshold Value Sorting Criteria](#) on page 84.

- 18 Select the **Cached Objects - Object Type** from the list. The types are:

- **All tables**
- **System Tables**
- **Temporary Table**
- **User Tables**

- 19 Click **Save**.

- 20 Go to [Setting Benchmark Properties](#) to set sample frequency agent properties.

Limit Incoming Rows Criteria

The **Limit Incoming Rows List** table displays a set of criteria which limits the amount of data written to the Foglight database.

Table 70. Limit Incoming Rows List Criteria

Group	Limit Criteria	Value
Blocked Users	Max rows returned	20
Blocked Users	Minimum time blocked (seconds)	30
Databases	Max rows returned	20
Databases	Minimum Database space percentage threshold	80%
Databases	Minimum Transaction Log space percentage threshold	80%
Databases	Minimum Segment space percentage threshold	80%
Disk I/O	Max rows returned	20
Disk I/O	Minimum percentage of waits over requests	1%
Error Log	Max rows returned	10
Named Caches	Max rows returned	10
Named Caches	Maximum Cache Hit Rate threshold	80
Engines	Max rows returned	10

Table 70. Limit Incoming Rows List Criteria

Engines	Minimum CPU busy percentage threshold	80%
Long Running Transaction	Max rows returned	10
Long Running Transaction	Length of a long running transaction	-1
OS Processes	Max rows returned	10
OS Processes	Minimum OS Process CPU Busy threshold	80
System Waits	Minimum wait event threshold	80

Minimum Threshold Value Sorting Criteria

The Minimum Threshold Value list displays a set of criteria used to specify a quantity which orders the list of:

- [Top Applications](#)
- [SQL Statements](#)
- [Users](#)
- [Cached Objects](#)

Top Applications

Table 71. Top Applications Description

Criteria	Description
CPU Time (seconds)	The CPU time (in seconds) used by the application. This maps to the <code>cpu_time</code> field in the TopApplications Table.
Logical Reads	The number of buffers read from cache. This maps to the <code>logical_reads</code> field in the TopSQL Table and TopUsers Table.
Memory Usage (kB)	The amount of memory (in kilobytes) allocated to the application. This maps to the <code>mem_usage_kb</code> field in the TopApplication Table.
Physical Reads	The number of buffers read from disk. This maps to the <code>physical_reads</code> field in the TopApplication Table.
Physical Writes	The number of buffers written to disk. This maps to the <code>physical_writes</code> field in the TopUsers Table.

SQL Statements

Table 72. SQL Statements Description

Criteria	Description
CPU Time (seconds)	The number of seconds (ms) of CPU time used by the SQL statement. This maps to the <code>cpu_time</code> field in the TopSQL Table.
Duration	The amount of time the SQL statement took to run in seconds. This maps to the <code>duration</code> field in the TopSQL Table.
Logical Reads	The number of buffers read from cache. This maps to the <code>logical_reads</code> field in the TopSQL Table and TopUsers Table.
Memory Usage (kB)	The amount of memory (in kilobytes) used for SQL statement execution. This maps to the <code>mem_usage_kb</code> field in the TopSQL Table.
Packets Received	The number of packets received by Adaptive Server. This maps to the <code>packets_received</code> field in the TopSQL Table.
Packets Sent	The number of network packets sent by Adaptive Server. This maps to the <code>packets_sent</code> field in the TopSQL table.

Table 72. SQL Statements Description

Criteria	Description
Physical Reads	The number of buffers read from disk. This maps to the <code>physical_reads</code> field in the TopSQL and TopUsers table.
Physical Writes	The number of buffers written to disk. This maps to the <code>physical_writes</code> field in the TopSQL table.

Users

Table 73. Users Description

Criteria	Description
CPU Time (seconds)	The CPU time (in seconds) used by the process. This maps to the <code>cpu_time</code> field in the TopUsers table.
Locks	The number of locks the process currently holds. This maps to the <code>locks_held</code> field in the TopUsers table.
Logical Reads	The number of buffers read from cache. This maps to the <code>logical_reads</code> field in the TopSQL and TopUsers table.
Memory Usage (kB)	The amount of memory used by the top users in kilobytes.
Physical Reads	The number of buffers read from disk. This maps to the <code>physical_reads</code> field in the TopSQL and TopUsers table.
Physical Writes	The number of buffers written to disk. This maps to the <code>physical_writes</code> field in the TopUsers table.
Transactions	The number of transactions being processed by the top users.

Cached Objects

Table 74. Cached Objects Description

Criteria	Description
Cached (kB)	The amount of memory used by the cached objects in kilobytes.
Processes Accessing	The number of processes accessing the cached objects.

Setting Benchmark Properties

Use the Benchmark set of agent properties to configure *BenchMarkTimes* table collection parameters.

To configure the benchmark properties:

- 1 Click the **Read Benchmark SQL from a File** radio button, True or False, to enable or disable this function.
To have Foglight read the Benchmark SQL from a file, select true.
- 2 Specify the SQL Benchmark in the **Benchmark SQL or Path** box. Either:
 - Enter the full path to the file; that is, a file containing an SQL statement that you want to run on a regular basis to track response time.
 - Enter the SQL statement directly.

i **NOTE:** Do NOT select true for the **Read Benchmark SQL from a File** property if you enter the SQL statement directly.
- 3 Enter the maximum amount of time in seconds that the SQL query should run in the **Benchmark Response Time (seconds)** box.

The rules associated with this agent property are based on a percentage of this time. For example, if the SQL query runs at or above the specified time a Warning level alert is issued. If it runs at 90 percent of the specified time, a Critical level alert is issued.

The default is 1 second.

- 4 Click **Save**.
- 5 Go to [Setting Configuration Properties](#) to set sample frequency agent properties.

Setting Configuration Properties

Use the Configuration set of agent properties to set the log collection configuration properties.

The configuration properties allow you to define the logs that the Sybase_MDA Agent monitors. Additionally, it allows you to define the messages within the logs that the Sybase_MDA agent traps.

To set the log collection properties:

- 1 Enter the full path to the Adaptive Server error log in the **ASE Error Log Path** box.
- 2 Enter the full path to the Backup Server error log in the **Backup Server Error Log Path** box.
- 3 Click the list from the **Message Trap List** that you want to update.

Foglight ships with a default message trap list. The message trap list is a filter that traps specific error log messages. The default name of this list is *ASE_Error_Log_Search_Patterns*.

The number of occurrences of each message, along with the severity and general description, is relayed back to Foglight, resulting in an alarm being generated. Check the error log file for the full text and content of the error message detected.

It is also possible to use error log messages in triggered or stored procedures by printing them to the log adding them to the message trap list. Error messages in the log appear with the `debug` prefix and result in an alarm being fired with the severity level specified in the message trap list.

- 4 Click **Edit**.
A secondary property list appears.
- 5 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
 - a **Search Pattern**—Indicates the regular expression pattern used to match the log message.
 - b **Action**—When set to `Include`, any error messages that match the search pattern are printed to the log. When set to `Exclude`, the matched messages are excluded from the log.
 - c **Severity**—The severity of the alarm that is triggered by any error messages that match the search pattern.
- 6 Click **Save Changes**.
- 7 Select whether or not databases are to be included or excluded from the **Use (Include DBs or Exclude DBs or None)** list. Click one of the following:
 - **None** to indicate that neither of the **Exclude DBs** or **Include DBs** lists is to be used. Every database is treated according to standard behavior for the **Limit Incoming Rows** agent property. For more information, see [Setting Data Retrieval Properties](#) on page 81. The default is none.
 - Select **Exclude DBs** to monitor most databases and exclude a few databases.
 - Select **Include DBs** to exclude most databases and monitor a few databases.

Note the following:

- In the **Include DBs** field, change the name of the of the agent property list from *DB_Include_List_Dflt* to *SYBM_DBInclusion_<Modified Agent Instance Name>*.

- For inclusion, <Modified Agent Instance Name> is a placeholder for the agent instance name as Foglight understands it. The Modified Agent Instance Name usually looks like `Sybase_MDA@dbhost1`. The '@' in the instance name is converted to a '_', so the modified name looks like `Sybase_MDA_dbhost1` and the full list name looks like `SYBM_DBInclusion_Sybase_MDA_dbhost1` for the inclusion list.
- If inclusion is activated and the list is named incorrectly, the exact name required by the agent is posted in the alert message. Then you can cut and paste the name into the appropriate list name field.

CAUTION: The agent pre-populates the inclusion list. The **Add** and **Delete** buttons should not be used.

8 Click the exclusion database list from the **Exclude DBs** list that you want to update.

9 Click **Edit**.

A secondary property list appears.

10 Add an entry to the list by clicking **Add Row**.

Fill in the dialog box fields.

- a For the **Status Alerting**, **Database Space Alerting**, **Transaction Log Alerting**, or **Space Trend** lists, the options are:

Standard—Foglight returns a row for this database only when it crosses the threshold. A row for the respective database may not be returned when it crosses a threshold if the number of databases crossing the threshold exceeds the “max rows returned” value during that poll cycle.

Always Include—Foglight always returns a row for this database when it crosses the threshold. The row for the respective database is always returned, even when the number of databases crossing the threshold exceeds the “max rows returned” value of that poll cycle.

Exclude—Foglight does not return a row for this database even if the threshold is crossed.

- b In the **Space Trend** list, the options are:

Exclude—Foglight does not return a row for this database even if the threshold is crossed.

Include—Foglight returns a row for this database.

- c Click **True** or **False** on the **Disable All** radio button to enable or disable this function. Set this value to true to exclude this database from Status Alerting, Database Alerting, Transaction Log, Alerting and Space Trend Alerting.

11 Click **Save Changes**.

12 Click the inclusion database list from the **Include DBs** list that you want to update.

This criteria is only applied if **Include DBs** is selected in the **Use (Include DBs or Exclude DBs or None)** property.

IMPORTANT: The agent pre-populates the inclusion and exclusion lists. The **Add** and **Delete** buttons should not be used.

13 Click **Edit**.

A secondary property list appears.

14 Add an entry to the list by clicking **Add Row**.

Fill in the dialog box fields.

- a For the **Status Alerting**, **Database Space Alerting**, or **Transaction Log Alerting** lists, the options are:

Standard—Foglight returns a row for this database only when it crosses the threshold. A row for the respective database may not be returned when it crosses a threshold if the number of databases crossing the threshold exceeds the “max rows returned” value during that poll cycle.

Always Include—Foglight always returns a row for this database when it crosses the threshold. The row for the respective database is always returned, even when the number of databases crossing the threshold exceeds the “max rows returned” value of that poll cycle.

Exclude—Foglight does not return a row for this database even if the threshold is crossed.

- b In the **Space Trend** list, the options are:

Exclude—Foglight does not return a row for this database even if the threshold is crossed.

Include—Foglight returns a row for this database.

- c Click **True** or **False** on the **Disable All** radio button to enable or disable this function. Set this value to true to exclude the database from Status Alerting, Database Alerting, Transaction Log, Alerting and Space Trend Alerting.

15 Click **Save Changes**.

16 Click **Save**.

Sybase_MDA Investigations

This section describes the following investigations:

- [Investigating Application Usage](#)
- [Investigating CPU Capacity](#)
- [Investigating Data Cache Performance](#)
- [Investigating Database Space](#)
- [Investigating Disk Performance](#)
- [Investigating Index Performance](#)
- [Investigating Locking](#)
- [Investigating Meta Data Cache Contention](#)
- [Investigating Network Activity](#)
- [Investigating Parallel Queries](#)
- [Investigating Procedure Cache](#)
- [Investigating Resource Utilization](#)
- [Investigating Sybase Performance](#)
- [Investigating the User Log Cache](#)

Investigating Application Usage

The Sybase_MDA Agent allows you to quickly view and determine which applications are using the most CPU resources.

To start with the SYBM_TopApplications table view:

- Open the SYBM_TopApplications table view which displays application performance metrics such as CPU, I/O, and memory for the top N applications running on the Adaptive Server.

It is critical to know which applications, on any Adaptive server, are consuming resources. This understanding allows the DBA to tune the Adaptive Server resources to the priority of each application.

You can configure how many applications are returned to Foglight, the criteria by which they are sorted, and the minimum threshold that the selected criteria must cross before being returned. This is completed through the [Setting Data Retrieval Properties](#) agent property for the Sybase_MDA Agent.

TIP: If no applications are being displayed, used the procedure below.

To ensure that applications are displayed:

- 1 Go to the Sybase_MDA Agent properties.
- 2 Go to the **Data Management** set of agent parameters.
- 3 Ensure that the **Performance** parameter has the **True** button selected, so that performance data is collected.
- 4 Ensure that the **Performance Monitoring** list item **Top Applications** is set to True.

Investigating CPU Capacity

The Foglight Cartridge for Sybase assists your investigation into Sybase ASE CPU use and availability.

This cartridge tracks internal allocation of CPU resources by the Adaptive Server as well as the external performance characteristics of the Adaptive Server processes from the perspective of the host Operating System.

To start with the SYBM_OS_Cpu graph:

- Open the SYBM_OS_Cpu graph view which displays how busy the Adaptive Server engines are during the time that the CPU is available to the Adaptive Server.

This graph shows total CPU utilization, as measured from within the ASE, charted against what the operating system is measuring. Remember that the ASE reports CPU usage as idle even though it may be performing housekeeping tasks which, in turn, are reflected as CPU activity from the Operating System.

What to look for when viewing this chart:

- If the system CPU use is high, investigate the I/O performance of the Adaptive Server. If the Adaptive Server is spending an excessive amount of time performing I/O, then the device layout and allocation of the server may need to be tuned.
- If the CPU utilization is high for users, but the connections in use are not increasing, go to the SYBM_TopUsers table view to see if a single user is consuming disproportional amount of resources. This view identifies users who have crossed your predefined threshold.
- Use this chart to watch for the point where connections saturate the CPU resource available. You should be able to project the point where the number of connections causes the CPU use to rise to levels that may cause overall ASE performance degradation.

To see more detailed information:

- Open the SYBM_OS_CpuDetail graph view. This graph shows Operating System CPU utilization broken down by each respective ASE engine process.

This lets you monitor the CPU performance reported within the Adaptive Server. This view shows the CPU activity broken down by user, system, and idle usage level charted against the number of connections in use.

- Open the SYBM_EnginePerfDetail graph view, where the CPU utilization levels are broken down by Adaptive Server engine.

From the SYBM_EnginePerfDetail graph view, you can drill down to the SYBM_EngineInfoDetail table view for general information on each engine, such as engine status or start time.

Investigating Data Cache Performance

Sybase Adaptive Server reserves memory to hold the data, index, log pages currently in use, and pages recently used by the Adaptive Server. This region of the memory is called the data cache. The initial installation of Adaptive Server creates a single default data cache that is used for all data, index, and log activity.

A DBA can create and configure additional regions for data cache which are commonly referred to as named caches.

The named data caches can only be used by databases or database objects that are explicitly bound to them. All objects not explicitly bound to named data caches use the default data cache.

Data caches can be sub-divided into regions called data pools. These data pools can be of varying page sizes. The amount of space allocated to a data pool and the corresponding page size allows a DBA to tune the cache for larger I/Os and, in turn, more efficient I/O utilization than the default settings would provide.

To start with SYBM_DataCacheSummary graph:

Open the SYBM_DataCacheSummary graph view which summarizes the effectiveness of the cache design on the server as a whole. This view displays the total number of hits and misses for all caches during the collection period.

- From the SYBM_DataCacheSummary graph, drill down on the Hits line to investigate where the cache misses are taking place with the SYBM_DataCacheDetail graph view.
- From the SYBM_DataCacheDetail graph view, drill down on the Hits bar to see the performance characteristics of the data pools with the SYBM_DataCachePools graph view.

To view more detailed information on MetaData Cache or Procedure Cache:

- See the [Investigating Meta Data Cache Contention](#) and [Investigating Procedure Cache](#) topics.

Investigating Database Space

The Foglight Cartridge for Sybase assists your investigation into database space with the SYBM_DbSpace graph and SYBM_TempDbSpace graph views. Both are identical except that the temporary databases are reported separately.

These views show the data space utilization for the respective database.

Transaction log space is reported separately in SYBM_DbSpaceTransLog graph and the SYBM_TempDbSpaceTransLog graph views.

Understanding what databases are returned

Both sets of views only report on databases that have crossed a threshold defining how much space is used. This threshold is set by editing the respective agent properties.

To set the threshold:

- 1 Go to the agent properties by **Dashboards: Administration > Agents > Agent Status**.
- 2 Select the Sybase_MDA Agent and click **Edit Properties**.
- 3 Go to the Data Retrieval set of properties and click the **Edit** button.
- 4 Double-click a **Foglight Table Group**.
- 5 Enter a **Limit Criteria Value** and click the **Save** button.

This new value increases or decreases the number of databases that are returned to Foglight and reported in this view.

i **NOTE:** The Limit Criteria “Max Rows Returned” limits the number of rows when the percentage threshold is low and/or the number of databases returned to Foglight increases. Once the number of Max Rows is reached, Foglight discards the rows over the Max number allowed.

See the Foglight online help for more detailed information on editing agent properties.

For each database returned, you can drill down to the *SYBM_DbSpaceSegment* graph view or the *SYBM_TempDbSpaceSegment* graph view to see space allocations per segment.

To see database utilization over a period of time, go to the *SYBM_DbSpaceTrendChart* graph view.

NOTE: You must set **Space Trending** to True (default is off) in the Data Management agent properties for this agent.

Investigating Disk Performance

The Foglight Cartridge for Sybase assists your investigation into Sybase I/O use and availability by allowing you to:

- See overall server I/O activity. Use this to determine peak performance and adjust disk resources or job scheduling.
- See individual device utilization and make changes to the I/O distribution. SybaseASE allows you to control the placement of databases, tables, and indexes across your physical storage devices. This can improve performance by equalizing the reads and writes to disk across many devices and controllers.
- Review when your Adaptive Server is experiencing any semaphore contention problems on a device.

To start with the *SYBM_DiskIOSummary* graph:

- Open the *SYBM_DiskIOSummary* graph view, you can drill down to view the number of I/O reads and writes for each specific device.

To view more detailed information related to the *SYBM_DiskIOSummary* graph:

- Drill down on the Reads line to see the detail I/O activity for each device in the *SYBM_DiskIODetail* graph view.

To ensure optimal I/O performance in the Adaptive Server:

- 1 Spread data across disks to avoid I/O contention.
- 2 Isolate the server-wide I/O from the database I/O.
- 3 Separate data storage and log storage for frequently updated databases.
- 4 Keep the random disk I/O away from the sequential disk I/O.
- 5 Mirror devices on separate physical disks.
- 6 Partition tables to match the number of physical devices in a segment.

To start with the *SYBM_DiskIODetail* graph:

- Use the *SYBM_DiskIODetail* graph view to help you determine if the I/O is evenly distributed across the devices as planned:

To view more detailed information related to the *SYBM_DiskIODetail* graph:

- Drill down on the Reads lines to see the *SYBM_DiskIOSemaphores* graph view to check if any semaphore contention exists on that device.

This view helps track the number of I/Os requested against the number of requests that were forced to wait for synchronization of an I/O request. A high number of semaphores forced to wait could indicate I/O saturation for that device.

Investigating Index Performance

The Foglight Cartridge for Sybase assists your monitoring of general index performance metrics. You can use this information to analyze your application workload.

To start with the SYBM_IndexUse graph:

- The SYBM_IndexUse graph view provides an overview of all processes and their summarized index performance. All processes are summarized by calculating the number of pages accessed with and without the use of indexes. This is presented as a server-wide percentage that is tracked.

This view helps you determine if it is time to perform index maintenance, or at least investigate your index strategy. Low values in these metrics negatively impact overall user response times and indicate unnecessary resource usage.

To determine more specifically where improvements can be made:

- Use a tool like Quest Software's Spotlight on Sybase ASE or Sybase's own DBA tools to view individual processes and objects that are being accessed.
- Review Sybase's Performance and Tuning manuals for a full discussion on Indexes and how their performance impacts user response times.

Investigating Locking

The Adaptive Server protects the tables or data pages currently used by active transactions by locking them. Locking is a concurrency control mechanism: it ensures the consistency of data across transactions. Locking is needed in a multi-user environment, since several users may be working with the same data at the same time.

Locking affects performance when one process holds locks that prevent another process from accessing needed data. The process that is blocked by the lock sleeps until the lock is released.

The Foglight Cartridge for Sybase assists your investigation into locking performance in the Adaptive Server by measuring specific locking characteristics of your server.

To start with the SYBM_LockSummary graph:

- Open the SYBM_LockSummary graph view to display the number of locks requested on the system, locks that had to wait to be granted, locks that timed out, and the number of deadlocks on the system.

To view more detailed information:

- Drill down on the locks time out to view the SYBM_LockWaitTime graph view to display the average wait time for locks. Use this chart to discover any specific period of time where the performance noticeably degrades. Investigate respective user and application activity during this time frame using one of the SYBM_ResourcesInUse or one of the SYBM_Top Foglight views.

A deadlock occurs when two user processes each have a lock on a separate data page, index page, or table, and each wants to acquire a lock on the other process's page or table. It is possible to encounter deadlocks when many long-running transactions are executed at the same time in the same database. Deadlocks become more common as the lock contention increases between those transactions, which decreases concurrency.

If deadlocks occur often, it severely affects the throughput of applications. Deadlocks cannot be completely avoided. However, redesigning the way transactions access the data can help reduce their frequency.

The full details of a deadlock appear in the ASE error log. Use Quest's Spotlight on Sybase ASE product or review the messages in the Adaptive Server's error log to track down which object and users were involved.

Tracking the number of locks used by the Adaptive Server

The number of locks required by a query can vary widely, depending on the locking scheme, the number of concurrent and parallel processes and the types of actions performed by the transactions. Configuring the correct number for your system is a matter of experience and familiarity with the system.

System administrators can use `sp_configure` to change this limit. For example:

```
sp_configure "number of locks", 25000
```

You may also need to adjust the `sp_configure` parameter total memory, since each lock uses memory.

A shortage of the number of locks available can have an impact on the overall locking scheme in place on the server.

Investigating Meta Data Cache Contention

The Foglight Cartridge for Sybase assists your investigation into the performance characteristics of the Metadata Cache.

The Metadata caches reside in the kernel and server structures portion of Adaptive Server memory. You configure space for each of these caches by defining the number of each type of cache that the server allocates. This is done with the respective properties to the Sybase system procedure `sp_configure`.

Tracking and managing individual metadata caches for databases, indexes, or objects is important for a database that contains a large number of indexes and objects and where there is an expectation of high concurrency among users.

To start with the *SYBM_MetaDataCache* and *SYBM_MetaDataCacheDetail* graphs:

- Use the *SYBM_MetaDataCache* graph view to track the percentage used for each type of Metadata cache.

The *SYBM_MetaDataCacheDetail* graph view tracks the number of used and free descriptors for each cache. Use these numbers as a guide if any changes to the memory configuration need to be adjusted due to high utilization (especially if the number of descriptors that are reused are greater than zero).

When descriptors need to be reused, there can be performance problems, particularly with open databases. An open database contains a substantial amount of metadata information, which means that to fill an open database, the Adaptive Server needs to access the metadata on the disk many times; the server can also have a spin lock contention problem. The Foglight rule *MetaData_Open* alerts you when any of the cache utilizations are over the defined threshold and when any descriptors are reused.

i IMPORTANT: The Adaptive Server uses the metadata cache: When the Adaptive Server opens a database or accesses an index or an object, it needs to read information in the respective system tables: *sysdatabases*, *sysindexes*, and *sysobjects*.

The metadata caches for databases, indexes, or objects let the Adaptive Server access the information that describes it in the *sysdatabases*, *sysindexes*, or *sysobjects* row that are directly in its in-memory structure. This improves performance by allowing the Adaptive Server to bypass expensive calls that require disk access. Synchronization and spin lock contention are also reduced when the Adaptive Server has to retrieve database, index, or object information during runtime.

Investigating Network Activity

The Foglight Cartridge for Sybase assists your investigation into network usage. Network activity refers to the number of bytes and packets that are sent and received by your network.

To start with the *SYBM_OverviewNetwork* graph:

- Use the *SYBM_OverviewNetwork* table view to get a server-wide look at network activity detailed by bytes and packets both sent and received.

Watch for packet errors. An increase in errors paired with a decrease in traffic can highlight network issues that may need attention by your network administrator.

Average packet size can help you size your Adaptive Server.

In most situations, the number of packets being transferred is more pertinent than the size of the packets. It should be noted that overall network performance also includes the time needed by the CPU and operating system to process a network packet. This overhead for each packet affects performance the most. Larger packets decrease the overall overhead costs and achieves higher physical throughput, provided that you have enough data to be sent.

To view more detailed information:

- Drill down on the Bytes Received or Packets Received columns to view the *SYBM_NetworkBytes* graph or *SYBM_NetworkPackets* graph views respectively. These graph views display the number of bytes or packets data received and sent through the network.

Investigating Parallel Queries

Parallel query optimization is the process of analyzing a query and choosing the best combination of parallel and serial access methods to yield the fastest response time for the query. In addition to the costing performed for serial query optimization, parallel optimization analyzes the cost of parallel access methods for each combination of join orders, join types, and indexes.

To start with the *SYBM_ProcessesParallel* graph:

- Use the graph to tune your server parameters to effectively carry out your parallel query strategy. The Foglight Cartridge for Sybase assists your investigation into parallel query limits and adjustments based on availability of threads with the *SYBM_ProcessesParallel* graph view.

A high value for the “number of plans altered” can indicate that the number of worked processes configured were not adequate for the Adaptive Server to complete the queries using the degree of parallelism determined by the optimizer.

Measure this against the number of parallel queries attempted to determine if the number of plans altered is an issue.

Use the max threads and the users with threads values to help size your environment needs for parallel queries.

Investigating Procedure Cache

The procedure cache is used for stored procedures, triggers, and short-term memory needs such as statistics and query plans for parallel queries.

If more than one user uses a procedure or trigger simultaneously, there will be multiple copies of it in cache. If the procedure cache is too small, a user trying to execute stored procedures or queries that fire triggers receives an error message and must resubmit the query. Space becomes available when unused plans age out of the cache.

The Foglight Cartridge for Sybase assists your investigation of Procedure Cache usage patterns. This information helps a DBA tune the size of memory allocated to the procedure cache in an efficient manner.

To start with the *SYBM_ProcedureCache* graph:

- The *SYBM_ProcedureCache* graph view displays the number of procedure load, requests, writes, and stalls encountered during the collection period.

On a production server, you want to minimize the procedure reads from disk. When a user needs to execute a procedure, Adaptive Server should be able to find an unused tree or plan in the procedure cache for the most common procedures. The percentage of times the server finds an available plan in cache is called the cache hit ratio. Keeping a high cache hit ratio for procedures in cache improves performance.

Investigating Resource Utilization

Foglight lets you track several key resource utilizations to help you plan Adaptive Server capacity.

To start with the *SYBM_ResourceOverview* graph:

- Open the *SYBM_ResourceOverview* graph view which shows the utilization percentage for locks, connections, and memory.

To view more detailed information:

- Drill down on the Connections (%) line to see the *SYBM_ResourceConnections* graph view. This view tracks used and free connections.

- Drill down on the Locks (%) line to see how many locks are used and free in the *SYBM_ResourceLocks* graph view. See [Investigating Locking](#) for more detail information on tracking Sybase Adaptive Server Locking effectiveness and how to handle configuring the number of locks on an Adaptive Server.

You can increase the “number of user connections”. If your Adaptive Server is running out of user connections, use the *SP_configure* system procedure to increase their number. Take note of the *@@max_connections* global variable to determine how many connections the server supports. Adding additional engines increases the *@@max_connections* value and, in turn, you can increase the number of user connections. Be aware that decreasing the number of engines requires that the number of user connections be lowered as well. Both of these parameters are not dynamic and require a reboot of the system. Also, increasing the number of connections has a direct impact on memory allocation for the Adaptive Server. Review this in the Sybase system administration guide for full details.

You can use the *SYBM_Processes* graph view to see the distribution of the connections on your Adaptive Server (idle, active, blocked, etc.).

- Drill down on the Logical Memory (%) line to the *SYBM_ResourceMemory* graph view to determine current allocation for physical and logical memory.

Investigating Sybase Performance

Performance is the measure of efficiency for an application or multiple applications running in the same environment.

The Foglight Cartridge for Sybase measures efficiency by performing the following tasks:

- Collects data about a Sybase Adaptive Server instance.
- Analyzes and converts the raw data into useful information.
- Compares data against a set of Foglight rules that you can customize to look for current or potential problems.
- Alerts designated individuals when a problem is identified.

To start with the *SYBM_Overview* graph:

- Open the *SYBM_Overview* graph view which displays an overview of the system performance of the Sybase Adaptive Server. Specifically, it shows the percentage of time the server was busy with User Requests, the percentage of time spent performing I/O requests, lock contention and cache hit rates for data and procedure caches along with user connection activity.

To inspect Adaptive Server CPU utilizations:

- Drill down on the CPU - Busy (%) line to investigate how busy all Adaptive Server engines were during the time that the CPU was available to Adaptive Server. This data is displayed in the *SYBM_EngineSummary* graph view. See [Investigating CPU Capacity](#) for additional related information in Foglight.

To Inspect Adaptive Server Device I/O:

- Drill down on the CPU - I/O (%) line to obtain a summary of I/O metrics for all Sybase logical devices during the collection period. This data is displayed in the *SYBM_DiskIOSummary* graph view. See [Investigating Disk Performance](#) for additional related information in Foglight.

To Inspect Adaptive Server Cache effectiveness:

- Drill down on the Data Cache Hit (%) line to investigate the number of hits and misses for all caches during the collection period. This data is displayed in the *SYBM_DataCacheSummary* graph view. See [Investigating Data Cache Performance](#) for additional related information in Foglight.

To Inspect Adaptive Server Procedure Cache effectiveness:

- Drill down on the Procedure Cache (%) line to investigate the number of Requests and Loads for the Procedure cache during the collection period. This data is displayed in the *SYBM_ProcedureCache* graph view. See [Investigating Procedure Cache](#) for additional related information in Foglight.

To Inspect locking:

- Drill down on the Lock Contention (%) line to investigate the total number of deadlocks, locks timed out, etc. This data is displayed in the SYBM_LockSummary graph view, see [Investigating Locking](#) for additional related information in Foglight.

To Inspect User connection activity:

- Drill down on the Connections Used (%) line to investigate the number of active, idle and blocked connections. This data is displayed in the SYBM_Processes graph view.

Investigating the User Log Cache

The Foglight Cartridge for Sybase assists your investigation on the proper sizing of the user log cache.

To start with the SYBM_UserLogCache graph:

- Open the SYBM_UserLogCache graph view which displays the percentage of times the user log cache was flushed because it was full.

A high user log cache value indicates that the Adaptive Server is flushing the User Log Cache (ULC) more than once per transaction, negating some performance benefits of user log caches. If the percentage is greater than 20%, consider increasing the size of the user log cache size parameter.

Increasing the ULC size increases the amount of memory required for each user connection, so you do not want to configure the ULC size to suit a small percentage of large transactions.

About the Sybase_RS Agent

The Sybase_RS Agent monitors the Replication Server on the host it is deployed on. The agent collects internal performance and availability metrics by connecting to the Replication Server directly and from its related RSSD database via a connection to the respective Adaptive Server. The agent monitors Replication Server health, internal threads, partitions (space and status), connections, and exceptions. It also scans the external error log of the Replication Server looking for any default or user defined errors.

In addition, the Sybase_RS Agent also gathers availability information of the Replication Agent (RepAgent) from the primary sites (Adaptive Servers and databases) connected directly to this Replication Server as well as latency information from the Replicate Sites (Adaptive Servers and databases) controlled directly from this Replication Server.

Sybase_RS Agent Properties

When an agent connects to the Foglight Management Server, it is provided with sets of properties that it uses to configure its correct running state.

Each agent is provided with a combination of two types of properties: agent properties and shareable properties.

Default versions of these properties are installed with the Foglight for Sybase. However, you can edit the default shareable and agent properties, configure agent properties that apply only to a specific agent instance, and create edited clones of shareable properties that are used by a subset of the agents of a certain type.

For more information about working with agent properties, see the *Foglight Administration and Configuration Guide*.

The Sybase_RS Agent is shipped with default properties that can be modified to suit your system requirements.

Agent properties described in this section comprise:

- [Setting Connection Details](#) on page 97
- [Setting Data Management Properties](#) on page 98

- [Setting Configuration Properties](#) on page 100
- [Setting Data Retrieval Properties](#) on page 101

To modify agent properties:

- 1 Ensure that the navigation panel on the left is open.
To open the navigation panel, click the right-facing arrow on the left ►.
- 2 Open the dashboard that lets you navigate to the agent properties by completing one of the following steps:
 - On the navigation panel, under Dashboards, click **Administration > Agents > Agent Properties**.
In the Agent Properties dashboard, in the Namespace > Type pane, select **Legacy > Sybase_RS**.
 - On the navigation panel, under Dashboards, click **Administration > Agents > Agent Status**.
In the Agent Status dashboard, select the instance of the Sybase_RS Agent whose properties you want to modify and click **Edit Properties**.

A list of agent properties appears in the **Sybase_RS** pane.

Figure 48. Agent Properties List

The screenshot shows a web-based form for configuring agent properties. It is divided into two main sections: **Connection Details** and **Data Management**.

Connection Details: This section contains several input fields for configuring the agent's connection to the database. The fields are: Replication Server Home Directory (\$SYBASE), Replication Server Name, Replication Server Login Name, Replication Server Login Password (masked with asterisks), RSSD Login, RSSD Password (masked with asterisks), Primary ASE (RepAgent) Login Name, Primary ASE (RepAgent) Login Password (masked with asterisks), Replicate ASE (Latency) Login Name, Replicate ASE (Latency) Login Password (masked with asterisks), Maximum Connection Retries (set to 3), and Connection Timeout (set to 5).

Data Management: This section contains settings for data management. It includes a dropdown menu for 'Purge Data After' (set to 'Purge_Data'), an 'Edit' button, a 'Clone' button, and a 'Delete' button. Below this, there is a 'Monitor Availability Data' section with radio buttons for 'True' (selected) and 'False'. At the bottom, there is an 'Availability Sample Frequency (seconds)' field set to 60.

The position of the Properties pane depends on the dashboard you used to access agent properties. If you used the Agent Properties dashboard, the Properties pane appears to the right of the Namespace > Type pane in the display area. If you used the Agent Status dashboard, the Properties pane appears across the display area.

Setting Connection Details

Use the Connection Details set of agent properties for a Sybase_RS Agent to designate the type of connection properties to use.

NOTE: You do not need to create any new OS users. The Sybase RS agent only requires the database login to function correctly.

To set the connection details:

- 1 Go to the **Connection Parameters** set of properties.

- 2 Enter the location of the interfaces or the *ini/sql.ini* file, in the **Replication Server Home Directory (\$SYBASE)** box.
 - i** | **NOTE:** It is recommended that the Sybase_RS Agent be deployed on the server where the Replication Server is running.
If the Sybase_RS Agent is deployed on a remote server then the user needs to enter the Sybase Home directory for the remote server where the Sybase_RS Agent is deployed.
- 3 Enter the name of the Replication Server to be monitored, in the **Replication Server Name** box.
- 4 Enter a login name that provides access to the replication server, in the **Replication Server Login Name** box.
 - i** | **NOTE:** The login needs to be a user of the replication server with no special permissions.
- 5 Enter the current password associated with the Replication Server login name, in the **Replication Server Login Password** box.
- 6 Enter the RSSD login name in the **RSSD Login** box.
 - i** | **NOTE:** The login needs to be a database user in the *rssd* database and a member of *rs_systabgroup*.
- 7 Enter the current RSSD password associated with the RSSD login name, in the **RSSD Password** box.
- 8 Enter the primary ASE login name in the **Primary ASE (RepAgent) Login Name** box.
 - i** | **NOTE:** The login needs to be a database user in the Primary database.
This login should also have the *replication_role* assigned to it in order to execute the *sp_help_rep_agent* command.
If you have a warm standby server for the Primary ASE, and you want to monitor the warm standby server, ensure you create the primary login name inside the warm standby server.
- 9 Enter the current primary ASE password associated with the primary ASE login name, in the **Primary ASE (RepAgent) Login Password** box.
- 10 Enter the replicate ASE login name in the **Replicate ASE (Latency) Login Name** box.
 - i** | **NOTE:** The login needs to be a database user in the Replicate database with select permission on *rs_lastcommit* in the Replicate databases.
If you have a warm standby server for the Replicate ASE, and you want to monitor the warm standby server, ensure you create the replicate login name inside the warm standby server.
- 11 Enter the current replicate ASE password associated with the replicate ASE login name in the **Replicate ASE Latency) Login Password** box.
- 12 Enter the number of times you want the Sybase_RS Agent to attempt to connect to the Sybase Replication Server, in the **Maximum Connection Retries** box.

This is the maximum number of retries the Sybase_RS Agent attempts when it encounters an error with the Adaptive Server before the agent stops.

The default is three retries.
- 13 Enter number of seconds allowed for an individual connection attempt, in the **Connection Timeout** box.

The connection aborts after waiting to connect for this time out value.

The default is five seconds.
- 14 Click **Save**.
- 15 Go to [Setting Data Management Properties](#) to set the sample frequency.

Setting Data Management Properties

Use the Data Management set of agent properties to set the sample frequency.

To set the data management properties:

- 1 Ignore the **Purge Data After** box. This field is not functional.
- 2 Click the **Monitor Availability Data** radio button, True or False, to turn the data filtering on or off. When this field is set to true, the agent monitors Replication Server availability.
The default value is true.
- 3 Enter the sample frequency (in seconds) to monitor the availability data, in the **Availability Sample Frequency (seconds)** box.
The Availability Sample Frequency is the rate at which the agent tests for the *RepServer* connectivity, up/down status, etc. The Sybase_RS Agent sample frequency is controlled at the agent level.
The default is 60 seconds.
- 4 Click the **Monitor Performance Data** radio button, True or False, to turn the monitoring of performance data on or off. When this field is set to true, the agent monitors the Replication Server performance data.
The default value is true.
- 5 Enter the sample frequency (in seconds) to monitor the performance data, in the **Performance Sample Frequency (seconds)** box.
The sample frequency, or collection rate, is the number of seconds between the end of a collection period and the start of the next. The Sybase_RS Agent sample frequency is controlled at the agent level.
The default is 300 seconds.
- 6 Click the list from the **Performance List** that you want to update. Select a list for the Performance Foglight Tables to monitor.
The default list name is *Performance_List*.
- 7 Click **Edit**.
A secondary property list appears.
- 8 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
- 9 Click the dialog box **Save** to save your entries and close the dialog box.
- 10 Click the **Monitor Space Trend Data** radio button, True or False, to monitor the Replication Server space trend data. When this field is set to true, the agent monitors the Replication Server space trend data.
The default value is false.
- 11 Enter the sample frequency (in minutes) to monitor the space trend data, in the **Space Trend Sample Frequency (minutes)** box.
The sample frequency, or collection rate, is the number of minutes between the end of a collection period and the start of the next. The Sybase_RS Agent sample frequency is controlled at the agent level.
The default is 60 minutes. The minimum value that can be entered is 15 minutes.
- 12 Click the list from the **Space Trend List** that you want to update. Select a list for the space trend tables to monitor.
The default secondary list name is *Space_Trend_List*.
- 13 Click **Edit**.
A secondary property list appears.
- 14 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
- 15 Click the dialog box **Save** to save your entries and close the dialog box.
- 16 Click **Save**.

- 17 Go to [Setting Configuration Properties](#) to filter system error messages.

Setting Configuration Properties

Use the Configuration set of agent properties to filter the collection of system error messages by including or excluding collection.

The Sybase_RS Agent collects specific system error messages, which are defined in the Message Trap List, and then stored in the *RepServer* Error Log.

To set the error message collection properties:

- 1 Click the list from the **Latency Include/Exclude List** that you want to update.


 **NOTE:** The Latency Exclusion List is not auto discovered when the agent is started. You must enter the Replicate ASE/database that you do not wish to monitor in the exclude list.
If you do not want to monitor latency for a warm standby server, then you must enter the warm standby server in the Latency Exclude List in the agent property.
- 2 Click **Edit**.
A secondary property list appears.
- 3 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
 - a **Replicate ASE**—Enter the Replicate ASE to exclude.
 - b **Replicate Database**—Enter the Replicate Database to exclude.
 - c **Exclude** radio button—Select true to exclude the Replicate ASE.
 - d **Critical Threshold Limit (s)**—Enter the preferred threshold limit for latency
- 4 Click the dialog box **Save** to save your entries and close the dialog box.
- 5 Click the list from the **Primary ASE Exclude List** that you want to update.
- 6 Click **Edit**.
A secondary property list appears.
- 7 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
- 8 Click the dialog box **Save** to save your entries and close the dialog box.
- 9 Click the list from the **Message Trap List** that you want to update.
The Log Filters list allows you to enter the log message type, search pattern, and whether to include or exclude. The Sybase_RS Agent records these messages in an event log.
The Sybase_RS Agent default log filter name is *LogFileFilter*.
- 10 Click **Edit**.
A secondary property list appears.
- 11 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
 - a **Log Message Type**—Allows you to specify the message type to be filtered.

Table 75. Log Message Types

Message Type Value	Message Type	Description
F	Fatal Error	This is a serious error which could cause the Replication Server to exit.
N	Internal Error	These errors are caused by anomalies in the Replication Server software.
E	Error	An error that does not prevent further processing, such as a site that is unavailable.
W	Warning	These are warnings about a condition that has not yet caused an error, but may require attention.
H	Thread Died	A Replication Server thread has unexpectedly shut down

- b **Search Pattern**—Indicates the regular expression pattern used to match the log message.
- c **Include Message**—If selected, it includes the particular Log Message Type (E, F, or W).
- d **Exclude Message**—If selected, it excludes the particular Log Message Type (E, F, or W).

12 Click the dialog box **Save** to save your entries and close the dialog box.

13 Click **Save**.

14 Go to [Setting Data Retrieval Properties](#) to set the maximum number of rows returned.

Setting Data Retrieval Properties

To set the data retrieval properties:

- 1 Click the maximum rows list from the **Max Row List** that you want to update.
The default list name is *Max_Row_List*. This list can be renamed to better describe the changes made.
- 2 Click **Edit**.
A secondary property list appears.
- 3 Add an entry to the list by clicking **Add new row**.
Fill in the dialog box fields.
 - a **Table**—The table to limit the rows returned.
 - b **Maximum Rows Returned**—Enter the maximum rows to return.
The maximum value for rows returned is 100.
- 4 Click the dialog box **Save** button to save your entries and close the dialog box.
- 5 Click **Save**.

Sybase_RS Agent Investigations

This section describes the following investigations:

- [Investigating Connections](#)
- [Investigating Health and Availability of the Replication Server](#)
- [Investigating Threads Down](#)
- [Investigating Stable Queues](#)
- [Investigating Partition Status](#)

- [Investigating Partition Space](#)
- [Investigating Error Log Messages](#)
- [Investigating Latency](#)

Investigating Connections

The Foglight Cartridge for Sybase assists your investigation into replication server connections with the *RS_ConnectionSummary* graph view and *RS_ConnectionDetail* table view.

Connections can be either a route to another Replication Server, or a data server connection to an Adaptive Server managed by the monitored Replication Server. The *RS_ConnectionSummary* graph view gives a summary of each connection state to provide an overall view of all the connections managed by the monitored Replication Server.

The *RS_ConnectionDetail* table view provides more in-depth information on each of the connections that are returned to Foglight. Review the Sybase replication Server Administration guide for a description of each state and the procedures on how to address them.

Foglight also raises an alert to notify the Replication Server administrator when a connection is returned to Foglight.

To start with the *RS_ConnectionSummary* graph view:

- Open this graph directly or from the *RS_Internal_Overview* graph view after clicking on closed connections.
- The *RS_ConnectionSummary* graph view shows each connection status type and the number of connections associated with each state.

To view more detailed information:

- Drill down on the Connecting line to view detailed information on Sybase Replication Server connections with the *RS_ConnectionDetail* table view. This view shows more specific information on each connection returned to Foglight.

Investigating Health and Availability of the Replication Server

The Sybase_RS Agent determines the health and availability of the monitored Replication Server several ways. It first attempts to open a new connection to the Replication Server each poll cycle. Once a connection has been established, it requests a confirmation that it is responding. Lastly, the agent requests the health status of the Replication Server.

Each of the steps is recorded by the Foglight agent and returned to the Foglight Management Server. Use the views below to investigate the status and health of the monitored Replication Server.

Foglight raises an alert if the health or availability of the Replication Server changes to an undesirable state.

To start with the *RS_State_Overview* graph view:

- Open the *RS_State_Overview* graph view to display a graphical representation of the Replication Servers health and availability.

To view more detailed information related to the *RS_State_Overview* graph view:

- Drill down on the Health Status line to view each discrete Health state with the *RS_Health* graph view.
- Drill down on the Health Status line to view a textual description of each state in a tabular format.
- Drill down on the Connect column to get to the *RS_Availability* graph view and inspect a graphical chart of the availability states of the monitored Replication Server.

To view more detailed information related to the RS_Availability graph view:

- Drill down on the Connect line to view a textual description of each state in a tabular format.

Investigating Threads Down

The Foglight Cartridge for Sybase assists your investigation into Sybase Replication Server internal status. Use the *RS_ThreadsDownSummary* table view to get an overview of the overall state of the critical internal components of the Replication Server. The *RS_ThreadsDown_Detail* table view provides more detailed information on each thread that is in a down state as indicated by the monitored Replication Server.

Use the thread name and the related information to identify the component. Check the Replication Server error log (use the *RS_ErrorlogDetail* table view to see if there is additional debug information available. If the component in a down state is a route or connection then connect to the respective Adaptive or Replication Server to further diagnose the issue.

Foglight also raises an alert to notify the Replication Server administrator when any down thread is detected.

To start with the RS_ThreadsDownSummary table:

- Open the *RS_ThreadsDownSummary* table view to display thread state summary totals of the Replication Server.

To view more detailed information:

- Drill down on the Name column to view information on Replication Server threads that are down with the *RS_ThreadsDown_Detail* table view.

Investigating Stable Queues

The Foglight Cartridge for Sybase assists your investigation into replication server stable queues with the *RS_SQM_Trend* table view, *RS_SQM_Activity* graph view, *RS_SQM_QueueSize* graph view, and the *RS_SQM_TrailSize* graph view.

Stable Queues are Store-and-forward queues where Replication Server stores messages destined for a route or database connection. Messages written into a stable queue remain there until they can be delivered to the destination Replication Server or database. Replication Server builds and allocates space for the stable queues from the disk partitions created for the Replication Server. The most referenced types of queues are inbound (transactions from a Primary Server) and outbound (transaction messages for another Replication Server via a route or to a direct Data Server connection - DSI).

Foglight allows a Replication Server administrator to monitor the overall performance of the stable queues over time by gathering the respective metrics and allowing these metrics to be viewed in pertinent graphs.

Monitoring the stable queues can be a measure of how much activity or volume is passing through the Replication Server. The Replication Server administrator can use this information to understand the performance profile of the monitored Replication Server. This information also serves as a tool for capacity planning.


The trending of SQM performance can be turned on or off via the Sybase_RS agent properties. Refer to the [Sybase_MDA Agent Properties](#) for this agent.

To start with the RS_SQM_Trend table view:

- Open the *RS_SQM_Trend* table view to display activity totals of SQM threads.

To view more detailed information:

- Drill down on the Messages Written or Messages Read line to view SQM messages read and written within the last polling period with the *RS_SQM_Activity* graph view.

 **TIP:** This is a good indication of the volume of transactions that the Replication Server is handling.

- Drill down on the Queue Size line to view SQM queue size with the *RS_SQM_QueueSize* graph view.

- Drill down on the Data Not Processed line to view SQM queue size with the RS_SQM_TrailSize graph view.
Use this to watch if the Stable Queue is keeping up with the incoming transactions. An increase in the size indicates the stable queue is falling behind in forwarding the transactions to their destination. Research downstream of the stable queue to find the cause.

Investigating Partition Status

The Foglight Cartridge for Sybase assists your investigation into partition status with the *RS_PartitionStatusSummary* graph view and the *RS_PartitionStatusDetail* table view.

Partitions in Replication Server are a raw disk partition or operating system file that Replication Server uses for stable queue storage.

Foglight monitors the Replication Server for any partition that is not in an online state. The *RS_PartitionStatusSummary* graph view provides a quick overview of partition status on the monitored Replication Server. Foglight returns more detailed information on any partition that is not online.

Foglight raises an alert if any partition is found in any state other than online.

Review the Replication Server error log for any additional information about an off-line partition. Check the location of the physical device and see if there are any reported errors on the physical volume.

To start with the *RS_PartitionStatusSummary* graph view:

- Open this graph directly, or from the RS_Internal_Overview **graph** view after clicking Partitions Offline.

To view more detailed information:

- Drill down on the Offline value to view detailed status information for partitions with the *RS_PartitionStatusDetail* table view.

Investigating Partition Space

The Foglight Cartridge for Sybase assists your investigation into partition space with the *RS_PartitionSpaceUsedSummary* table view and the *RS_PartitionSpace_Trend* table view. Partitions in Replication Server are a raw disk partition or operating system file that Replication Server uses for stable queue storage.

Foglight monitors the replication Server for overall Partition Space that is allocated to the stable queues. The *RS_PartitionSpaceUsedSummary* table view gives a quick overview of partition space on the monitored Replication Server.

The *RS_PartitionSpace_Trend* table view provides information to track individual partition space allocation over time. Trending must be turned on in order for this information to be available. Review the agent properties for the Sybase_RS Agent for more information on how to configure this.

To start with the *RS_PartitionSpaceUsedSummary* table view:

- 1 Open this view directly or from the RS_Internal_Overview graph view after clicking on Space Used Percentage. The *RS_PartitionSpaceUsedSummary* table view shows a breakdown of the space used by all partitions on the Replication Server. The Space percent free should be a good indicator when it may be needed to increase the amount of partition space needed by the Replication Server. If the Replication Server runs out of space to store messages, there is a negative impact on the Replication Server.
- 2 Use the “create partition” command to make more partition space available to the Replication Server. If the partition is to be an operating system file, then before using the “Create partition” command, you must first create the file (in UNIX use ‘touch’).

To view more detailed information:

- Drill down on the Total Size (Mb) line to view detail information on each partition with the *RS_PartitionSpace_Trend* table view.

The trending of Partitions space metrics can be turned on or off via the [Sybase_MDA Agent Properties](#). Refer to the [Setting Data Management Properties](#) for this agent.

Investigating Error Log Messages

The Foglight Cartridge for Sybase assists your investigation into the status of the monitored Replication Server by returning pertinent messages Foglight finds in the Replication Server error log.

Replication Server errors are recorded in the Replication Server error log. A Replication Server error log contains informational, warning, thread-terminated, fatal, and internal error messages.

To start with the *RS_ErrorlogSummary* table view:

- Open the *RS_ErrorlogSummary* table view to display summary information from the Replication Server error log. This view provides summary totals for each type of message in the error log found in the last poll cycle for the Sybase_RS Agent.

To view more detailed information:

- Drill down on the Informational column to view detailed row information from the Replication Server error log with the *RS_ErrorlogDetail* table view.

Investigating Latency

The Foglight Cartridge for Sybase assists your investigation into latency with the *RS_LatencyTabular* table view and the *RS_Latency_Graph* view.

Latency is how long a transaction takes from being committed on the Primary Server to being committed on the Replicated Data Server. Foglight connects to each replicated server controlled by the monitored Replication Server. Foglight retrieves the latency time from each Replicated database for each Primary site for the last committed transaction from that site and returns this information to Foglight.

If a transaction from that Primary site has not been committed in more than 24 hours, and if the last commit time for that transaction is earlier than the last poll cycle of the Foglight agent, this transaction is marked as .stale. information for that Primary site.

Foglight needs to be able to connect to each Replicated Server to collect this information. In addition, Foglight needs to be able to connect to each Primary server in order to determine the time zone factor when calculating the latency times.

To start with the *RS_Latency_Tabular* table view:

- Open the *RS_LatencyTabular* table view to display the Replication Server's latency information in a table format. There should be one row for each Primary database -Replicate database combination found in each Replicate Database that the monitored Replication Server connects to directly.

To view more detailed information:

- Drill down on the Latency(s) column to view the latency for the Replication Server in a graphical format with the *RS_Latency_Graph* view.
- Drill down on the Latency (s) line to view the Replication Server's latency information in a tabular format with the *RS_LatencyTabular* table view.

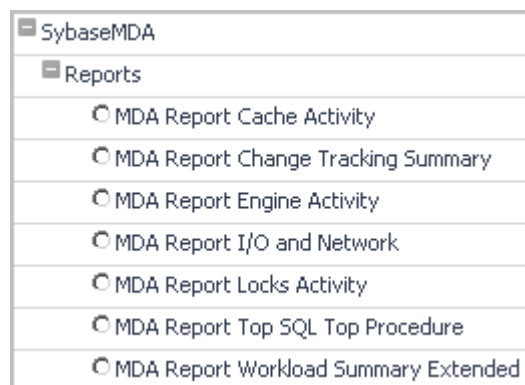
Generating SybaseMDA Reports

Foglight for Sybase Foglight for Sybase includes a report generation and scheduling ability. This allows you to create reports using a set of predefined templates to report on the various aspects of your database environment. For more information about reports in Foglight for Sybase, how to generate and view them using the browser interface, see the *Foglight for SybaseUser Help*.

In general, all SybaseMDA reports begin with a summary page that shows overall statistics appearing in a combination of bar gauge and time plot charts. The summary page is followed by a graph showing the performance of a system wait event class and a list of the related events that occur during the selected time range.

The Foglight for Sybase report templates can be found in the following location, when listed by module:
SybaseMDA >Reports.

Figure 49. Reports



The cartridge comes equipped with the following report templates:

- [Cache Activity Report](#) (see page 106)
- [Change Tracking Summary Report](#) (see page 107)
- [Engine Activity Report](#) (see page 108)
- [I/O and Network Activity Report](#) (see page 108)
- [Locks Activity Report](#) (see page 109)
- [Top SQL Top Procedure Report](#) (see page 109)
- [Workload Summary Extended Report](#) (see page 109)

Cache Activity Report

The Cache Activity report contains details about the data and procedure cache metrics monitored by a selected SybaseMDA agent for a given time range. It contains the following sections:

- *Cache Activity Summary* shows the cache activity represented by the related metrics in bar charts. This section contains the number of named data caches, their total size, the size of the procedure cache, the maximum available memory, the available physical memory, and the data and procedure cache hit rates. The Data Cache and Procedure Cache hit rates show hit rate averages for the selected time period, along with the related severity level. Different-colored vertical lines in the graphs indicate threshold levels.
- *Waiting for a memory or buffer summary* displays a graph that shows average wait times per sample for a specific class, and a table that lists all the events related to that class that occur during the selected time range.
- *Cache Graphs Summary* contains graphs that indicate the data cache hit rate average, hits, misses and dirty buffers, as well as the procedure cache hit rate average, requests, loads, and writes.
- *Metrics Statistics Summary* shows the a table listing the metrics related to the cache activity. For each metric, the table shows their previous and current values, the baseline minimum and maximum values, and the deviation of the current value from the baseline. The current values appear in bold text for better visibility. Metric groups appear in alternate text color for the same reason.

NOTE: Shorter data sampling period can cause some values to be zero. This is because the retention policy process can clean the data. With periods longer than 30 minutes this is not the case.

- *Data Cache Details* show the metrics for each named cache, such as the size in MB, average hit rate, and others.
- *Data Pools Details* show the various metrics for each pool in cache, such as the I/O size, pool size, and others.

Change Tracking Summary Report

The Change Tracking Summary report helps you understand the impact of various changes on the Sybase Server during a given time range. The report shows two periods: one before the change and one after the change. It shows this information in graph and table form. The cartridge can become aware of changes in any of the following ways:

- It is automatically notified of any configuration parameter changes
- It can become aware of other types of changes using the following stored procedure:

```
exec sp_fgl_add_event "event description"
```

i | IMPORTANT: Run this procedure as the Foglight Sybase user.

This report takes two special input parameters:

- *stepMinutes*: Set this parameter to the period of time in minutes that passes before and after the
- *beforeAfterMinutes*: Set this parameter to the offset in minutes from the change time. This value is very important because it can eliminate a possible negative impact of the retention policy cycle in case the change happens in the middle of retention policy period. In this case, the period average is calculated using the values from the opposite sides of the change. The default retention policy period is 15 min. For example:

The change occurs at 10:00. *stepMinutes* is set to 15 and *beforeAfterMinutes* is set to 10 minutes. As a result, two periods appear: 9:35-9:50 and 10:10-10:25.

The report contains a table of contents and sections listed below. In addition to the list of the reports sections, the table of contents also shows all change events that occurred during the selected time period.

- *Instance Information* shows the information about the selected database server, such as its health, user connections, engine, cache, databases running, and others.
- *Change Tracking Events Compare Graphs* list all changes and contain several graphs. The first graph displays the count of changes over time, while the second graph shows the wait events which shows the impact of the changes on the monitored system. For example, an index creation change can help reduce the wait events in general. Other graphs in this section also illustrate how the changes affect the system, such as their impact on the engine and IO activity, and others.
- *Important Metrics Statistics Summary* lists various metrics, grouped into several collections, such as *Engines*, *Processes*, *SQLs*. For each metric, it gives its value before the change and after the change, along with the respective baseline values for the period after the change. It can help understand the impact of one metric change on another. For example, a dramatic decrease in cache misses can be caused by an increased use of APF reads.
- *System Wait Classes Compare* shows in list form the wait classes for the period before the change and after.
- *System Wait Events* shows in list form the system wait events for the period before the change and after.
- *Change Tracking Events Details Compare* contains additional details about the change tracking events.

Engine Activity Report

The Engine Activity report shows performance details for Sybase engine operations. It contains the following sections:

- *Engine Activity Summary* indicates the number of Sybase engines and connections. It also shows in graph form the average percentage of times the engine, engine I/O and engine CPU were busy over a selected time range. The Total Busy Percent and CPU and I/O busy percentages show the hit rate averages for the selected time period, along with the related severity level. Different-colored vertical lines in the graphs indicate threshold levels.
- *Schedule (CPU) Waits Summary* shows in graph form the schedule (CPU) waits and a table that lists all the related events that occur during the selected time range.
- *Engines Details* lists all Sybase engines and shows the average percentage of times each engine was busy, including I/O and CPU busy times, among other metrics.
- *Metrics Statistics Summary* shows the a table listing the metrics related to the engine activity. For each metric, the table shows their previous and current values, the baseline minimum and maximum values, and the deviation of the current value from the baseline. The current values appear in bold text for better visibility. Metric groups appear in alternate text color for the same reason.

i | NOTE: Shorter data sampling period can cause some values to be zero. This is because the retention policy process can clean the data. With periods longer than 30 minutes this is not the case.

- *Metrics Statistics Summary* lists various metrics, grouped into several collections, such as *Engines*, *Processes*, *SQLs*. For each metric, it gives its average values, the expected range, and the deviation (if applicable).
- *Engine and Processes Summary* contains several graphs with engine and process summary information along with the metric baseline.

I/O and Network Activity Report

The I/O and Network Activity report shows network activity details as well as various device I/O metrics. It contains the following sections:

- *I/O Activity Summary* shows the numbers of I/O devices and database, and various I/O read and write rates. It also lists various metrics, grouped into several collections, such as *I/O* and *Pages*. For each metric, it gives its average values, the expected range, and the deviation (if applicable).
- *I/O Waits Summary* shows in graph form the amount of time spent on disk write and read waits over a selected time range. It also lists class-related events and shows wait-related metrics for each event. There are two wait classes that appear in this section: *Disk Read Wait* with the class ID of 2, and *Disk Write Wait*, with the class ID of 3.
- *Device I/O Summary* shows in graph form the rate of device reads, writes, and Asynchronous Prefetch Activity (APF) reads, and then the same information in comparison with the expected baseline.
- *Metrics Statistics Summary* lists various metrics, grouped into several collections, such as *Caches*, *Disk I/O*, *Network*. For each metric, it gives its average values, the expected range, and the deviation (if applicable).
- *Device I/O Details* lists all I/O devices, and for each device it shows its logical name, physical name, total I/O amount, the counts and averages of the device writes, reads, and APF reads, and other metrics.
- *Temporary Database Spaces Usage* shows the space usage for temporary databases, including the total amount of database space, and the available space. The same information appears in graph form, for a given time range.
- *Network Activity Summary* shows the average numbers of bytes received and sent per data packet, and their rate. The rates also appear in graph form for a selected time range.
- *Network Waits Summary* shows in graph form the amount of time spent on network input and output waits over a selected time range. It also lists class-related events and shows wait-related metrics for each event.

- *Network I/O Details* lists various metrics, grouped into several collections, such as *Packets Received* and *Packets Sent*. For each metric, it gives its average values, the expected range, and the deviation (if applicable).

Locks Activity Report

The Lock Activity report shows lock activity details as well as various lock metrics. It contains the following sections:

- *Locks Waits Summary* shows lock wait statistics. It contains two graph views that display the amount of time spent on lock waits and average lock waits per connection during a given time range. It also lists class-related events and shows wait-related metrics for each event.
- *Locks Summary* shows the average counts of blocks and deadlocks, and also the average wait times, all over a selected time range.
- *Locks Details* lists various metrics, grouped into several collections, such as *Connections*, *Locks (Max)*, *Locks (Average)*, *Locks (Total)*, and *Deadlocks*. For each metric, it gives its average values.
- *Blocked Tree* lists information on blocked trees, for each entry, shows the SPID, login name, database name, application, locked table, SQL text, and other metrics.

Top SQL Top Procedure Report

The Top SQL Top Activity report provides detailed information about the SQL activities that utilized the highest amount of processing time or the highest amount of wait time over a given time range. It contains the following sections:

- *Important Metrics Statistics Summary* lists various metrics, grouped into several collections, such as *Engines*, *Processes*, *SQLs*. For each metric, it gives its value before the change and after the change, along with the baseline value for the period after the change. It can help understand the impact of one metric change on another. For example, a dramatic decrease in cache misses can be caused by an increased use of APF reads.
- *Locks Details* lists various metrics, grouped into several collections, such as *Connections*, *Locks (Max)*, *Locks (Average)*, *Locks (Total)*, and *Deadlocks*. For each metric, it gives its average values.
- *Top SQL* provides detailed information about the SQL statements that experienced the longest duration or highest number of total wait events during the specified time range. The Top SQL information is sorted by the Avg Duration (ms) column.
- *Top Procedure* provides detailed information about the specific SQL procedure that experienced the longest duration or highest number of total wait events during the specified time range. The Top Procedure information is sorted by the Avg Duration (ms) column.

This report template provides the Max Rows input parameter, which allows you to specify the number of rows to be shown in each SQL table.

Workload Summary Extended Report

The Workload Summary Extended report provides detailed information about all major aspects of your monitored environment. It contains the following sections:

- *Instance Information* shows the information about the selected database server, such as its health, user connections, engine, cache, databases running, and others.
- *Last Alarms* lists the alarms generated against the monitored system, and for each alarm, shows the time it was created, its duration, source name, and the alarm message.
- *Important Metrics Statistics Summary* lists various metrics, grouped into several collections, such as *Engines*, *Processes*, *SQLs*. For each metric, it gives its value before the change and after the change,

along with the baseline value for the period after the change. It can help understand the impact of one metric change on another. For example, a dramatic decrease in cache misses can be caused by an increased use of APF reads.

- *System Wait Classes* shows in graph form the amount of time spent on various system processes, such as system events, disk writes, and others, for the system wait classes, during a selected time period.
- *System Wait Events*
- *Engine Activity Summary* indicates the number of Sybase engines and connections. It also shows in graph form the average percentage of times the engine, engine I/O and engine CPU were busy over a selected time range.
- *Engines Details* lists all Sybase engines and shows the average percentage of times each engine was busy, including I/O and CPU busy times, among other metrics.
- *Engine and Processes Summary* contains several graphs, showing the average percentage of busy engine I/O and CPU times, and the counts of active and blocked processes, over a selected time range.
- *Locks Details* lists various metrics, grouped into several collections, such as *Connections*, *Locks (Max)*, *Locks (Average)*, *Locks (Total)*, and *Deadlocks*. For each metric, it gives its average values.
- *Blocked Tree* lists information on blocked trees, for each entry, shows the SPID, login name, database name, application, locked table, SQL text, and other metrics.
- *Cache Activity Summary* shows the number of named data caches, their total size, the size of the procedure cache size, the maximum available memory, the available physical memory, and the data and procedure cache hit rates.
- *Cache Graphs Summary* contains graphs that indicate the data cache hit rate average, hits, misses and dirty buffers, as well as the procedure cache hit rate average, requests, loads, and writes.
- *Data Cache Details* show the metrics for each named cache, such as the size in MB, average hit rate, and others.
- *Data Pools Details* show the various metrics for each named cache, such as the I/O size, pool size, and others.
- *I/O Activity Summary* shows the numbers of I/O devices and database, and various I/O read and write rates. It also lists various metrics, grouped into several collections, such as *I/O* and *Pages*. For each metric, it gives its average values, the expected range, and the deviation (if applicable).
- *I/O Waits Summary* shows in graph form the amount of time spent on disk write and read waits over a selected time range. It also lists class-related events and shows wait-related metrics for each event. There are two wait classes that appear in this section: *Disk Read Wait* with the class ID of 2, and *Disk Write Wait*, with the class ID of 3.
- *Device I/O Summary* shows in graph form the rate of device reads, writes, and Asynchronous Prefetch Activity (APF) reads, and then the same information in comparison with the expected baseline.
- *Device I/O Details* lists all I/O devices, and for each device it shows its logical name, physical name, total I/O amount, the counts and averages of the device writes, reads, and APF reads, and other metrics.
- *Temporary Database Spaces Usage* shows the space usage for temporary databases, including the total amount of database space, and the available space. The same information appears in graph form, for a given time range.
- *Network Activity Summary* shows the average numbers of bytes received and sent per data packet, and their rate. The rates also appear in graph form for a selected time range.
- *Network Waits Summary* shows in graph form the amount of time spent on network input and output waits over a selected time range. It also lists class-related events and shows wait-related metrics for each event.
- *Network I/O Details* lists various metrics, grouped into several collections, such as *Packets Received* and *Packets Sent*. For each metric, it gives its average values, the expected range, and the deviation (if applicable).
- *Top SQL* provides detailed information about the SQL statements that experienced the longest duration or highest number of total wait events during the specified time range. The Top SQL information is sorted by the Avg Duration (ms) column.

- *Top Procedure* provides detailed information about the specific SQL procedure that experienced the longest duration or highest number of total wait events during the specified time range. The Top Procedure information is sorted by the Avg Duration (ms) column.

This report template provides the Max Rows input parameter, which allows you to specify the number of rows to be shown in each SQL table.

Reference

This chapter contains reference information about views, rules, and data tables that are included with the cartridge. Read this chapter to find out details about these components.

Views

Foglight displays monitoring data in views that group, format, and display data. The main types are described below.

Dashboards are top-level views that contain lower-level views. The dashboards supplied with Foglight, as well as those created by users, are accessible from the navigation panel.

Lower-level views in Foglight can be added to dashboards or can be accessed by drilling down from a dashboard. They receive and display data directly from the Foglight Management Server or from other views. Some views filter or select data that appears in other views in the same dashboard. Some are tree views with expandable nodes for selecting servers, applications, or data.

This cartridge includes the following views:

- Sybase_MDA Agent SYBM Overview Graph View
- Sybase_MDA Agent SYBM OverviewNetwork Table View
- Sybase_MDA Agent SYBM ResourceOverview Graph View
- Sybase_MDA Agent SYBM ASE Availability Graph View
- Sybase_MDA Agent SYBM ASE Config Table View
- Sybase_MDA Agent SYBM ASE Info Table View
- Sybase_MDA Agent SYBM BenchmarkTimes Graph View
- Sybase_MDA Agent SYBM DBErrorLog Table View
- Sybase_MDA Agent SYBM DataCacheDetail Graph View
- Sybase_MDA Agent SYBM DataCachePoolInfo Table View
- Sybase_MDA Agent SYBM DataCachePools Graph View
- Sybase_MDA Agent SYBM DataCacheSummary Graph View
- Sybase_MDA Agent SYBM DbSpace Graph View
- Sybase_MDA Agent SYBM DbSpaceSegment Graph View
- Sybase_MDA Agent SYBM DbSpaceTransLog Graph View
- Sybase_MDA Agent SYBM DbSpaceTrendBar Graph View
- Sybase_MDA Agent SYBM DbSpaceTrendChart Graph View
- Sybase_MDA Agent SYBM DbStatus Table View
- Sybase_MDA Agent SYBM DiskIODetail Graph View
- Sybase_MDA Agent SYBM DiskIOSemaphores Graph View
- Sybase_MDA Agent SYBM DiskIOSummary Graph View

- Sybase_MDA Agent SYBM DiskSpaceTrendBar Graph View
- Sybase_MDA Agent SYBM DiskSpaceTrendChart Graph View
- Sybase_MDA Agent SYBM EngineInfoDetail Table View
- Sybase_MDA Agent SYBM EnginePerfDetail Graph View
- Sybase_MDA Agent SYBM EngineSummary Graph View
- Sybase_MDA Agent SYBM IndexUse Graph View
- Sybase_MDA Agent SYBM LockSummary Graph View
- Sybase_MDA Agent SYBM LockWaitTime Graph View
- Sybase_MDA Agent SYBM MetaDataCache Graph View
- Sybase_MDA Agent SYBM MetaDataCacheDetail Graph View
- Sybase_MDA Agent SYBM NetworkBytes Graph View
- Sybase_MDA Agent SYBM NetworkPackets Graph View
- Sybase_MDA Agent SYBM OS Cpu Graph View
- Sybase_MDA Agent SYBM OS CpuDetail Graph View
- Sybase_MDA Agent SYBM ProcedureCache Graph View
- Sybase_MDA Agent SYBM Processes Graph View
- Sybase_MDA Agent SYBM ProcessesParallel Graph View
- Sybase_MDA Agent SYBM ResourceConnections Graph View
- Sybase_MDA Agent SYBM ResourceLocks Graph View
- Sybase_MDA Agent SYBM ResourceMemory Graph View
- Sybase_MDA Agent SYBM SybaseConfigChanges Table View
- Sybase_MDA Agent SYBM TempDbSpace Graph View
- Sybase_MDA Agent SYBM TempDbSpaceSegment Graph View
- Sybase_MDA Agent SYBM TempDbSpaceTransLog Graph View
- Sybase_MDA Agent SYBM TopApplications Table View
- Sybase_MDA Agent Sybase MDA SYBM TopSQL Table View
- Sybase_MDA Agent SYBM TopUsers Table View
- Sybase_MDA Agent SYBM UserLogCache Graph View
- Sybase_MDA Agent SYBM UsersBlocked Table View
- Sybase_MDA Agent SYBM UsersBlocking Table View
- Sybase_RS Agent RS Internal Overview Graph View
- Sybase_RS Agent RS State Overview Graph View
- Sybase_RS Agent RS Availability Graph View
- Sybase_RS Agent RS Availability Text Table View
- Sybase_RS Agent RS ConnectionDetail Table View
- Sybase_RS Agent RS ConnectionSummary Graph View
- Sybase_RS Agent RS ErrorlogDetail Table View
- Sybase_RS Agent RS ErrorlogSummary Table View
- Sybase_RS Agent RS ExceptionDetail Table View
- Sybase_RS Agent RS Health Graph View

- Sybase_RS Agent RS HealthText Table View
- Sybase_RS Agent RS Information Table View
- Sybase_RS Agent RS LatencyTabular Table View
- Sybase_RS Agent RS Latency Graph View
- Sybase_RS Agent RS PartitionSpaceUsedSummary Table View
- Sybase_RS Agent RS PartitionSpace Trend Table View
- Sybase_RS Agent RS PartitionStatusDetail Table View
- Sybase_RS Agent RS PartitionStatusSummary Graph View
- Sybase_RS Agent RS RepAgent Recovery Table View
- Sybase_RS Agent RS RepAgentStatus Table View
- Sybase_RS Agent RS SQM Activity Graph View
- Sybase_RS Agent RS SQM QueueSize Graph View
- Sybase_RS Agent RS SQM TrailSize Graph View
- Sybase_RS Agent RS SQM Trend Table View
- Sybase_RS Agent RS ThreadsSummary Table View
- Sybase_RS Agent RS ThreadsDown Detail Table View

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit <https://www.quest.com/company/contact-us.aspx> or call +1-949-754-8000.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.

Table 1. List of third-party contributions

Component	License or acknowledgment