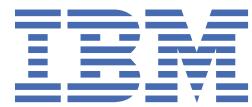


Monitoring Agent for Microsoft SQL Server
08.21.03.00

Reference



Note

Before using this information and the product it supports, read the information in [“Notices” on page 253.](#)

This edition applies to version 8.1.4.0.11 of the Monitoring Agent for Microsoft SQL Server and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2013, 2020.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Monitoring Agent for Microsoft SQL Server.....	1
Chapter 2. Dashboard.....	3
Default dashboard pages.....	3
Widgets for the Default dashboard pages.....	4
Custom views.....	32
Chapter 3. Thresholds.....	33
Predefined thresholds.....	33
Customized thresholds.....	41
Chapter 4. Attributes.....	43
Data sets for the monitoring agent.....	44
Attribute descriptions.....	49
MS SQL Additional Availability Database Details data set.....	49
MS SQL Additional Database Detail data set.....	51
MS SQL Analysis Services Data Mining data set.....	52
MS SQL Analysis Services MDX Processing data set.....	54
MS SQL Analysis Services Memory Statistics data set.....	59
MS SQL Analysis Services Rows Processing data set.....	69
MS SQL Analysis Services Storage Engine Query data set.....	71
MS SQL Audit Details data set.....	79
MS SQL Availability Database Details data set.....	81
MS SQL Availability Database Details In Cluster data set.....	84
MS SQL Availability Database Statistics data set.....	86
MS SQL Availability Database Summary data set.....	88
MS SQL Availability Group Listener Details data set.....	91
MS SQL Availability Groups Details data set.....	93
MS SQL Availability Groups Details In Cluster data set.....	96
MS SQL Availability Groups Summary data set.....	97
MS SQL Availability Replicas Details data set.....	100
MS SQL Availability Replicas Details In Cluster data set.....	102
MS SQL Availability Replicas Statistics data set.....	104
MS SQL Availability Replicas Status data set.....	106
MS SQL Availability Replicas Status Summary data set.....	108
MS SQL Batch Stats data set.....	111
MS SQL Blocked Queries Details data set.....	113
MS SQL Configuration data set.....	116
MS SQL Customized SQL Query Details data set.....	117
MS SQL Customized SQL Query Result data set.....	119
MS SQL Database Detail data set.....	124
MS SQL Database Mirroring data set.....	133
MS SQL Database Summary data set.....	138
MS SQL Device Detail data set.....	141
MS SQL Error Event Detail data set.....	144
MS SQL Expensive Query Plans data set.....	146
MS SQL Filegroup Detail data set.....	149
MS SQL FileTable Detail data set.....	150
MS SQL Individual Queries Details data set.....	152
MS SQL Integration Service Details data set.....	156

MS SQL Job Detail data set.....	157
MS SQL Job Summary data set.....	162
MS SQL Lock Conflict Detail data set.....	164
MS SQL Lock Detail data set.....	166
MS SQL Lock Resource Type Summary data set.....	168
MS SQL Lock Summary data set.....	170
MS SQL Log Shipping DB Details data set.....	173
MS SQL Log Shipping Errors data set.....	174
MS SQL Log Shipping Jobs Detail data set.....	176
MS SQL Memory Manager data set.....	177
MS SQL Problem Detail data set.....	180
MS SQL Problem Summary data set.....	182
MS SQL Process Detail data set.....	184
MS SQL Process Summary data set.....	189
MS SQL Remote Servers data set.....	193
MS SQL Resource Pool Stats data set.....	194
MS SQL Running Queries Details data set.....	197
MS SQL Server Detail data set.....	201
MS SQL Server Enterprise View data set.....	206
MS SQL Server Properties data set.....	211
MS SQL Server Summary data set.....	215
MS SQL Server Transactions Summary data set.....	223
MS SQL Service Broker Activation data set.....	226
MS SQL Service Broker Statistics data set.....	228
MS SQL Service Broker Transport data set.....	231
MS SQL Services Detail data set.....	233
MS SQL Statistics Detail data set.....	235
MS SQL Statistics Summary data set.....	237
MS SQL Table Detail data set.....	240
MS SQL Table Summary data set.....	243
MS SQL Text data set.....	246
MS SQL Workload Group Stats data set.....	247
Accessibility features.....	251
Notices.....	253
Trademarks.....	254
Terms and conditions for product documentation.....	255
IBM Online Privacy Statement.....	255
Index.....	257

Chapter 1. Monitoring Agent for Microsoft SQL Server

The Monitoring Agent for Microsoft SQL Server offers a central point of management for your Microsoft SQL Server environment or application.

The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Microsoft SQL Server agent you can easily collect and analyze Microsoft SQL Server specific information.

Installing and configuring the agent

Install the monitoring agent on the system where the application that you want to monitor is located.

For more information, see the agent installation and configuration topics in IBM Knowledge Center:

- [IBM Cloud Application Performance Management](#)
- [IBM Cloud Application Performance Management, Private](#)

For supported operating systems, see [System Requirements](#) in the APM Developer Center.

Chapter 2. Dashboard

Open the Application Performance Dashboard in the Cloud APM console to see a status summary of all your applications. As you drill down to dashboard pages for specific applications and their supporting elements, more details are available about the selected item. Use the Microsoft SQL Server agent dashboard pages to proactively monitor your Microsoft SQL Server deployment. Each page contains views with key performance indicators.

When an application that includes Microsoft SQL Server *managed resources* is selected, the navigator and the **Status Overview** tab show Microsoft SQL Server in the Components group:

- Click **Components** to see a single Microsoft SQL Server group widget that is displayed along with a group widget for every other data source type in the application.
- Click the Microsoft SQL Server subgroup to see a group widget for each managed resource in the application.
- Click inside a Microsoft SQL Server group widget or click a Microsoft SQL Server managed resource from the navigator **Instances** section to open a dashboard page with KPIs from the selected managed resource.

For more information about the KPIs, click  in the view or click  in the dashboard banner.

Default dashboard pages

Microsoft SQL Server

The Components page provides an overall status of all the SQL Server instances that are being monitored. Click anywhere in the group widget to get detailed information about the status and performance of SQL Server instances. If the Events tab shows a warning or critical status indicator, open the tab to see the open events and status for the application.

Alerts Detail

Use the Alerts Detail dashboard to get detailed information about the errors that are logged in the current SQL error log file.

Availability Groups - Detail

Use the Availability Groups - Detail dashboard to do the following tasks:

- View the current status of the availability databases and replicas by selecting the availability group.
- Determine if any resources of the availability databases are experiencing problems.
- View the details of the availability groups.

Database - Detail

Use the Database - Detail dashboard to know the following details of the selected database:

- Total usage of database space and log.
- Accessibility properties of the database.
- Historical data related to database transactions.
- Information about the database resources, such as the number of devices and tables.
- Information about the database recovery methods.

Databases - Detail

Use the Databases - Detail dashboard to know the following information about all the SQL databases:

- The current status and space usage of all the databases.
- Determine if any of the databases are experiencing problems.

Jobs - Detail

Use the Jobs - Detail dashboard to know details of the jobs, such as status of the jobs and the historical data related to the failure and success of jobs during a specific duration. Also, use this dashboard to determine if any of the jobs are experiencing problems.

Locks - Detail

Use the Locks - Detail dashboard to know the details of locks, such as status, resource type and conflict details. Also, use this dashboard to determine if any of the locks are experiencing problems. Click anywhere on the Locks group widget to drill down to additional details about locks.

Overview

The Monitoring Agent for Microsoft SQL Server dashboard shows the current status and availability of each monitored SQL application. Review the resource usage and investigate any warning or critical thresholds. If the Events tab shows a warning or critical status indicator, open the tab to see the open events and status for the application. Click anywhere on a group widget to drill down to detailed information about the application instance.

Processes - Detail

Use the Processes - Detail dashboard to view detailed information about all the SQL Server processes, such as the state of the processes and the number of blocked processes.

Server - Detail

Use the Server - Detail dashboard to know the SQL Server details, such as server statistics and services details. Also, determine if any of the SQL Server resources are experiencing problems. Click anywhere on the Availability group widget to drill down to detailed information about the availability groups, replicas, and databases.

Server Performance - Detail

Use the Server Performance - Detail dashboard to know the following details of the SQL servers:

- The CPU and memory usage.
- Historical data of the server performance parameters, such as memory usage and cache usage during a specific duration.
- The top 10 cached query plans according to the performance statistics of the SQL server.

Tables - Detail

Use the Tables - Detail dashboard to know the details of tables, such as properties of the fragmented tables and table optimization statistics. Also, use this dashboard to determine if any of the table resources are experiencing problems. Click the Database name in the Databases group widget to get detailed information about the tables in that database. The detailed information about the tables in a database appears in the Tables Status group widget.

Additional group widgets

These pop-up group widgets are displayed after you click a group widget for more details. Some group widgets have links to more granular information in a popup widget, described here.

Widgets for the Default dashboard pages

Microsoft SQL Server

The Components page provides an overall status of all the SQL Server instances that are being monitored. Click anywhere in the group widget to get detailed information about the status and performance of SQL Server instances. If the Events tab shows a warning or critical status indicator, open the tab to see the open events and status for the application.

The following widgets are available in this dashboard page:

MS SQL Server

The MS SQL Server group widget provides an HTML table that summarizes the overall health of the SQL server.

The following KPIs for the widget are derived from the attributes as described:

- Cache hit ratio (%): The current ratio of data cache hits to total requests. Use this attribute to check the effectiveness of the data cache. For optimal performance, the returned value must be

approximately 90% or greater. The value format is an integer. This data is derived from the Cache Hit Ratio attribute in the MS SQL Server Summary data set.

- **Collector status:** Indicates the status of the data collector on a remote node. The data collector is the part of the product that collects information about the SQL server. This data is derived from the Collection Status attribute in the MS SQL Server Summary data set.
- **Databases:** The cumulative status of all the databases based on the total data and log space usage, status of databases, and age of the oldest open transaction that is running on the database.
- **Jobs:** The status of processed jobs in the current sampling interval.
- **Processes blocked (%):** The percentage of processes that are being blocked. The value format is a percentage with two decimal places allowed; for example, 4.00. This value is based on information from the sysprocesses table. The value includes all processes currently in a waiting state. This data is derived from the Percent Processes Blocked attribute in the MS SQL Process Summary data set.
- **Server errors:** The number of error messages that occurred during the current interval. The value format is an integer; for example, 1. This data is derived from the Total Errors Current Interval attribute in the MS SQL Problem Summary data set.
- **Server status:** Indicates the status of the SQL server. The server status is displayed as Unknown when the collector process of the SQL Server agent is not running. The server status is displayed as Inactive when the collector process of the SQL Server agent is running, but the SQL Server is not responding to the request. This data is derived from the Server Status attribute in the MS SQL Server Summary data set.

Alerts Detail

Use the Alerts Detail dashboard to get detailed information about the errors that are logged in the current SQL error log file.

The following widgets are available in this dashboard page:

Alert Details

The Alert Details group widget provides a grid widget that gives the details of the alerts.

The following KPIs for the widget are derived from the attributes as described:

- **Error ID:** The ID of the error message. The value format is an integer; for example, 2520. The unnumbered errors with text explanations are written to the SQL server message logs. This data is derived from the Error ID attribute in the MS SQL Problem Detail data set.
- **Error SPID:** The ID of the session when an event occurs. It is prefixed with spid. For example, spid34s, spid57. When the ID of the session is unavailable, the source of error is showed as it is. For example, logon. This attribute is not applicable for eventing thresholds. This data is derived from the Error SPID attribute in the MS SQL Problem Detail data set.
- **Message Age (minutes):** The number of minutes that have elapsed since the error occurred. The value format is an integer; for example, 2. Monitor this value to track current messages. This data is derived from the Message Age attribute in the MS SQL Problem Detail data set.
- **Message Text:** The message text. This attribute is not available for use in eventing thresholds. This data is derived from the Message Text U attribute in the MS SQL Problem Detail data set.
- **Message Timestamp:** The timestamp that indicates the date and time the error occurred. The format is MM/DD/YY HH:mm:ss, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. This data is derived from the Message Timestamp attribute in the MS SQL Problem Detail data set.
- **Severity Level:** Indicates the severity level of the error. This data is derived from the Severity Level attribute in the MS SQL Problem Detail data set.

- SQL State Code: The SQL state value for the error message. The value format is an integer; for example, 37. This data is derived from the SQL State Code attribute in the MS SQL Problem Detail data set.

Alert Properties

The Alert Properties group widget provides an HTML table that gives details of the alert properties.

The following KPIs for the widget are derived from the attributes as described:

- Age of last error (minutes): The number of minutes that have elapsed since the last error message occurred. The value format is an integer; for example, 2. This data is derived from the Age of Last Error attribute in the MS SQL Problem Summary data set.
- Error messages in current interval: The number of error messages that occurred during the current interval. The value format is an integer; for example, 1. This data is derived from the Total Errors Current Interval attribute in the MS SQL Problem Summary data set.
- Error messages since startup: The number of error messages that have occurred since the SQL server was started. The value format is an integer; for example, 3. This data is derived from the Total Error Messages attribute in the MS SQL Problem Summary data set.
- Highest severity timestamp: The timestamp that indicates the date and time the error message with the highest severity level occurred. The format is MM/DD/YY HH:mm:SS, where MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. This data is derived from the Maximum Sev Timestamp attribute in the MS SQL Problem Summary data set.
- Low severity errors: The number of error messages with a severity level of less than 17 that have occurred since the SQL server was started. The value format is an integer; for example, 3. This data is derived from the Total Errors Other attribute in the MS SQL Problem Summary data set.

Errors In Current Interval (history)

The Errors In Current Interval (history) group widget provides a line graph that gives the details of the alert errors during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Problem Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Errors: The number of error messages that occurred during the current interval. The value format is an integer; for example, 1. This data is derived from the Total Errors Current Interval attribute in the MS SQL Problem Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Problem Summary data set.

Availability Groups - Detail

Use the Availability Groups - Detail dashboard to do the following tasks:

- View the current status of the availability databases and replicas by selecting the availability group.
- Determine if any resources of the availability databases are experiencing problems.
- View the details of the availability groups.

The following widgets are available in this dashboard page:

Availability Databases Status

The Availability Databases Status group widget provides a grid widget that gives the status of the availability databases.

The following KPIs for the widget are derived from the attributes as described:

- Group Name: The name of the availability group to which the availability database belongs. This data is derived from the Group Name U attribute in the MS SQL Availability Database Details data set.

- Local: Indicates whether the availability replica is hosted by the local SQL Server instance. This data is derived from the Local attribute in the MS SQL Availability Database Details data set.
- Name: The name of the availability database that is hosted by the SQL Server instance. This data is derived from the Database Name U attribute in the MS SQL Availability Database Details data set.
- Role: Indicates whether the availability database is a primary or a secondary database. This data is derived from the Role attribute in the MS SQL Availability Database Details data set.
- State: Indicates the current state of the availability database. This data is derived from the Database State attribute in the MS SQL Availability Database Details data set.
- Suspend Reason: The reason for a database to enter the suspended state. This data is derived from the Suspend Reason attribute in the MS SQL Availability Database Details data set.
- Synchronization Health: Indicates the health of the availability database replica. The replica is healthy if the synchronized and the asynchronized replicas are in the Synchronized state. This data is derived from the Synchronization Health attribute in the MS SQL Availability Database Details data set.
- Synchronization State: The current synchronization state of the availability database replica. This data is derived from the Synchronization State attribute in the MS SQL Availability Database Details data set.
- Synchronized Commit: Indicates whether the transaction commit is synchronized with the database replica. This data is derived from the Synchronized Commit attribute in the MS SQL Availability Database Details data set.

Availability Groups

The Availability Groups group widget provides a grid widget that gives details of the availability groups.

In the Availability Groups group widget, when you click a particular availability group the status of availability replicas and availability databases is displayed in the Availability Replicas Status and Availability Databases Status group widgets.

The following KPIs for the widget are derived from the attributes as described:

- Name: The name of the availability group to which the replica belongs. The value format is an alphanumeric string with a maximum of 384 characters. This data is derived from the Group Name U attribute in the MS SQL Availability Groups Details data set.
- Status: A summary of the synchronization health of all the replicas in the availability group. This data is derived from the Synchronization Health attribute in the MS SQL Availability Groups Details data set.

Availability Replicas Status

The Availability Replicas Status group widget provides a grid widget that gives the status of the availability replicas.

The following KPIs for the widget are derived from the attributes as described:

- Connection State: The current connection state of the availability replica. This data is derived from the Connection State attribute in the MS SQL Availability Replicas Status data set.
- Databases Count: The total number of databases that are hosted by the availability replica. This data is derived from the Total Databases attribute in the MS SQL Availability Replicas Status data set.
- ID: The ID of the replica. The value format is an alphanumeric string with a maximum of 64 characters. This data is derived from the Replica ID attribute in the MS SQL Availability Replicas Status data set.
- Local: Indicates whether the replica is hosted by the local instance. This data is derived from the Local attribute in the MS SQL Availability Replicas Status data set.

- Operational State: The current operational state of the availability replica. This attribute shows the value No Data on the replica that is not local. This data is derived from the Operational State attribute in the MS SQL Availability Replicas Status data set.
- Recovery Health: Indicates whether the databases that are connected to the availability group are online or are being recovered after a failover. This attribute shows the value No Data on the replica that is not local. This data is derived from the Recovery Health attribute in the MS SQL Availability Replicas Status data set.
- Role: Indicates whether the replica is a primary or a secondary replica. This data is derived from the Role attribute in the MS SQL Availability Replicas Status data set.
- Synchronization Health: The synchronization state of all the databases that are connected to the availability group on the availability replica. This data is derived from the Synchronization Health attribute in the MS SQL Availability Replicas Status data set.

Database - Detail

Use the Database - Detail dashboard to know the following details of the selected database:

- Total usage of database space and log.
- Accessibility properties of the database.
- Historical data related to database transactions.
- Information about the database resources, such as the number of devices and tables.
- Information about the database recovery methods.

The following widgets are available in this dashboard page:

Database Properties

The Database Properties group widget provides an HTML table that gives information about the database accessibility.

The following KPIs for the widget are derived from the attributes as described:

- Access validation: Access validation of the database. When the database is in offline, restoring, or transition state, the value of this attribute is displayed as No. This data is derived from the Access Validation attribute in the MS SQL Database Detail data set.
- Database owner: The SQL server-assigned user ID for the owner of the database. Use the create database command to establish this identifier. This data is derived from the DB Owner U attribute in the MS SQL Database Detail data set.
- Database owner only access: Indicates whether the database has a status of DBO only. This data is derived from the DBO Only Access attribute in the MS SQL Database Detail data set.
- Free space accounting suppressed: Indicates whether the free space accounting option is disabled for the database. The no-free-space-actg option turns off free space accounting on non-log segments only. Information about free space is inaccurate when free space accounting is turned off. Use the no-free-space-actg option and the checkpoint command to speed recovery. No time is needed to count free space for non-log segments. This data is derived from the Free Space Accounting Supresed attribute in the MS SQL Database Detail data set.
- Read only access: Indicates whether the database has a status of read only. This data is derived from the Read Only Access attribute in the MS SQL Database Detail data set.
- Single user access: Indicates whether the database has a status of single user. A database with a status of single user can only be accessed by one user at a time. This data is derived from the Single User Access attribute in the MS SQL Database Detail data set.

Database Recovery

The Database Recovery group widget provides an HTML table that gives information about the recovery methods.

The following KPIs for the widget are derived from the attributes as described:

- Page verify: The option that identifies and notifies incomplete I/O transactions that have occurred because of disk I/O errors. This data is derived from the Page Verify attribute in the MS SQL Database Detail data set.
- Recovery model: The method used to maintain the transaction log. The three types of recovery models include Simple, Full, and Bulk-logged. This data is derived from the Recovery Model attribute in the MS SQL Database Detail data set.

Database Resources

The Database Resources group widget provides an HTML table that gives information about the database resources.

The following KPIs for the widget are derived from the attributes as described:

- Device count: The number of devices allocated for the database. The value format is an integer in the range 1-128; for example, 4. This data is derived from the Total Devices attribute in the MS SQL Database Detail data set.
- Table count: The number of tables that exist in the database. The value format is an integer. This data is derived from the Table Count attribute in the MS SQL Database Detail data set.

Database Transaction

The Database Transaction group widget provides an HTML table that gives information about the database transactions.

The following KPIs for the widget are derived from the attributes as described:

- Long running transaction name: The long running transaction name. This data is derived from the Long Running Transaction Name attribute in the MS SQL Database Detail data set.
- Long running transaction process ID: The longest running transaction. The Long running transactions are transactions that are open for longer than the LongRunningTransColl standard collector parameter. The format is an integer. This data is derived from the Long Running Transaction Process ID attribute in the MS SQL Database Detail data set.
- Long running transaction time (seconds): The time for which the longest running transaction is active (in seconds). The format is an integer. This data is derived from the Long Running Transaction Time Per Sec attribute in the MS SQL Database Detail data set.

Database Used

The Database Used group widget provides an HTML table that summarizes the status of the database space utilization.

The following KPIs for the widget are derived from the attributes as described:

- Aggregate data free space (MB): The total amount of freespace (in MB) that is available in all the data files and on the storage device. The value format is a decimal number with two decimal places allowed. This data is derived from the Aggregate Database Freespace attribute in the MS SQL Database Detail data set.
- Data maximum growth (MB): The maximum size to which the database can grow in MB. This data is derived from the Database Maximum Growth Size attribute in the MS SQL Database Detail data set.
- Data size (MB): The number of megabytes (MB) allocated for the data only segments of the database. The value format is a decimal number with two decimal places allowed; for example, 50.00. This data is derived from the Data Size attribute in the MS SQL Database Detail data set.
- Database growth (%): The percentage of growth for the database from the last sample to the current sample. This data is derived from the Database Growth Percent attribute in the MS SQL Database Detail data set.

Database Write Transactions(history)

The Database Write Transactions(history) group widget provides a Sparkline that gives information about the number of transactions written to the database and committed per second.

The following KPIs for the widget are derived from the attributes as described:

- Write transactions per second: The number of transactions written to the database and committed in the last second. The format is an integer. When the number of written transactions collected by the agent in the last interval is greater than or equal to the number of written transactions collected in the current interval, then 0 (zero) is displayed as the attribute value on the console. When the number of written transactions collected in the last interval is less than the number of written transactions collected in the current interval, then the value is calculated as follows: The number of written transactions collected in the last interval is subtracted from the number of written transactions collected in the current interval, and the value after subtraction is divided by the time interval between these two collections. When the value of written transactions, which are collected in the last or the current interval, is negative, then Not_Collected is displayed on the console. This data is derived from the Write Transactions per sec attribute in the MS SQL Database Detail data set.

Device Status

The Device Status group widget provides a grid widget that summarizes the status of devices.

The following KPIs for the widget are derived from the attributes as described:

- Auto Growth: Indicates whether the autogrowth feature is enabled for a device. If the autogrowth feature is enabled, the maximum file size is displayed as restricted or unrestricted. This data is derived from the Autogrowth attribute in the MS SQL Device Detail data set.
- Auto Growth Unit: Indicates the unit of file growth. If the autogrowth feature is enabled for a device, the file growth value is displayed either in a percentage or MB. This data is derived from the Autogrowth Unit attribute in the MS SQL Device Detail data set.
- Auto Growth Value: Indicates the growth value of a file. This data is derived from the Autogrowth Value attribute in the MS SQL Device Detail data set.
- Capacity (MB): The number of megabytes (MB) on the device. The value format is an integer; for example, 20. Use this value when analyzing the amount of free space. This data is derived from the Device Size attribute in the MS SQL Device Detail data set.
- Capacity Used (%): The percentage of space that is used on the device. This data is derived from the Device Free Space Percent attribute in the MS SQL Device Detail data set.
- File Name: The name of the physical device allocated for the database. The value format is an alphanumeric string with a maximum of 30 characters; for example, DATA_1. This data is derived from the Physical Device Name attribute in the MS SQL Device Detail data set.
- Freespace (MB): The number of megabytes (MB) of free space on a device. The value format is a decimal with two decimal places allowed; for example, 10.00. Adequate space is required to support database and system administration activities. Set alerts for abnormal conditions. Refer to this value when estimating space needed for tables, indexes, logs, and system administration. This data is derived from the Device Free Space attribute in the MS SQL Device Detail data set.
- Name: The Device name. This data is derived from the Device Name U attribute in the MS SQL Device Detail data set.
- Status: Indicates the status of the device.
- Type: Indicates the type of device allocated for the database. The value indicates the type of data that is stored on the device. Databases are frequently spread across several devices due to size, performance, and recoverability issues. This data is derived from the Device Type attribute in the MS SQL Device Detail data set.

Log Used

The Log Used group widget provides a grid widget that summarizes the status of the log space utilization.

The following KPIs for the widget are derived from the attributes as described:

- Log free space (MB): The number of megabytes (MB) of free space in the transaction log for the database. The value includes the number of megabytes of free space on the log only partitions. The value format is a decimal number with two decimal places, for example, 8. 00. Various types of transactions, such as mass updates and bulk copying can involve extensive logging. This data is derived from the Log Freespace attribute in the MS SQL Database Detail data set.

- Log growths: The total number of times the database transaction log is expanded since the SQL Server was started. The format is an integer. This data is derived from the Log Growths attribute in the MS SQL Database Detail data set.
- Log maximum growth (MB): The maximum size to which the log can grow in MB. This data is derived from the Log Maximum Growth Size attribute in the MS SQL Database Detail data set.
- Log size (MB): The number of megabytes (MB) allocated for the transaction log for the database. The value includes the number of megabytes allocated for the transaction log on the log only partitions. The value format is a decimal number with two decimal places allowed; for example, 500. The appropriate size for a transaction log depends on how the database is used. Several factors, such as the number of transactions, type of transactions, and number of users, effect sizing. Evaluate the need to truncate the transaction log to prevent it from filling up. This data is derived from the Log Size attribute in the MS SQL Database Detail data set.

Databases - Detail

Use the Databases - Detail dashboard to know the following information about all the SQL databases:

- The current status and space usage of all the databases.
- Determine if any of the databases are experiencing problems.

The following widgets are available in this dashboard page:

Database Status

The Database Status group widget provides a grid widget that summarizes the status of the database.

To open the Database Used (history) group widget, click the Database Used column of the Database Status group widget.

To see a line graph in the Log Used (history) group widget, click the Log Utilization column of the Database Status group widget.

The following KPIs for the widget are derived from the attributes as described:

- Active Transactions: The number of active transactions for the database. This data is derived from the Active Transactions attribute in the MS SQL Database Detail data set.
- Database Used (%): The amount of space (in megabytes) used in the database as a percentage of total space allowed. This data is derived from the Database Space Percent Used attribute in the MS SQL Database Detail data set.
- Error Status: Indicates whether the database has an error status. A database with an error status has a status of suspect, crashed, or recovery. Check the status bits in the sysdatabases table to determine the cause of the error. Use the database consistency checker (dbcc) to verify the database integrity. This data is derived from the Error Status attribute in the MS SQL Database Detail data set.
- Log Used (%): The percentage of the transaction log that is full. This data is derived from the Log Space Percent Used attribute in the MS SQL Database Detail data set.
- Name: The name of the database. Each database name is unique. The SQL server also assigns each database its own identification number. This data is derived from the Database Name U attribute in the MS SQL Database Detail data set.
- Oldest Open Transactions (minutes): The age (in minutes) of the oldest open transaction in the database transaction log. Use this attribute to determine up to which point in time the transaction log can be truncated and backed up. This data is derived from the Oldest Open Transaction attribute in the MS SQL Database Detail data set.
- Replicated Transactions: The number of transactions in the publication database transaction log that are marked for replication but have not yet been delivered to the distribution database. Use this attribute to discover if there are any bottlenecks in the replication process. Bottlenecks cause delays in the published transactions reaching the subscriber database. This data is derived from the Replicated Transactions attribute in the MS SQL Database Detail data set.

- Replication Latency (ms): The number of milliseconds between the time a transaction marked for replication is entered into the publication database transaction log and the time it is read out of the log and delivered to the distribution database. Use this attribute to discover if there are any bottlenecks in the replication process. Bottlenecks cause delays in the published transactions reaching the subscriber database. This data is derived from the Replication Latency attribute in the MS SQL Database Detail data set.
- State: Reports the database state. This data is derived from the Database State attribute in the MS SQL Database Detail data set.
- Status: Reports when the database is unavailable. If a database is offline, you cannot access it. Use this attribute to warn that the database has become unavailable. When the database is in transition state, the database status is displayed as Available. This data is derived from the Database Status attribute in the MS SQL Database Detail data set.
- Transactions Per Second: The number of transactions started for the database per second for the current interval. When the number of transactions collected by the agent in the last interval is greater than or equal to the number of transactions collected in the current interval, then 0 (zero) is displayed as the attribute value on the console. When the number of transactions collected in the last interval is less than the number of transactions collected in the current interval, then the value is calculated as follows: The number of transactions collected in the last interval is subtracted from the number of transactions collected in the current interval, and the value after subtraction is divided by the time interval between these two collections. When the value of transactions collected in the last or the current interval, is negative, then Not_Collected is displayed on the console. This data is derived from the Transactions per Second attribute in the MS SQL Database Detail data set.

Database Used - Top 5

The Database Used - Top 5 group widget provides a grid widget that contains the top five database used in percentage.

The following KPIs for the widget are derived from the attributes as described:

- Database Used (%): The amount of space used in the database as a percentage of total space allowed. This data is derived from the Database Space Percent Used attribute in the MS SQL Database Detail data set.
- Name: The name of the database. This data is derived from the Database Name U attribute in the MS SQL Database Detail data set.

Databases Configuration Summary

The Databases Configuration Summary group widget provides an HTML table that summarizes the configuration details.

The following KPIs for the widget are derived from the attributes as described:

- Database owner only: The number of databases with a status of DBO only. A database with a status of database owner can be accessed only by users with DBO authority. The value format is an integer; for example, 5. This data is derived from the Total DBs DBO Only attribute in the MS SQL Database Summary data set.
- No freespace accounting: The number of databases that have the free space accounting option disabled. The value format is an integer; for example, 5. This data is derived from the Total DBs No Free Space Accounting attribute in the MS SQL Database Summary data set.
- Read only: The number of databases with a status of read only. The value format is an integer; for example, 5. This data is derived from the Total DBs Read Only attribute in the MS SQL Database Summary data set.
- Single user: The number of databases with a status of single user. A database with a status of single user can be accessed by only one user at a time. The value format is an integer; for example, 2. This data is derived from the Total DBs Single User attribute in the MS SQL Database Summary data set.

Databases in Errors (history)

The Database in Errors (history) group widget provides a line graph that summarizes the total number of databases and the number of databases in error during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Database Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- **Databases:** The number of databases for the SQL server. The value format is an integer; for example, 10. This data is derived from the Total Databases attribute in the MS SQL Database Summary data set.
- **Databases in error:** The number of databases with an error status. A database with an error status is a database with a status of suspect, crashed, or recovery. The value format is an integer; for example, 2. This data is derived from the Total DBs in Error attribute in the MS SQL Database Summary data set.
- **Time:** The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Database Summary data set.

Jobs - Detail

Use the Jobs - Detail dashboard to know details of the jobs, such as status of the jobs and the historical data related to the failure and success of jobs during a specific duration. Also, use this dashboard to determine if any of the jobs are experiencing problems.

The following widgets are available in this dashboard page:

Job Count

The Job Count group widget provides a grid widget that lists the job owners and the number of jobs that are started by the respective owners.

The following KPIs for the widget are derived from the attributes as described:

- **Count:** The total number of jobs that are started by the owner. This data is derived from the Job Id attribute in the MS SQL Job Detail data set.
- **Owner:** The name of the owner of the job. The value format is an alphanumeric string with a maximum of 128 characters. This data is derived from the Job Owner attribute in the MS SQL Job Detail data set.

Job Status

The Jobs Status group widget provides a grid widget that gives information about the status of the jobs.

The following KPIs for the widget are derived from the attributes as described:

- **Completion Duration:** The amount of time it took for the job to complete (in seconds). This data is derived from the Job Duration attribute in the MS SQL Job Detail data set.
- **Enable:** Whether or not the job is enabled to run. This data is derived from the Enabled attribute in the MS SQL Job Detail data set.
- **Error Code:** Error code for the last completion of the job. This data is derived from the Job Error Code attribute in the MS SQL Job Detail data set.
- **Last Run Outcome:** The last job execution status. This data is derived from the Last Run Outcome attribute in the MS SQL Job Detail data set.
- **Last Run Timestamp:** The timestamp of last job execution. If the job has not been run at all, the Last Run Timestamp value is 'N/P' (Not Present). If there is an error retrieving the job information, the Last Run Timestamp value is 'N/P'. This data is derived from the Last Run Timestamp attribute in the MS SQL Job Detail data set.
- **Name:** The SQL Server job name. This data is derived from the Job Name attribute in the MS SQL Job Detail data set.
- **Retry Attempts:** The current number of retry attempts that are done on a step of a running job. This data is derived from the Current Retry Attempt attribute in the MS SQL Job Detail data set.

- Status: Current status of the job. This data is derived from the Job Status attribute in the MS SQL Job Detail data set.

Jobs Failed (history)

The Jobs Failed (history) group widget provides a line graph that contains the number of jobs that failed during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Job Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Failed jobs: The number of jobs that failed in the current sample interval. This data is derived from the Jobs Failed Current Interval attribute in the MS SQL Job Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Job Summary data set.

Jobs Since Server Startup

The Jobs Since Server Startup group widget provides an HTML table that summarizes the status of jobs.

The following KPIs for the widget are derived from the attributes as described:

- Active: The number of active jobs since the agent was started. The format is an integer. This data is derived from the Active Jobs attribute in the MS SQL Job Summary data set.
- Failed: The number of failed jobs since the SQL Server agent was started. The format is an integer. This data is derived from the Failed Jobs attribute in the MS SQL Job Summary data set.
- Queued: The number of jobs in a queue since the agent was started. The format is an integer. This attribute is not supported in MSSQL Linux agent. This data is derived from the Queued Jobs attribute in the MS SQL Job Summary data set.
- Successful: The number of successful jobs since the agent was started. The format is an integer. This data is derived from the Successful Jobs attribute in the MS SQL Job Summary data set.

Jobs Success Rate (history)

The Jobs Success Rate (history) group widget provides a line graph that summarizes the success rate of the jobs over during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Job Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Success Rate: The rate at which jobs are successful since the agent was started. The format is an integer. This data is derived from the Job Success Rate attribute in the MS SQL Job Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Job Summary data set.

Locks - Detail

Use the Locks - Detail dashboard to know the details of locks, such as status, resource type and conflict details. Also, use this dashboard to determine if any of the locks are experiencing problems. Click anywhere on the Locks group widget to drill down to additional details about locks.

The following widgets are available in this dashboard page:

Granted Lock Summary

The Granted Lock Summary group widget provides an HTML table that contains a list of types of locks.

The following KPIs for the widget are derived from the attributes as described:

- Exclusive locks: The number of granted Exclusive locks. This data is derived from the Exclusive Locks attribute in the MS SQL Lock Summary data set.

- Intent locks: The number of granted Intent locks. This data is derived from the Intent Locks attribute in the MS SQL Lock Summary data set.
- Page locks: The number of granted Page locks. This data is derived from the Page Locks attribute in the MS SQL Lock Summary data set.
- Shared locks: The number of granted Shared locks. This data is derived from the Shared Locks attribute in the MS SQL Lock Summary data set.
- Table locks: The number of granted Table locks. This data is derived from the Table Locks attribute in the MS SQL Lock Summary data set.

Lock Conflict Details

The Lock Conflict Details group widget provides a grid widget that contains detailed information about a selected lock conflict.

The following KPIs for the widget are derived from the attributes as described:

- Blocking Process ID: The identifier for the process that is blocking a request for a lock. The value format is an alphanumeric string with a maximum of 8 characters; for example, 12. Some blocking processes can become stranded. Investigate eventing thresholds in which a process is being blocked for an extended period of time. For more information on a blocking process, query the sysprocesses and use the sp-lock procedure. This data is derived from the Blocking Process ID attribute in the MS SQL Lock Conflict Detail data set.
- Client Group ID: The group ID of the user executing the process. This value does not apply to Microsoft SQL Server version 7. The value format is an alphanumeric string with a maximum of 8 characters; for example, ACCT_1. This data is derived from the Client Group ID attribute in the MS SQL Lock Conflict Detail data set.
- Client User ID: ID of the user who executed the command. The value format is an alphanumeric string with a maximum of 8 characters. This data is derived from the Client User ID attribute in the MS SQL Lock Conflict Detail data set.
- Database Name: The name of the database. This data is derived from the Database Name U attribute in the MS SQL Lock Conflict Detail data set.
- Requester process ID: The ID of the blocked process. The value format is an alphanumeric string with a maximum of 16 characters; for example, 21. This data is derived from the Requestor Process ID attribute in the MS SQL Lock Conflict Detail data set.

Lock Details

The Locks Details group widget provides a grid widget that summarizes the details of the locks.

The following KPIs for the widget are derived from the attributes as described:

- Database ID: The ID of the database that is locked. The value format is an integer. This data is derived from the Database Id attribute in the MS SQL Lock Detail data set.
- Database Name: The name of the locked database. This data is derived from the Database Name U attribute in the MS SQL Lock Detail data set.
- Lock Request Status: Enumeration that identifies the Lock Request Status. The format is an integer. This data is derived from the Lock Request Status E attribute in the MS SQL Lock Detail data set.
- Lock Resource Type: Enumeration that identifies the Lock Resource Type. The format is an integer. This data is derived from the Lock Resource Type E attribute in the MS SQL Lock Detail data set.
- Lock Type: Indicates the type of lock on the resource that is being requested. This data is derived from the Type attribute in the MS SQL Lock Detail data set.
- Page Number: The page number of the table being locked. The value format is an integer. This data is derived from the Page Number attribute in the MS SQL Lock Detail data set.
- Process Holding Lock: The ID of the process holding the lock. The value format is an integer. This data is derived from the Process Holding Lock attribute in the MS SQL Lock Detail data set.

- Table Name: The name of the table being locked. The value format is an alphanumeric string with a maximum of 16 characters. This data is derived from the Table Name attribute in the MS SQL Lock Detail data set.

Lock Resource Type Summary

The Lock Resource Type Summary group widget provides a grid widget that contains the details of the Lock Resource Type.

The following KPIs for the widget are derived from the attributes as described:

- Lock Requests Per Second: The number of new locks and lock conversions per second requested from the lock manager. This value is calculated on an interval, and does not reflect a total value. This data is derived from the Lock Requests per Second attribute in the MS SQL Lock Resource Type Summary data set.
- Lock Resource Type: An enumeration of the resources that the SQL Server can lock. This data is derived from the Lock Resource Type attribute in the MS SQL Lock Resource Type Summary data set.
- Lock Timeouts Per Second: The number of lock requests per second that timed out, including requests for NOWAIT locks. This data is derived from the Lock Timeouts per Second attribute in the MS SQL Lock Resource Type Summary data set.
- Lock Wait Time (ms): The total wait time (in milliseconds) for locks in the last second. This data is derived from the Lock Wait Time attribute in the MS SQL Lock Resource Type Summary data set.
- Number of Deadlocks Per Second: The number of Deadlocks per second for the current sample interval. This data is derived from the Number Deadlocks per Second attribute in the MS SQL Lock Resource Type Summary data set.

Locks

The Locks group widget provides an HTML table that contains the status of the locks. For more details about the locks, click in the group widget area to open the Additional Locks Details group widget.

The following KPIs for the widget are derived from the attributes as described:

- Blocking lock requests: The total number of current locks blocking other processes. A blocking lock must be released before other processes requesting competing locks can progress. (For more information on locks, see the Microsoft SQL Server Books online.) Use this attribute to show the number of blocking locks active during server activity. This attribute can indicate that processes are being held up through lock contention rather than hardware performance issues. The value format is an integer. This data is derived from the Total Blocking Lock Requests attribute in the MS SQL Server Summary data set.
- Current locks: The number of current locks for the SQL server. The value format is an integer; for example, 73. This data is derived from the Current Locks attribute in the MS SQL Server Detail data set.
- Locks remaining: The total number of locks that can still be taken out. The maximum number of locks is configurable. In Microsoft SQL Server Version 8.0, you can have the number of locks configured automatically. Use this attribute to show the number of locks active during server activity. The value format is an integer. This data is derived from the Total Locks Remaining attribute in the MS SQL Server Summary data set.
- Maximum locks allowed: The maximum number of allowable locks. The value format is an integer; for example, 2000. This value is a configuration parameter. This data is derived from the Max Locks Allowed attribute in the MS SQL Server Detail data set.

Overview

The Monitoring Agent for Microsoft SQL Server dashboard shows the current status and availability of each monitored SQL application. Review the resource usage and investigate any warning or critical thresholds. If the Events tab shows a warning or critical status indicator, open the tab to see the open events and status for the application. Click anywhere on a group widget to drill down to detailed information about the application instance.

The following widgets are available in this dashboard page:

Databases

The Databases group widget provides a Status Summary widget that summarizes the percentage of the space that is used by the databases.

The following KPIs for the widget are derived from the attributes as described:

- Databases (Critical): Indicates that the databases is in the critical state. This data is derived from the Total Databases Critical attribute in the MS SQL Database Summary data set.
- Databases (Normal): The number of databases with a normal status. A database is in the normal state when the database state is online, and when the database status is available, and when the age of the oldest open transaction on the database is less than 6 minutes, and when the percentage space usage of database log files is less than 70, and when the percentage space usage of database data files is less than 70.
- Databases (Warning): Indicates that the databases is in the warning state. This data is derived from the Total Databases Warning attribute in the MS SQL Database Summary data set.

Errorlog Alert

The Errorlog Alert group widget provides an HTML table that summarizes the details of the error logs.

The following KPIs for the widget are derived from the attributes as described:

- Error log size (bytes): The number of bytes in the error log file. The value format is an integer; for example, 50000. This data is derived from the Error Log Size attribute in the MS SQL Problem Summary data set.
- Highest severity level: The error message of the highest severity level encountered during the current interval. The value format is an integer in the range 10 through 24; for example, 19. This data is derived from the Maximum Sev Current Interval attribute in the MS SQL Problem Summary data set.

Jobs In Current Interval

The Jobs In Current Interval group widget provides a status summary widget that summarizes the status of the jobs that are processed in the current sampling interval.

The following KPIs for the widget are derived from the attributes as described:

- Jobs (Critical): The number of jobs that failed in the current sample interval. This data is derived from the Jobs Failed Current Interval attribute in the MS SQL Job Summary data set.
- Jobs (Succeeded): The number of jobs that are processed successfully in the current sample interval.
- Jobs (Warning): The number of processed jobs in the current sampling interval with the status as unknown or canceled. This data is derived from the Warning Jobs Current Interval attribute in the MS SQL Job Summary data set.

Locks

The Locks group widget provides a volume gauge that summarizes the details of the Locks.

The following KPIs for the widget are derived from the attributes as described:

- Locks used (%): The percentage of locks on resources of the maximum number of locks allowed by the SQL server. The value format is a percentage with two decimal places allowed; for example, 10.00. Setting a higher limit for the maximum number of locks does not impair performance. If your operations exceed the number of available locks, you can increase this limit. This data is derived from the Percent Max Locks attribute in the MS SQL Server Summary data set.

Processes

The Processes group widget provides an HTML table that summarizes the details of the server processes.

The following KPIs for the widget are derived from the attributes as described:

- Log suspend: The number of processes with a status of log suspend. A process with a status of log suspend is a process that cannot complete until there is free space in the transaction log. The value format is an integer; for example, 1. This data is derived from the Total Log Suspend attribute in the MS SQL Process Summary data set.
- Processes: The number of processes. The value includes background processes, processes for applications, and user processes. The value format is an integer; for example, 50. This value includes background processes, processes for applications, and user processes. This data is derived from the Total Processes attribute in the MS SQL Process Summary data set.

Server Details

The Server Details group widget provides a group of HTML table and volume gauge that summarizes the details of the server.

The following KPIs for the widget are derived from the attributes as described:

- Cache average free scan: The average number of buffers scanned by the LazyWriter when the LazyWriter searches the data cache for an unused buffer to replenish the free buffer pool. If Microsoft SQL Server must read a large number of buffers to find a free one, server performance might degrade. A low value indicates optimal performance. The value format is an integer. This data is derived from the Cache Average Free Scan attribute in the MS SQL Server Summary data set.
- Cache maximum free page scan: The maximum value for the number of buffers scanned by the LazyWriter when the LazyWriter searches the data cache for an unused buffer to replenish the free buffer pool. Use this attribute to check the effectiveness of the data cache. The value format is an integer. This data is derived from the Cache Maximum Free Page Scan attribute in the MS SQL Server Summary data set.
- Client count used (%): The number of client workstations currently connected to the database server and returns the ratio percentage of the number of possible connections. For example, if you have 10 client connections and eight are currently in use, this attribute returns a value of 80%. Running out of client connections can result in server access problems. The value format is an integer. This data is derived from the Client Count Percent Used attribute in the MS SQL Statistics Summary data set.
- Procedure buffer and cache: The procedure and buffer cache.

Server Performance

The Server Performance group widget provides a group of volume gauges and an HTML table that summarizes the details of the server performance.

The following KPIs for the widget are derived from the attributes as described:

- Active connections: The number of active connections (logons). The value format is an integer; for example, 50. This data is derived from the Current Logons attribute in the MS SQL Statistics Summary data set.
- Active transactions: The number of transaction enlistments (local, dtc, and bound) that are currently active. The value format is an integer. This data is derived from the Total Active Transactions attribute in the MS SQL Server Transactions Summary data set.
- CPU used (%): The percentage of CPU time the SQL server process is using on the host. The value format is a percentage with two decimal places allowed; for example, 20.00. This data is derived from the Server CPU Percent attribute in the MS SQL Server Summary data set.
- Logins per second: The number of login operations that started per second in the current interval. The value format is a decimal with two decimal places allowed; for example, 15.00. This data is derived from the Total Logins Per Sec attribute in the MS SQL Statistics Summary data set.
- Logouts per second: Total number of logout operations started per second. This data is derived from the Total Logouts per Second attribute in the MS SQL Statistics Summary data set.
- Maximum logons active (%): The percentage of active connections (logons) of the maximum number of active connections allowed for the SQL server. The value format is a percentage with

two decimal places allowed; for example, 50.00. This data is derived from the Percent Max Logons Active attribute in the MS SQL Statistics Summary data set.

- Page life expectancy: The duration (in seconds) for which an SQL Server block or page is stored in the memory. This data is derived from the Page Life Expectancy attribute in the MS SQL Server Summary data set.
- Queries (Critical): The state of the query. The critical status is displayed in red color. This data is derived from the Query State attribute in the MS SQL Expensive Query Plans data set.
- Queries (Normal): The state of the query. The normal status is displayed in green color. This data is derived from the Query State attribute in the MS SQL Expensive Query Plans data set.

Tables

The Tables group widget provides a status summary widget that gives information about the table statuses.

The following KPIs for the widget are derived from the attributes as described:

- Tables (Critical): Number of tables in the critical state. This data is derived from the Fragmentation attribute in the MS SQL Table Detail data set.
- Tables (Normal): Number of tables in the normal state. This data is derived from the Fragmentation attribute in the MS SQL Table Detail data set.
- Tables (Warning): Number of tables in the warning state.

Processes - Detail

Use the Processes - Detail dashboard to view detailed information about all the SQL Server processes, such as the state of the processes and the number of blocked processes.

The following widgets are available in this dashboard page:

Process Details

The Process Details group widget provides a grid widget that gives the status of the processes.

The following KPIs for the widget are derived from the attributes as described:

- Blocking Process ID: The identifier for the process that is blocking a request for a lock. The value format is an alphanumeric string with a maximum of 8 characters; for example, 12. Some blocking processes can become stranded. Investigate eventing thresholds in which a process is being blocked for an extended period of time. For more information on a blocking process, query the sysprocesses and use the sp-lock procedure. This data is derived from the Blocking Process ID attribute in the MS SQL Process Detail data set.
- Client Process ID: The ID the client assigned to the process. The value format is an alphanumeric string with a maximum of 16 characters; for example, amc_2236. This data is derived from the Client Process ID attribute in the MS SQL Process Detail data set.
- Client User ID: The ID of the user executing the process. The value format is an alphanumeric string with a maximum of 8 characters; for example, J_Kelly. This data is derived from the Client User ID attribute in the MS SQL Process Detail data set.
- Command: The name of the command being executed by the process. The value format is an alphanumeric string with a maximum of 16 characters; for example, CREATE VIEW. This data is derived from the Command attribute in the MS SQL Process Detail data set.
- Database Name: The database name. The value format is an alphanumeric string with a maximum of 384 bytes. This data is derived from the Database Name U attribute in the MS SQL Process Detail data set.
- Disk IO: The number of accesses to hard disk since the process started. The value includes accesses to hard disk for physical reads and physical writes. The value format is an integer; for example, 10. This data is derived from the Total Disk IO attribute in the MS SQL Process Detail data set.
- Duration: The time (in seconds) for which the process has been running. This data is derived from the Duration attribute in the MS SQL Process Detail data set.

- ID: The ID of the process. The value format is an alphanumeric string with a maximum of 10 characters; for example, 42168. This data is derived from the Process ID attribute in the MS SQL Process Detail data set.
- Memory Allocated (KB): The number of KB allocated for this process, based on the number of pages in the procedure cache. A negative number indicates that the process is freeing memory allocated by another process. This data is derived from the Total Memory Alloc attribute in the MS SQL Process Detail data set.
- Program Name: The Program Name. The value format is an alphanumeric string with a maximum of 384 bytes. This data is derived from the Program Name U attribute in the MS SQL Process Detail data set.
- Server User ID: The SQL server-assigned ID for the user executing the process. The value format is an alphanumeric string with a maximum of 8 characters; for example, S. This data is derived from the Server User ID attribute in the MS SQL Process Detail data set.
- Status: Indicates the status of the process.
- Total CPU Time (seconds): The amount of CPU time, in seconds, the process has used on the host since the process started. The value format is an integer; for example, 60. This value is based on the statistics collected by the SQL server. Use this value to check for processes that use abnormal amounts of CPU time. This data is derived from the Total CPU Time attribute in the MS SQL Process Detail data set.

Process Summary

The Process Summary group widget provides a stacked bar that gives the details of the processes in percentage.

The following KPIs for the widget are derived from the attributes as described:

- Bad: The number of processes with a status of bad. The value format is an integer; for example, 1. This data is derived from the Total Processes Bad attribute in the MS SQL Process Summary data set.
- Blocked: The number of processes that are being blocked. The value format is an integer; for example, 2. This data is derived from the Total Processes Blocked attribute in the MS SQL Process Summary data set.
- Infected: The number of processes with a status of infected. A process with a status of infected is a process that cannot be completed. The value format is an integer; for example, 1. This data is derived from the Total Processes Infected attribute in the MS SQL Process Summary data set.
- lock sleep: The number of processes with a status of locksleep. A process with a status of locksleep is a process waiting for a lock on a resource to be released. The value format is an integer; for example, 5. This data is derived from the Total Processes Locksleep attribute in the MS SQL Process Summary data set.
- Other sleep: The number of processes with a status of othersleep. The value format is an integer; for example, 2. This data is derived from the Total Other Sleep attribute in the MS SQL Process Summary data set.
- Stopped: The number of processes with a status of stopped. The value format is an integer; for example, 1. This data is derived from the Total Processes Stopped attribute in the MS SQL Process Summary data set.

Processes Blocked (history)

The Processes Blocked (history) group widget provides line graph that gives the details of the blocked processes during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Process Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Blocked: The number of processes that are being blocked. The value format is an integer; for example, 2. This data is derived from the Total Processes Blocked attribute in the MS SQL Process Summary data set.

- **Time:** The local time at the agent when the data was collected. This data is derived from the **Timestamp** attribute in the MS SQL Process Summary data set.
- **Total:** The number of processes. The value includes background processes, processes for applications, and user processes. The value format is an integer; for example, 50. This value includes background processes, processes for applications, and user processes. This data is derived from the **Total Processes** attribute in the MS SQL Process Summary data set.

Server - Detail

Use the Server - Detail dashboard to know the SQL Server details, such as server statistics and services details. Also, determine if any of the SQL Server resources are experiencing problems. Click anywhere on the Availability group widget to drill down to detailed information about the availability groups, replicas, and databases.

The following widgets are available in this dashboard page:

Additional Server Properties

The Addition Server Properties group widget provides an HTML table that contains additional properties for the server.

The following KPIs for the widget are derived from the attributes as described:

- **Current locks:** The number of current locks for the SQL server. The value format is an integer; for example, 73. This data is derived from the **Current Locks** attribute in the MS SQL Server Detail data set.
- **Data cache size (KB):** The number of kilobytes (KB) allocated for the data cache memory. The SQL server uses the data cache to store data and index pages. The value format is an integer; for example, 1000. The cache is sometimes referred to as the buffer cache. This data is derived from the **Data Cache Size** attribute in the MS SQL Server Detail data set.
- **Error log size (Bytes):** The number of bytes in the error log file. The value format is an integer; for example, 50000. The error log contains the fatal error and kernel error messages issued by the SQL server. This data is derived from the **Error Log Size** attribute in the MS SQL Server Detail data set.
- **Maximum locks:** The maximum number of allowable locks. The value format is an integer; for example, 2000. This value is a configuration parameter. This data is derived from the **Max Locks Allowed** attribute in the MS SQL Server Detail data set.
- **Procedure cache size (KB):** The number of kilobytes (KB) that are allocated for the procedure cache. The SQL server uses the procedure cache to compile queries and store procedures that are compiled. The value format is an integer; for example 1000. This data is derived from the **Procedure Cache Size** attribute in the MS SQL Server Detail data set.

Availability

The Availability group widget provides an HTML table that summarizes the details of the availability of the groups.

The following KPIs for the widget are derived from the attributes as described:

- **Non healthy databases:** The total number of the availability databases that have been non-healthy since the agent startup. This data is derived from the **Total Non Healthy Databases** attribute in the MS SQL Availability Database Summary data set.
- **Non healthy groups:** The total number of availability groups that are currently not healthy in the current interval. This data is derived from the **Total Non Healthy Groups** attribute in the MS SQL Availability Groups Summary data set.
- **Non healthy replicas:** The total number of non-healthy replicas in all the availability groups that are hosted on the local server instance. This data is derived from the **Total Non Healthy Replicas** attribute in the MS SQL Availability Replicas Status Summary data set.
- **Total groups:** The total number of availability groups that are available on the SQL Server instance in the current interval. This data is derived from the **Total Groups** attribute in the MS SQL Availability Groups Summary data set.

Procedure Buffer And Cache

The Procedure Buffer and Cache group widget provides a group of volume gauges that summarizes the details of the buffer and cache size of the procedure.

The following KPIs for the widget are derived from the attributes as described:

- Buffer active (%): The percentage of slots with a procedure that is currently executing. Use this attribute to see how much of the cache these procedures use in relation to the cache size as a whole. The value format is an integer. This data is derived from the Procedure Buffers Percent Active attribute in the MS SQL Server Detail data set.
- Buffer used (%): The percentage of slots that have a procedure in them. A procedure buffer is considered used when it is associated with a procedure cache entry. A used procedure buffer can be active or not active. Use this attribute to see how much of the cache is being used by currently compiled procedures in relation to the cache size as a whole. The value format is an integer. This data is derived from the Procedure Buffers Percent Used attribute in the MS SQL Server Detail data set.
- Cache active (%): The total size of the procedure cache in pages. The size of the procedure cache can fluctuate depending on the activity of other database server processes that might require procedure cache slots, such as query plans. Use the attribute to see the current size of the procedure cache. The value format is an integer. This data is derived from the Procedure Cache Percent Active attribute in the MS SQL Server Detail data set.
- Cache used (%): The percentage of the procedure cache that has procedures in it. Use this attribute to see how much of the cache currently executing procedures use in relation to the cache size as a whole. The value format is an integer. This data is derived from the Procedure Cache Percent Used attribute in the MS SQL Server Detail data set.

Server Properties

The Server Properties group widget provides an HTML table that summarizes the properties of the server. For more details about the server, click in the group widget area to open the Server Statistics group widget.

The following KPIs for the widget are derived from the attributes as described:

- Name: The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. This data is derived from the Server attribute in the MS SQL Server Summary data set.
- Operating system type: The operating system for the SQL server. The value format is an alphanumeric string; for example, AIX. This data is derived from the OS Type attribute in the MS SQL Server Detail data set.
- Operating system version: The version of the operating system for the SQL server. The value format is the version in the format version.release; for example, 2. 5. This data is derived from the OS Version attribute in the MS SQL Server Detail data set.
- Startup time: The timestamp that indicates the date and time the SQL server was started. The format is YYYY/MM/DD HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. This data is derived from the Startup Timestamp attribute in the MS SQL Server Detail data set.
- Version: The version of the SQL Server. The value format is the version in the format w. x. y. z; for example, 10. 0. 2531. 0. This data is derived from the SQL Server Ver attribute in the MS SQL Server Summary data set.

Server Services Detail

The Server Services Detail group widget provides a grid widget, which contains detailed information about the services.

The following KPIs for the widget are derived from the attributes as described:

- Name: The service name for the SQL Server. This data is derived from the Service Name attribute in the MS SQL Services Detail data set.
- Start Mode: Defined start mode for the service. This data is derived from the Service Start Mode attribute in the MS SQL Services Detail data set.
- State: The current service state.
- Status: The current service status.
- Type: The service type for this service to the Microsoft SQL Server. This data is derived from the Service Type attribute in the MS SQL Services Detail data set.

Server Statistics Since Startup

The Server Statistics Since Startup group widget provides a grid widget that summarizes the statistics of the server since start up.

The following KPIs for the widget are derived from the attributes as described:

- Average Per Second: The average value per second for the statistic since the SQL server was started. The value format is a decimal with two decimal places allowed; for example, 5.00. This value indicates the norm for the statistic during the current interval. This data is derived from the Average Value per Second attribute in the MS SQL Statistics Detail data set.
- Maximum Seen: The greatest value per second for the statistic since the SQL server was started. The value format is an integer; for example, 2000. This is a benchmark value. This data is derived from the Maximum Seen attribute in the MS SQL Statistics Detail data set.
- Minimum Seen: The smallest value per second for the statistic since the SQL server was started. The value format is an integer; for example, 10. This is a benchmark value. This data is derived from the Minimum Seen attribute in the MS SQL Statistics Detail data set.
- Statistics Name: The name of the statistic. The value format is an alphanumeric string with a maximum of 32 characters; for example, cpu_busy. This data is derived from the Statistic Name attribute in the MS SQL Statistics Detail data set.

Server Performance - Detail

Use the Server Performance - Detail dashboard to know the following details of the SQL servers:

- The CPU and memory usage.
- Historical data of the server performance parameters, such as memory usage and cache usage during a specific duration.
- The top 10 cached query plans according to the performance statistics of the SQL server.

The following widgets are available in this dashboard page:

Cache Hit Ratio (history)

The Cache Hit Ratio (history) group widget provides a line graph that gives the details of the cache hit ratio during a specified period.

The following KPIs for the widget are derived from the attributes as described:

- (%): The current ratio of data cache hits to total requests. Use this attribute to check the effectiveness of the data cache. For optimal performance, the returned value must be approximately 90% or greater. The value format is an integer. This data is derived from the Cache Hit Ratio attribute in the MS SQL Server Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Server Summary data set.

Cache Used (history)

The Cache Used (history) group widget provides a line graph that gives details of the cache used during a specified period.

The following KPIs for the widget are derived from the attributes as described:

- Data: The number of kilobytes (KB) allocated for the data cache memory. The SQL server uses the data cache to store data and index pages. The value format is an integer; for example, 1000.

The cache is sometimes referred to as the buffer cache. This data is derived from the Data Cache Size attribute in the MS SQL Server Summary data set.

- **Procedure:** The number of kilobytes (KB) allocated for the procedure cache. The SQL server uses the procedure cache to compile queries and store procedures that are compiled. The value format is an integer; for example, 1000. This data is derived from the Procedure Cache Size attribute in the MS SQL Server Summary data set.
- **Time:** The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Server Summary data set.

Connection Activity (history)

The Connection Activity (history) group widget provides a line graph that gives details of the connection activity during a specified period.

The following KPIs for the widget are derived from the attributes as described:

- **Current:** The number of active connections (logons). The value format is an integer; for example, 50. This data is derived from the Current Logons attribute in the MS SQL Statistics Summary data set.
- **Maximum:** The maximum number of active connections (logons) allowed for the SQL server. The value format is an integer in the range 5 through 2147483647; for example, 100. This is a configuration value for the SQL server. To determine the number of connections that can be configured for the SQL server, use the select @@max_connections command. This data is derived from the Max User Connections Allowed attribute in the MS SQL Statistics Summary data set.
- **Time:** The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Statistics Summary data set.

CPU Used

The CPU Used group widget provides a volume gauge that gives details of the CPU usage.

The following KPIs for the widget are derived from the attributes as described:

- **Application processes (%):** The percentage of CPU time the SQL server application processes are using on the host. The value format is a percentage with two decimal places allowed; for example, 20.00. This data is derived from the Server CPU Percent Application attribute in the MS SQL Process Summary data set.
- **IO operations (%):** The percentage of time used for I/O operations during the current monitoring interval. Use this attribute to gauge how much of the CPU resource the database server uses for I/O so you can allocate resources more efficiently. You also can use this attribute to determine how I/O resource-intensive certain operations are. The value format is an integer. This data is derived from the Server CPU Percent IO attribute in the MS SQL Server Summary data set.
- **Operating system (%):** The percentage of CPU time being used by all processes on the host. The value format is a percentage with two decimal places allowed; for example, 40. 00. This data is derived from the Total OS CPU Percent attribute in the MS SQL Server Summary data set.
- **System processes (%):** The percentage of CPU time the SQL server processes are using on the host. The value format is a percentage with two decimal places allowed; for example, 25.00. This data is derived from the Server CPU Percent System attribute in the MS SQL Process Summary data set.

Expensive Queries

The Expensive Queries group widget provides a grid widget that gives the top 10 cached query plans according to the performance statistics of the SQL server.

The following KPIs for the widget are derived from the attributes as described:

- **Age (minutes):** The time interval (in minutes) between the last time that the query was executed and the time when the query was created. This data is derived from the Age attribute in the MS SQL Expensive Query Plans data set.

- Avg. CPU Time: The average CPU time (in milliseconds) that is required to execute the plan since the plan was compiled. This data is derived from the Average CPU Time attribute in the MS SQL Expensive Query Plans data set.
- Avg. Duration: The average elapsed time (in milliseconds) that is required to execute this plan. This data is derived from the Average Duration attribute in the MS SQL Expensive Query Plans data set.
- Avg. Logical Reads: The average logical reads that are performed to execute this plan since the plan was compiled. This data is derived from the Average Logical Reads attribute in the MS SQL Expensive Query Plans data set.
- Avg. Logical Writes: The average logical writes that are performed to execute this plan since the plan was compiled. This data is derived from the Average Logical Writes attribute in the MS SQL Expensive Query Plans data set.
- Avg. Physical Reads: The average physical reads that are performed to execute this plan since the plan was compiled. This data is derived from the Average Physical Reads attribute in the MS SQL Expensive Query Plans data set.
- Database: The name of the database against which the query is executed. This data is derived from the Database Name attribute in the MS SQL Expensive Query Plans data set.
- Execution Rate: The value in percentage that is calculated by dividing the execution count by the value that is obtained after subtracting the creation time of the query from the last execution time of the query. This data is derived from the Execution rate attribute in the MS SQL Expensive Query Plans data set.
- Expensive Queries Based On: The sorting criteria for top N expensive queries. This data is derived from the Sorting Criteria attribute in the MS SQL Expensive Query Plans data set.
- Query Text: The SQL text of the expensive query plan. This data is derived from the Query Text attribute in the MS SQL Expensive Query Plans data set.
- State: The state of the query that is based on the threshold (Normal/Critical). This data is derived from the Query State attribute in the MS SQL Expensive Query Plans data set.
- This Group Widget Shows Top: The number of expensive queries that are displayed in the group widget. This data is derived from the Query State attribute in the MS SQL Expensive Query Plans data set.

Memory Used (history)

The Memory Used (history) group widget provides a line graph that gives the details of the memory usage during a specified period.

The following KPIs for the widget are derived from the attributes as described:

- KB: The total amount of dynamic memory (in KB) that the server is using currently. This data is derived from the Total Server Memory attribute in the MS SQL Server Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Server Summary data set.

Memory Used By Server Components

The Memory Used By Server Components group widget provides an HTML table that gives details of the memory usage by server components.

The following KPIs for the widget are derived from the attributes as described:

- Connection (MB): The total amount of dynamic memory the server is using for maintaining connections in megabytes. The format is an integer. This data is derived from the Connection Memory attribute in the MS SQL Memory Manager data set.
- Lock memory (MB): The total amount of dynamic memory the server is using for locks in megabytes. The format is an integer. This data is derived from the Lock Memory attribute in the MS SQL Memory Manager data set.

- Optimizer (MB): The total amount of dynamic memory the server is using for query optimization in megabytes. The format is an integer. This data is derived from the Optimizer Memory attribute in the MS SQL Memory Manager data set.
- SQL cache (MB): The total amount of dynamic memory the server is using for the dynamic SQL cache in megabytes. The format is an integer. This data is derived from the SQL Cache Memory attribute in the MS SQL Memory Manager data set.
- Target server (MB): The total amount of dynamic memory the server is willing to consume in megabytes. The format is an integer. This data is derived from the Target Server Memory attribute in the MS SQL Memory Manager data set.

Network Read And Write Per Second (history)

The Network Reads And Writes Per Second (history) group widget provides a line graph that gives details of the network reads and writes per second during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Statistics Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Read: The rate at which tabular data stream (TDS) packets are read from the network, in packets per second. This statistic is an indicator of network throughput. When this statistic is high, it indicates heavy network traffic. The value format is an integer. This data is derived from the Network Read Rate attribute in the MS SQL Statistics Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Statistics Summary data set.
- Write: The rate at which tabular data stream (TDS) packets are written to the network, in packets per second. This statistic is an indicator of network throughput. When this statistic is high, it indicates heavy network traffic. The value format is an integer. This data is derived from the Network Write Rate attribute in the MS SQL Statistics Summary data set.

Physical Reads And Writes Per Second (history)

The Physical Reads And Writes Per Second (history) group widget provides a line graph that gives details of the physical reads and writes per second during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Statistics Summary attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Read: The average number of physical reads per second during the current interval. The value format is a decimal with two decimal places allowed; for example, 5.00. This data is derived from the Physical Reads per Second attribute in the MS SQL Statistics Summary data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Statistics Summary data set.
- Write: The average number of physical writes per second during the current interval. The value format is a decimal with two decimal places allowed; for example, 5.00. This data is derived from the Physical Writes per Second attribute in the MS SQL Statistics Summary data set.

Tables - Detail

Use the Tables - Detail dashboard to know the details of tables, such as properties of the fragmented tables and table optimization statistics. Also, use this dashboard to determine if any of the table resources are experiencing problems. Click the Database name in the Databases group widget to get detailed information about the tables in that database. The detailed information about the tables in a database appears in the Tables Status group widget.

The following widgets are available in this dashboard page:

Databases

The Databases group widget provides a grid widget that lists the database names.

The following KPIs for the widget are derived from the attributes as described:

- Name: The name of the database. Each database name is unique. The SQL server also assigns each database its own identification number. This data is derived from the Database Name U attribute in the MS SQL Database Detail data set.

Fragmented Tables - Top 5

The Fragmented Tables - Top 5 group widget provides a grid widget that gives information about the fragmentation details.

The following KPIs for the widget are derived from the attributes as described:

- Database Name: The Database name. The value format is an alphanumeric string with a maximum of 30 characters. This data is derived from the Database Name attribute in the MS SQL Table Detail data set.
- Fragmentation (%): The degree of fragmentation for the table. Use this attribute to determine if fragmentation is reaching a level that causes performance degradation. Performance degradation is due to non-contiguous table reads that require additional extent switches. The format is an integer. This data is derived from the Fragmentation attribute in the MS SQL Table Detail data set.
- Name: The table name. The value format is an alphanumeric string with a maximum of 30 characters. This data is derived from the Table Name attribute in the MS SQL Table Detail data set.

Optimized Age Tables - Top 5

The Optimized Age Tables - Top 5 group widget provides a grid widget that gives statistics about optimization.

The following KPIs for the widget are derived from the attributes as described:

- Database Name: The Database name. The value format is an alphanumeric string with a maximum of 30 characters. This data is derived from the Database Name attribute in the MS SQL Table Detail data set.
- Name: The table name. The value format is an alphanumeric string with a maximum of 30 characters. This data is derived from the Table Name attribute in the MS SQL Table Detail data set.
- Optimized Age (minutes): The time (in minutes) since statistics were last updated for the table. Use this attribute to ensure that queries base their query plans on up-to-date information. If query plans are based on old information, they might be inefficient. The value for this attribute is displayed as 0 (zero) when the time since statistics was last updated for the table is less than 1 minute and is displayed as Not Collected when the statistics for the table was never updated. The format is an integer. This data is derived from the Optimizer Statistics Age attribute in the MS SQL Table Detail data set.

Table Status

The Table Status group widget provides a grid widget that gives information about the top 50 tables with the highest percentage of fragmentation. When you click a database name in the Database group widget, this group widget displays the top 10 tables with the highest percentage of fragmentation in the database. This is the default behavior of the group widget. However, you can change the sorting criteria at the agent side by setting the value for COLL_TBLD_SORTBY environment variable.

The following KPIs for the widget are derived from the attributes as described:

- File Table: Indicates whether the current table is a FileTable. The format is an integer. This data is derived from the FileTable attribute in the MS SQL Table Detail data set.
- Fragmentation (%): The degree of fragmentation for the table. Use this attribute to determine if fragmentation is reaching a level that causes performance degradation. Performance degradation is due to non-contiguous table reads that require additional extent switches. The format is an integer. This data is derived from the Fragmentation attribute in the MS SQL Table Detail data set.

- Index ID: The ID number of this table index. This data is derived from the Index Id attribute in the MS SQL Table Detail data set.
- Index Name: The table index name. This data is derived from the Index Name attribute in the MS SQL Table Detail data set.
- Name: The table name. The value format is an alphanumeric string with a maximum of 30 characters. This data is derived from the Table Name attribute in the MS SQL Table Detail data set.
- Optimized Age (minutes): The time (in minutes) since statistics were last updated for the table. Use this attribute to ensure that queries base their query plans on up-to-date information. If query plans are based on old information, they might be inefficient. The value for this attribute is displayed as 0 (zero) when the time since statistics was last updated for the table is less than 1 minute and is displayed as Not Collected when the statistics for the table was never updated. The format is an integer. This data is derived from the Optimizer Statistics Age attribute in the MS SQL Table Detail data set.
- Row Count: The number of rows that are present in a database table. The format is an integer. This data is derived from the Number Of Rows attribute in the MS SQL Table Detail data set.
- Space Used (MB): The amount of space (in megabytes) used by the specified table. Use this attribute to determine how much of the database space is used by a table. You can monitor the growth of individual tables and compare the actual growth to the expected growth. The format is an integer. This data is derived from the Space Used attribute in the MS SQL Table Detail data set.

Additional group widgets

These pop-up group widgets are displayed after you click a group widget for more details. Some group widgets have links to more granular information in a popup widget, described here.

The following widgets are available in this dashboard page:

dummy widget

Description needed in JSON

The following KPIs for the widget are derived from the attributes as described:

Additional Job Detail

The Additional Job Detail group widget provides an HTML table that contains additional details of the job that you select on the Job Status group widget.

The following KPIs for the widget are derived from the attributes as described:

- Current status: The current job status. This data is derived from the Current Status attribute in the MS SQL Job Detail data set.
- Current step: The current step that is being executed in the job. The value format is an alphanumeric string with a maximum of 128 characters. This data is derived from the Current Execution Step attribute in the MS SQL Job Detail data set.
- Execution duration: The time (in seconds) that has elapsed since the job started. This data is derived from the Job Execution Duration attribute in the MS SQL Job Detail data set.

Additional Locks Details

The Additional Locks Details group widget provides an HTML table that summarizes the details of the Locks.

The following KPIs for the widget are derived from the attributes as described:

- Database having maximum blocks: The name of the database blocking the largest number of processes. This data is derived from the Database Max blocks U attribute in the MS SQL Server Enterprise View data set.
- Database having maximum locks: The name of the database locking the largest number of processes. This data is derived from the Database Max Locks U attribute in the MS SQL Server Enterprise View data set.

- Lock conflicts: The total number of processes involved in lock conflicts. The value format is an integer. This data is derived from the Total Lock Conflicts attribute in the MS SQL Server Enterprise View data set.
- Locks: The total number of locks for the server. The value format is an integer. This data is derived from the Total Locks attribute in the MS SQL Server Enterprise View data set.
- Table having maximum locks: The name of the table with largest number of locks. Valid values include text strings with a maximum of 32 bytes. This data is derived from the Table Max Locks U attribute in the MS SQL Server Enterprise View data set.

Database Used (history)

To open the Database Used (history) group widget, click the Database Used column of the Database Status group widget. The Database Used (history) group widget summarizes the percentage of space that is used by a database during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Database Detail attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Data: The amount of space used in the database as a percentage of total space allowed. This data is derived from the Database Space Percent Used attribute in the MS SQL Database Detail data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Database Detail data set.

Device Details

The Device Details group widget is a window that provides an HTML table with more information about a device.

The following KPIs for the widget are derived from the attributes as described:

- Mirror device name: The name of the mirror device for the database. The value format is an alphanumeric string with a maximum of 64 characters; for example, /dev/rsd2g . The mirror device duplicates the contents of a primary device. Refer to the logical and physical names of the device that is being mirrored. This data is derived from the Mirror Device Name attribute in the MS SQL Device Detail data set.
- Physical device name: The Name of the physical device. Valid values include text strings with a maximum of 192 bytes. This data is derived from the Physical Device Name U attribute in the MS SQL Device Detail data set.

Devices

The Devices group widget provides a status summary widget that summarizes the percentage of the space used by the devices.

The following KPIs for the widget are derived from the attributes as described:

- Devices (Warning): Indicates that the device is in the warning state. This data is derived from the Device Name U attribute in the MS SQL Device Detail data set.
- State: The current device state.

File Names

The File Names group widget provides an HTML table that shows the path to the data file and log file.

The following KPIs for the widget are derived from the attributes as described:

- Data file name: The names of all physical files for data that make up this database. This data is derived from the Data File Names attribute in the MS SQL Database Detail data set.
- Log file name: The names of all physical files for logs that make up this database. This data is derived from the Log File Names attribute in the MS SQL Database Detail data set.

Fragmentation Trend (history)

The Fragmentation Trend (history) group widget provides a line graph that displays the fragmentation trend of the tables during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Table Detail attribute group.

The following KPIs for the widget are derived from the attributes as described:

- (%) : The degree of fragmentation for the table. Use this attribute to determine if fragmentation is reaching a level that causes performance degradation. Performance degradation is due to non-contiguous table reads that require additional extent switches. The format is an integer. This data is derived from the Fragmentation attribute in the MS SQL Table Detail data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Table Detail data set.

Jobs

The Jobs group widget provides a status summary widget that summarizes the status of the jobs and a status gauge that displays the status of the failed jobs.

The following KPIs for the widget are derived from the attributes as described:

- Failed: Reports any jobs that are run by the SQL Server and failed in the last monitoring interval. This data is derived from the Job Id attribute in the MS SQL Job Detail data set.
- Failed jobs in last interval: Reports any jobs run by the SQLServerAgent service that have failed in the last monitoring interval. The first time you retrieve this attribute, it returns 0. The next time you retrieve it, it returns the total number of SQLServerAgent failed jobs found in the system history tables since the first run. These jobs include replication and user-defined jobs, such as maintenance or backup tasks. Use this attribute to alert you when scheduled tasks have failed. The value format is an integer. This data is derived from the SQLServerAgent Failed Jobs attribute in the MS SQL Server Summary data set.
- Succeeded/In-progress: Reports the number of jobs that have succeeded or are currently in progress. This data is derived from the Job Id attribute in the MS SQL Job Detail data set.
- Warning: Indicates the number of jobs that are in the warning state.

Log Used (history)

To see a line graph in the Log Used (history) group widget, click the Log Utilization column of the Database Status group widget. The Log Used (history) group widget summarizes the percentage of space that is used by the transaction log during a specified period.

The line chart is displayed only when the historical data collection is configured for the MS SQL Database Detail attribute group.

The following KPIs for the widget are derived from the attributes as described:

- Log Used(%): The percentage of the transaction log that has been used. This data is derived from the Log Space Percent Used attribute in the MS SQL Database Detail data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Database Detail data set.

Retry Attempts (history)

To open the Retry Attempts (history) group widget, click the Retry Attempts column of the Jobs Status group widget. The Retry Attempts (history) group widget provides a line graph that contains the number of retry attempts during a specified period.

The following KPIs for the widget are derived from the attributes as described:

- Count: The current number of retry attempts that are done on a step of a running job. This data is derived from the Current Retry Attempt attribute in the MS SQL Job Detail data set.
- Time: The local time at the agent when the data was collected. This data is derived from the Timestamp attribute in the MS SQL Job Detail data set.

Server Services

The Server Services group widget provides a status summary widget that summarizes the status of the SQL server services.

The following KPIs for the widget are derived from the attributes as described:

- Services (Critical): Indicates that the services status is critical. This data is derived from the Service Name attribute in the MS SQL Services Detail data set.
- Services: The current service status. This data is derived from the Service Name attribute in the MS SQL Services Detail data set.
- Services (Warning): Indicates that the services status is warning.

Server Statistics

The Server Statistics group widget provides an HTML table that summarizes the statistics of the server.

The following KPIs for the widget are derived from the attributes as described:

- Edition: The installed product edition of this instance of the SQL Server. This data is derived from the Edition attribute in the MS SQL Server Properties data set.
- Engine edition: The database engine edition of the instance of the SQL Server that is currently installed. The value format is an integer. This data is derived from the Engine Edition attribute in the MS SQL Server Properties data set.
- Filestream level: The current level of FileStream support that is enabled for the SQL Server instance. This data is derived from the FileStream Level attribute in the MS SQL Server Properties data set.
- FileStream share name: The Windows share name where the FileStream data is stored. The value format is an alphanumeric string. This data is derived from the FileStream Share Name attribute in the MS SQL Server Properties data set.
- Is clustered: Indicates whether the SQL Server instance is configured in the Windows Server Failover Clustering (WSFC) cluster. The value format is an integer. This data is derived from the Is Clustered attribute in the MS SQL Server Properties data set.
- Is integrated security only: The server is in the integrated security mode. The value format is an integer. This data is derived from the IsIntegrated Security Only attribute in the MS SQL Server Properties data set.
- Is single user: The server is currently in the single-user mode, in which only a single user can connect to the server. The value format is an integer. This data is derived from the Is Single User attribute in the MS SQL Server Properties data set.
- Licence type: The mode of this instance of the SQL Server. The mode can be per-seat or per-processor. This data is derived from the License Type attribute in the MS SQL Server Properties data set.
- Number of licences: The number of client licenses that are currently registered for this instance of the SQL Server if the SQL Server is in the per-seat mode. The number of processors that are currently licensed for this instance of the SQL Server if the SQL Server is in the per-processor mode. This data is derived from the Num Licenses attribute in the MS SQL Server Properties data set.
- Physical netBIOS name: The NetBIOS name of the machine where this instance of the SQL Server is currently running. This data is derived from the Physical NetBIOS Name attribute in the MS SQL Server Properties data set.
- Product level: The level of the version of this instance of the SQL Server, for example, the original release version and the service pack version. This data is derived from the Product Level attribute in the MS SQL Server Properties data set.
- Qualified server name: The name of the qualified SQL Server instance in the format hostname:SQL server instance. For the default SQL Server instance, the format is hostname. The value format is an alphanumeric string with a maximum of 128 characters. This data is derived from the Qualified Server Name attribute in the MS SQL Server Properties data set.

- Server collation: The name of the default collation for the server. The value format is a string. This data is derived from the Collation attribute in the MS SQL Server Properties data set.
- Time since startup (minutes): The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360. This data is derived from the Time Since Startup attribute in the MS SQL Server Summary data set.

Custom views

After you select an application that includes a Microsoft SQL Server managed resource, the **Custom Views** tab is available for displaying and building custom dashboard pages with attribute values from the Microsoft SQL Server agent. You can quickly build monitoring pages for an application and save them for viewing.

Only a subset of Microsoft SQL Server agent attributes, which are the most useful for reporting, are available for custom views. These attributes are shown in *italic* in [Chapter 4, “Attributes,” on page 43](#).

Chapter 3. Thresholds

Thresholds test for certain conditions on your managed resources, such as memory usage over 95%, and raise an event when the conditions have been met. The agent comes with predefined thresholds that you can use to monitor your Microsoft SQL Server environment. You can create additional thresholds for the areas of interest.

After you click  **System Configuration > Threshold Manager**, select **Microsoft SQL Server** as the data source type to see all the available thresholds.

Predefined thresholds

The thresholds are organized in the Cloud APM console **Threshold Manager** by the data set for which they were written. The Microsoft SQL Server agent has the following predefined thresholds:

MS_SQL_Cache_Avg_FreePage_Crit

Average number of cache buffers scanned to find a free buffer

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Cache_Average_Free_Scan *GT 15
```

This threshold is evaluated every 9 minutes 37 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Cache_Average_Free_Scan[SRVS.CACHAVGFCN].

MS_SQL_Cache_Hit_Ratio_Crit

Percentage of data cache hits to total data requests

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Cache_Hit_Ratio *LT 70 *AND *VALUE  
MS_SQL_Server_Summary.Cache_Hit_Ratio *GE 0
```

This threshold is evaluated every 3 minutes 30 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Cache_Hit_Ratio[SRVS.CACHHITR].

MS_SQL_Cache_Hit_Ratio_Warn

The cache hit ratio is 70 - 90.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Cache_Hit_Ratio *LT 90.00 *AND *VALUE  
MS_SQL_Server_Summary.Cache_Hit_Ratio *GE 70.00
```

This threshold is evaluated every 3 minutes 25 seconds.

The severity of this threshold is Warning.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Cache_Hit_Ratio[SRVS.CACHHITR].

MS_SQL_Client_Cnt_Pct_Used_Crit

Percentage of available client licenses being used

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Statistics_Summary.Client_Count_Percent_Used *GT 90.00
```

This threshold is evaluated every 9 minutes 47 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Statistics_Summary.Client_Count_Percent_Used[STATS.CLNPCT].

MS_SQL_Collstatus_Critical

The status of collector process is inactive.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Collection_Status *EQ Inactive
```

This threshold is evaluated every 3 minutes 5 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Collection_Status[SRVS.COLSTATUS].

MS_SQL_Collstatus_Warning

The status of collector process is unknown.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Collection_Status *EQ Unknown
```

This threshold is evaluated every 3 minutes 10 seconds.

The severity of this threshold is Warning.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Collection_Status[SRVS.COLSTATUS].

MS_SQL_DB_Status_Crit

Database is unavailable due to an error or from being offline

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Detail.Database_Status *NE Available
```

This threshold is evaluated every 4 minutes 43 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Database_Name_U attribute.

This threshold uses the following attributes:

MS_SQL_Database_Detail.Database_Status[DBD.DBSTAT],

MS_SQL_Database_Detail.Database_Name_U[DBD.UDBNAME].

MS_SQL_DB_Suspect_Crit

Database is inconsistent due to a recovery failure

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Detail.Suspect_Database *EQ True
```

This threshold is evaluated every 4 minutes 43 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Database_Name_U attribute.

This threshold uses the following attributes:

MS_SQL_Database_Detail.Suspect_Database[DBD.SUSPECTDB] (not visible in the UI),
MS_SQL_Database_Detail.Database_Name_U[DBD.UDBNAME].

MS_SQL_Errors_Critical

The number of SQL server errors in the current sampling interval is one or greater.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Problem_Summary.Total_Errors_Current_Interval *GE 1
```

This threshold is evaluated every 3 minutes 15 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Problem_Summary.Total_Errors_Current_Interval[PROBS.NUMCURR].

MS_SQL_EXPENSIVE_QRY_STATE_CRIT

Blank Description Text

The default configuration has the following SQL syntax:

```
*VALUE MS_SQL_Expensive_Query_Plans.Query_State *EQ Critical
```

This threshold is evaluated every 3 minutes.

The severity of this threshold is CRITICAL .

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Expensive_Query_Plans.Query_State[EXPQP.QUERYSTATE].

MS_SQL_IO_Disk_Errors_Crit

SQL Server disk read or write errors

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Statistics_Detail.Statistic_Name *EQ 'io errors' *AND  
*VALUE MS_SQL_Statistics_Detail.Current_Value *GT 0
```

This threshold is evaluated every 9 minutes 53 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Statistics_Detail.Statistic_Name[STATD.STATNAME],
MS_SQL_Statistics_Detail.Current_Value[STATD.STATCURR] (not visible in the UI).

MS_SQL_Jobs_Critical

The number of jobs that failed in the current sampling interval is one or greater.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Job_Summary.Jobs_Failed_Current_Interval *GE 1
```

This threshold is evaluated every 6 minutes 10 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Job_Summary.Jobs_Failed_Current_Interval[JOBS.NUMJBADCUR].

MS_SQL_Jobs_Warning

The number of jobs with canceled or unknown status in the current sampling interval is one or greater and there are no failed jobs in the current sampling interval.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Job_Summary.Jobs_Failed_Current_Interval *EQ 0 *AND *VALUE MS_SQL_Job_Summary.Warning_Jobs_Current_Interval *GE 1
```

This threshold is evaluated every 6 minutes 15 seconds.

The severity of this threshold is Warning.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Job_Summary.Jobs_Failed_Current_Interval[JOBS.NUMJBADCUR],
MS_SQL_Job_Summary.Warning_Jobs_Current_Interval[JOBS.WARNJOBCUR].

MS_SQL_LogonPct_Crit

Percent active user connection to available connections

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Statistics_Summary.Percent_Max_Logons_Active *GT 90
```

This threshold is evaluated every 29 minutes 49 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Statistics_Summary.Percent_Max_Logons_Active[STATS.LOGONPCT].

MS_SQL_Mirror_Failover_Critical

Mirroring State is Pending Failover

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Mirroring.Mirroring_State *EQ 3
```

This threshold is evaluated every 15 minutes.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Database_Mirroring.Mirroring_State[DBMIR.MIRSTATE] (not visible in the UI).

MS_SQL_Num_Process_Blocked_Crit

Number of blocked processes

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Process_Summary.Total_Processes_Blocked *GT 10
```

This threshold is evaluated every 9 minutes 59 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Process_Summary.Total_Processes_Blocked[PRCS.NUMBLOCK].

MS_SQL_Oldest_Transaction_Crit

Age (in minutes) of oldest database transaction

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Detail.Oldest_Open_Transaction *GT 15
```

This threshold is evaluated every 4 minutes 53 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Database_Name_U attribute.

This threshold uses the following attributes:

MS_SQL_Database_Detail.Oldest_Open_Transaction[DBD.OLDOPTRAN],
MS_SQL_Database_Detail.Database_Name_U[DBD.UDBNAME].

MS_SQL_PCT_MAX_Locks_Critical

Percent of max Locks is \geq 80.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Enterprise_View.Percent_Max_Locks *GE 80
```

This threshold is evaluated every 14 minutes 13 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Enterprise_View.Percent_Max_Locks[SRVRE.PCTLOCKS] (not visible in the UI).

MS_SQL_Pct_ProcBlock_Warn

The percentage of process blockages is 50 or greater.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Process_Summary.Percent_Processes_Blocked *GE 50.00
```

This threshold is evaluated every 3 minutes 20 seconds.

The severity of this threshold is Warning.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Process_Summary.Percent_Processes_Blocked[PRCS.PCTBLOCK].

MS_SQL_Proc_Buffs_Active_Crit

Percent of active procedure buffers

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Detail.Procedure_Buffers_Percent_Active *LT 75 *AND  
*VALUE MS_SQL_Server_Detail.Procedure_Buffers_Percent_Active *GE 0
```

This threshold is evaluated every 10 minutes 1 second.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Detail.Procedure_Buffers_Percent_Active[SRVD.PCTPROCBA].

MS_SQL_Proc_Buffs_Used_Crit

Percent of used procedure buffers

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Detail.Procedure_Buffers_Percent_Used *LT 75 *AND  
*VALUE MS_SQL_Server_Detail.Procedure_Buffers_Percent_Used *GE 0
```

This threshold is evaluated every 10 minutes 1 second.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Detail.Procedure_Buffers_Percent_Used[SRVD.PCTPROCBU].

MS_SQL_Proc_Cache_Active_Crit

Percent of procedure cache with executing procedures

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Detail.Procedure_Cache_Percent_Active *LT 75 *AND  
*VALUE MS_SQL_Server_Detail.Procedure_Cache_Percent_Active *GE 0
```

This threshold is evaluated every 10 minutes 7 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Detail.Procedure_Cache_Percent_Active[SRVD.PCTPROCCA].

MS_SQL_Proc_Cache_Used_Crit

Percent of procedure cache buffers with procedures

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Detail.Procedure_Cache_Percent_Used *LT 75 *AND  
*VALUE MS_SQL_Server_Detail.Procedure_Cache_Percent_Used *GE 0
```

This threshold is evaluated every 10 minutes 7 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Detail.Procedure_Cache_Percent_Used[SRVD.PCTPROCCU].

MS_SQL_Rem_Serv_Stat_Critical

Where Remote Server Status is Inactive.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Remote_Servers.Remote_Server_Status *EQ Inactive
```

This threshold is evaluated every 14 minutes 17 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Remote_Servers.Remote_Server_Status[SRVR.RSRVSTATUS] (not visible in the UI).

MS_SQL_Repl_Latency_Crit

Replication processing delay (in milliseconds)

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Detail.Replication_Latency *GT 900000
```

This threshold is evaluated every 5 minutes 7 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Database_Name_U attribute.

This threshold uses the following attributes:

MS_SQL_Database_Detail.Replication_Latency[DBD.REPTLAT],
MS_SQL_Database_Detail.Database_Name_U[DBD.UDBNAME].

MS_SQL_Status_Critical

Where Server status not Active.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Server_Status *NE Active
```

This threshold is evaluated every 3 minutes.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Server_Status[SRVS.SRVSTATUS].

MS_SQL_Total_Locks_Critical

Total Number Of Locks Critical >= 4000

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Enterprise_View.Total_Locks *GE 4000
```

This threshold is evaluated every 14 minutes 17 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Enterprise_View.Total_Locks[SRVRE.NUMLOCKS].

MS_SQL_V1_DB_Critical

Declares a critical condition when the number of databases that are in the critical status is one or greater. A database is in the critical state when the database state is offline or emergency, or the status of database is not available or error, or the percentage space usage of database log files is greater than 90, or the percentage space usage of database data files is greater than 90.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Summary.Total_Databases_Critical *GE 1
```

This threshold is evaluated every 6 minutes.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Database_Summary.Total_Databases_Critical[DBS.TDBCRIT].

MS_SQL_V1_DB_Spc_Pct_Usd_Crit

Declares a critical condition when the percentage usage of the database space compared to the total permitted space is greater than 90. The autogrowth property of the data files is used to calculate the total permitted space.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Detail.Database_Space_Percent_Used *GT 90
```

This threshold is evaluated every 5 minutes.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Database_Name_U attribute.

This threshold uses the following attributes:

MS_SQL_Database_Detail.Database_Space_Percent_Used[DBD.PCTDBUSE],

MS_SQL_Database_Detail.Database_Name_U[DBD.UDBNAME].

MS_SQL_V1_DB_Warning

Declares a warning condition when the number of databases that are in the warning status is one or greater. A database is in the warning state when the database state is restoring or recovering or

recovery pending or suspect, or the percentage space usage of database log file is greater than 70 but less than 90, or the percentage space usage of database data files is greater than 70 but less than 90.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Summary.Total_Databases_Critical *EQ 0 *AND  
*VALUE MS_SQL_Database_Summary.Total_Databases_Warning *GE 1
```

This threshold is evaluated every 6 minutes 5 seconds.

The severity of this threshold is Warning.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Database_Summary.Total_Databases_Critical[DBS.TDBCRIT],
MS_SQL_Database_Summary.Total_Databases_Warning[DBS.TDBWARN].

MS_SQL_V1_Fragmentation_Crit

Declares a critical condition when the percentage of database table fragmentation is greater than 90.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Table_Detail.Fragmentation *GT 90
```

This threshold is evaluated every 29 minutes 47 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Table_Name attribute.

This threshold uses the following attributes: MS_SQL_Table_Detail.Fragmentation[TBLD.PCTFRAG],
MS_SQL_Table_Detail.Table_Name[TBLD.UTBLNAME].

MS_SQL_V1_Log_Spc_Pct_Usd_Crit

Declares a critical condition when the percentage usage of the transaction log space compared to the maximum permitted log space is greater than 90. The maximum allowed log space is calculated considering autogrowth property of the log files. The autogrowth property of the log files is used to calculate the maximum permitted log space.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Database_Detail.Log_Space_Percent_Used *GT 90
```

This threshold is evaluated every 5 minutes.

The severity of this threshold is CRITICAL.

The threshold is evaluated for each distinct value of the Database_Name_U attribute.

This threshold uses the following attributes:

MS_SQL_Database_Detail.Log_Space_Percent_Used[DBD.PCTLOGUSED],
MS_SQL_Database_Detail.Database_Name_U[DBD.UDBNAME].

MS_SQL_V1_Netw_Read_Rate_Warn

Declares a warning condition when the number of read operations of the data stream packet per second exceeds 150.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Statistics_Summary.Network_Read_Rate *GT 150
```

This threshold is evaluated every 3 minutes.

The severity of this threshold is WARNING.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Statistics_Summary.Network_Read_Rate[STATS.NTWKRDRATE].

MS_SQL_V1_Netw_Write_Rate_Warn

Declares a warning condition when the number of write operations of the data stream packet per second exceeds 150.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Statistics_Summary.Network_Write_Rate *GT 150
```

This threshold is evaluated every 3 minutes.

The severity of this threshold is WARNING.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Statistics_Summary.Network_Write_Rate[STATS.NTWKWRATE].

MS_SQL_V1_Server_CPU_Critical

Declares a critical condition when 80 percent or more of the processor time is used by the SQL server process on the host, and when the age of the server is 10 minutes or greater.

The default configuration has the following SQL syntax:

```
*IF *VALUE MS_SQL_Server_Summary.Server_CPU_Percent *GE 80 *AND *VALUE  
MS_SQL_Server_Summary.Time_Since_Startup *GE 10
```

This threshold is evaluated every 13 minutes 59 seconds.

The severity of this threshold is CRITICAL.

The threshold is evaluated for the table.

This threshold uses the following attributes:

MS_SQL_Server_Summary.Server_CPU_Percent[SRVS.PCTCPU],

MS_SQL_Server_Summary.Time_Since_Startup[SRVS.SRVAGE].

Customized thresholds

You can use the predefined thresholds as a starting point for event monitoring, and create your own thresholds as conditions arise that you want to monitor.

The Microsoft SQL Server agent has many data sets that you can use to create thresholds to monitor for specific conditions. For descriptions of the data sets, see [Chapter 4, “Attributes,” on page 43](#).

Tip: The hover help for the **Threshold Editor Data set** field has a *Learn more* link to the attribute descriptions for the selected data set.

Chapter 4. Attributes

Attributes are the application properties that are being measured and reported by the Monitoring Agent for Microsoft SQL Server. Attributes make up the key performance indicators (KPIs) that are reported, and you can use them to create thresholds for conditions that you want to monitor.

About attributes

Attributes are organized into *data sets* (also referred to as *attribute groups*). The values can be selectively displayed in dashboard pages or used to define a threshold.

The most recent data sample of the attributes in the data set are used after you open a dashboard page or start a threshold.

Dashboard pages

Only a subset of Microsoft SQL Server agent attributes is displayed in the dashboard pages. Queries to the dashboard data provider specify which attribute values to request from the managed resource.

These attributes are shown in *italic* in this chapter. You can use these attributes to create the charts and tables in custom dashboard pages.

Thresholds

You can define thresholds that monitor the state of your operating system, database, or application and open an event when the threshold is exceeded. You use attributes to define thresholds that describe a condition that you want to test. After the threshold is started, the attribute values that are specified in the threshold are compared with the values collected by the Microsoft SQL Server agent. After the condition is met, an event is registered and you are alerted by indicators in the Application Performance Dashboard navigator, **All My Applications** summary boxes, and the **Events** tab.

The Microsoft SQL Server agent comes with *predefined thresholds* that are enabled and started with the agent. If you edit a predefined threshold, such as to change the condition or severity, it is no longer treated as a predefined threshold but considered a *custom threshold*.

All Microsoft SQL Server agent attributes, unless otherwise noted, can be used to create custom thresholds. The **Events** tab has a table of open events with information, including threshold name, severity, source, and display item. You can expand an event row to see the formula and drill down to the dashboard page for the managed resource.

Some attributes names display differently in the Threshold Editor, as shown in parentheses after the name, such as "Object Count (OBJECT_COUNT)".

Historical data configurations

The Microsoft SQL Server agent collects historical data for key data sets that are shown in the dashboard pages. A page that includes historical views from the managed resource instance has a time selector tool for adjusting the time range. With line charts, you can also compare the values with a previous day, up to the number of days that have been saved.

Additional information about attributes

Note the following conditions:

- When no data can be collected for a data set, an empty result is returned (no rows of data)
- When a specific attribute cannot be collected, the value 0 or "" is returned unless otherwise specified in a particular attribute (for example, "N/A")
- Any numeric attribute value that is greater than the largest (positive or negative) number that can be represented by that type returns the corresponding maximum or minimum value (for example, the maximum value for a 32-bit number is 2,147,483,647). These values are displayed as text values that are defined by the data set, such as "Value Exceeds Maximum" or "Value Exceeds Minimum".

Numeric attributes have characteristics that are indicated in parentheses after the data type, such as "(32-bit numeric property)". A numeric attribute value can be 32-bit or 64-bit or some other size. The

value type can be gauge, which means it varies, like a speedometer; counter, which counts and always increases; or numeric property, such as disk size.

For a list of the data sets, a list of the attributes in each data set, and descriptions of the attributes in the Microsoft SQL Server agent, see [“Data sets for the monitoring agent” on page 44](#) and [“Attribute descriptions” on page 49](#).

Data sets for the monitoring agent

The Microsoft SQL Server agent contains the following data sets.

- Data set name: MS SQL Additional Availability Database Details
 - Table name: KOQAADBD
 - Historical table name: KOQ_MS_SQL_ADDITIONAL_AVAILABILITY_DATABASE_DETAILS or KOQAADBD
- Data set name: MS SQL Additional Database Detail
 - Table name: KOQADDBD
 - Historical table name: KOQ_MS_SQL_ADDITIONAL_DATABASE_DETAIL or KOQADDBD
- Data set name: MS SQL Analysis Services Data Mining
 - Table name: KOQSASDM
 - Historical table name: KOQ_MS_SQL_ANALYSIS_SERVICES_DATA_MINING or KOQSASDM
- Data set name: MS SQL Analysis Services MDX Processing
 - Table name: KOQSASMX
 - Historical table name: KOQ_MS_SQL_ANALYSIS_SERVICES_MDX_PROCESSING or KOQSASMX
- Data set name: MS SQL Analysis Services Memory Statistics
 - Table name: KOQSASSM
 - Historical table name: KOQ_MS_SQL_ANALYSIS_SERVICES_MEMORY_STATISTICS or KOQSASSM
- Data set name: MS SQL Analysis Services Rows Processing
 - Table name: KOQSASPR
 - Historical table name: KOQ_MS_SQL_ANALYSIS_SERVICES_ROWS_PROCESSING or KOQSASPR
- Data set name: MS SQL Analysis Services Storage Engine Query
 - Table name: KOQSASSQ
 - Historical table name: KOQ_MS_SQL_ANALYSIS_SERVICES_STORAGE_ENGINE_QUERY or KOQSASSQ
- Data set name: MS SQL Audit Details
 - Table name: KOQAUDIT
 - Historical table name: KOQ_MS_SQL_AUDIT_DETAILS or KOQAUDIT
- Data set name: MS SQL Availability Database Details
 - Table name: KOQAVDBD
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_DATABASE_DETAILS or KOQAVDBD
- Data set name: MS SQL Availability Database Details In Cluster
 - Table name: KOQCAVDB
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_DATABASE_DETAILS_IN_CLUSTER or KOQCAVDB
- Data set name: MS SQL Availability Database Statistics
 - Table name: KOQADBST

- Historical table name: KOQ_MS_SQL_AVAILABILITY_DATABASE_STATISTICS or KOQADBST
- Data set name: MS SQL Availability Database Summary
 - Table name: KOQADBSU
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_DATABASE_SUMMARY or KOQADBSU
- Data set name: MS SQL Availability Group Listener Details
 - Table name: KOQAVGLD
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_GROUP_LISTENER_DETAILS or KOQAVGLD
- Data set name: MS SQL Availability Groups Details
 - Table name: KOQGRPDT
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_GROUPS_DETAILS or KOQGRPDT
- Data set name: MS SQL Availability Groups Details In Cluster
 - Table name: KOQCAVGD
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_GROUPS_DETAILS_IN_CLUSTER or KOQCAVGD
- Data set name: MS SQL Availability Groups Summary
 - Table name: KOQGRPSM
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_GROUPS_SUMMARY or KOQGRPSM
- Data set name: MS SQL Availability Replicas Details
 - Table name: KOQAVARD
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_REPLICAS_DETAILS or KOQAVARD
- Data set name: MS SQL Availability Replicas Details In Cluster
 - Table name: KOQCAVRD
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_REPLICAS_DETAILS_IN_CLUSTER or KOQCAVRD
- Data set name: MS SQL Availability Replicas Statistics
 - Table name: KOQAVRST
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_REPLICAS_STATISTICS or KOQAVRST
- Data set name: MS SQL Availability Replicas Status
 - Table name: KOQAVARS
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_REPLICAS_STATUS or KOQAVARS
- Data set name: MS SQL Availability Replicas Status Summary
 - Table name: KOQAVRSU
 - Historical table name: KOQ_MS_SQL_AVAILABILITY_REPLICAS_STATUS_SUMMARY or KOQAVRSU
- Data set name: MS SQL Batch Stats
 - Table name: KOQBTCHS
 - Historical table name: KOQ_MS_SQL_BATCH_STATS or KOQBTCHS
- Data set name: MS SQL Blocked Queries Details
 - Table name: KOQBLKQD
 - Historical table name: KOQ_MS_SQL_BLOCKED_QUERIES_DETAILS or KOQBLKQD
- Data set name: MS SQL Configuration
 - Table name: KOQSCFG
 - Historical table name: KOQ_MS_SQL_CONFIGURATION or KOQSCFG
- Data set name: MS SQL Customized SQL Query Details

- Table name: KOQCUSQL
- Historical table name: KOQ_MS_SQL_CUSTOMIZED_SQL_QUERY_DETAILS or KOQCUSQL
- Data set name: MS SQL Customized SQL Query Result
 - Table name: KOQCSQLR
 - Historical table name: KOQ_MS_SQL_CUSTOMIZED_SQL_QUERY_RESULT or KOQCSQLR
- Data set name: MS SQL Database Detail
 - Table name: KOQDBD
 - Historical table name: KOQ_MS_SQL_DATABASE_DETAIL or KOQDBD
- Data set name: MS SQL Database Mirroring
 - Table name: KOQDBMIR
 - Historical table name: KOQ_MS_SQL_DATABASE_MIRRORING or KOQDBMIR
- Data set name: MS SQL Database Summary
 - Table name: KOQDBS
 - Historical table name: KOQ_MS_SQL_DATABASE_SUMMARY or KOQDBS
- Data set name: MS SQL Device Detail
 - Table name: KOQDEVD
 - Historical table name: KOQ_MS_SQL_DEVICE_DETAIL or KOQDEVD
- Data set name: MS SQL Error Event Detail
 - Table name: KOQERROR
- Data set name: MS SQL Expensive Query Plans
 - Table name: KOQEXPQP
 - Historical table name: KOQ_MS_SQL_EXPENSIVE_QUERY_PLANS or KOQEXPQP
- Data set name: MS SQL Filegroup Detail
 - Table name: KOQFGRPD
 - Historical table name: KOQ_MS_SQL_FILEGROUP_DETAIL or KOQFGRPD
- Data set name: MS SQL FileTable Detail
 - Table name: KOQFTBLD
 - Historical table name: KOQ_MS_SQL_FILETABLE_DETAIL or KOQFTBLD
- Data set name: MS SQL Individual Queries Details
 - Table name: KOQINQRD
 - Historical table name: KOQ_MS_SQL_INDIVIDUAL_QUERIES_DETAILS or KOQINQRD
- Data set name: MS SQL Integration Service Details
 - Table name: KOQSSISD
 - Historical table name: KOQ_MS_SQL_INTEGRATION_SERVICE_DETAILS or KOQSSISD
- Data set name: MS SQL Job Detail
 - Table name: KOQJOBBD
 - Historical table name: KOQ_MS_SQL_JOB_DETAIL or KOQJOBBD
- Data set name: MS SQL Job Summary
 - Table name: KOQJOBS
 - Historical table name: KOQ_MS_SQL_JOB_SUMMARY or KOQJOBS
- Data set name: MS SQL Lock Conflict Detail
 - Table name: KOQLOCK

- Historical table name: KOQ_MS_SQL_LOCK_CONFLICT_DETAIL or KOQLOCK
- Data set name: MS SQL Lock Detail
 - Table name: KOQLOCKS
 - Historical table name: KOQ_MS_SQL_LOCK_DETAIL or KOQLOCKS
- Data set name: MS SQL Lock Resource Type Summary
 - Table name: KOQLRTS
 - Historical table name: KOQ_MS_SQL_LOCK_RESOURCE_TYPE_SUMMARY or KOQLRTS
- Data set name: MS SQL Lock Summary
 - Table name: KOQLOKSU
 - Historical table name: KOQ_MS_SQL_LOCK_SUMMARY or KOQLOKSU
- Data set name: MS SQL Log Shipping DB Details
 - Table name: KOQLSDBD
 - Historical table name: KOQ_MS_SQL_LOG_SHIPPING_DB_DETAILS or KOQLSDBD
- Data set name: MS SQL Log Shipping Errors
 - Table name: KOQLSERR
 - Historical table name: KOQ_MS_SQL_LOG_SHIPPING_ERRORS or KOQLSERR
- Data set name: MS SQL Log Shipping Jobs Detail
 - Table name: KOQLSJBD
 - Historical table name: KOQ_MS_SQL_LOG_SHIPPING_JOBS_DETAIL or KOQLSJBD
- Data set name: MS SQL Memory Manager
 - Table name: KOQMEMGR
 - Historical table name: KOQ_MS_SQL_MEMORY_MANAGER or KOQMEMGR
- Data set name: MS SQL Problem Detail
 - Table name: KOQPROBD
 - Historical table name: KOQ_MS_SQL_PROBLEM_DETAIL or KOQPROBD
- Data set name: MS SQL Problem Summary
 - Table name: KOQPROBS
 - Historical table name: KOQ_MS_SQL_PROBLEM_SUMMARY or KOQPROBS
- Data set name: MS SQL Process Detail
 - Table name: KOQPRCD
 - Historical table name: KOQ_MS_SQL_PROCESS_DETAIL or KOQPRCD
- Data set name: MS SQL Process Summary
 - Table name: KOQPRCS
 - Historical table name: KOQ_MS_SQL_PROCESS_SUMMARY or KOQPRCS
- Data set name: MS SQL Remote Servers
 - Table name: KOQSRVR
 - Historical table name: KOQ_MS_SQL_REMOTE_SERVERS or KOQSRVR
- Data set name: MS SQL Resource Pool Stats
 - Table name: KOQRPOOL
 - Historical table name: KOQ_MS_SQL_RESOURCE_POOL_STATS or KOQRPOOL
- Data set name: MS SQL Running Queries Details
 - Table name: KOQRUNQD

- Historical table name: KOQ_MS_SQL_RUNNING_QUERIES_DETAILS or KOQRUNQD
- Data set name: MS SQL Server Detail
 - Table name: KOQSRVD
 - Historical table name: KOQ_MS_SQL_SERVER_DETAIL or KOQSRVD
- Data set name: MS SQL Server Enterprise View
 - Table name: KOQSRVRE
- Data set name: MS SQL Server Properties
 - Table name: KOQSVRPR
 - Historical table name: KOQ_MS_SQL_SERVER_PROPERTIES or KOQSVRPR
- Data set name: MS SQL Server Summary
 - Table name: KOQSRVS
 - Historical table name: KOQ_MS_SQL_SERVER_SUMMARY or KOQSRVS
- Data set name: MS SQL Server Transactions Summary
 - Table name: KOQSTRNS
 - Historical table name: KOQ_MS_SQL_SERVER_TRANSACTIONS_SUMMARY or KOQSTRNS
- Data set name: MS SQL Service Broker Activation
 - Table name: KOQSBACT
 - Historical table name: KOQ_MS_SQL_SERVICE_BROKER_ACTIVATION or KOQSBACT
- Data set name: MS SQL Service Broker Statistics
 - Table name: KOQSBSTA
 - Historical table name: KOQ_MS_SQL_SERVICE_BROKER_STATISTICS or KOQSBSTA
- Data set name: MS SQL Service Broker Transport
 - Table name: KOQSBTRP
 - Historical table name: KOQ_MS_SQL_SERVICE_BROKER_TRANSPORT or KOQSBTRP
- Data set name: MS SQL Services Detail
 - Table name: KOQSRVCD
 - Historical table name: KOQ_MS_SQL_SERVICES_DETAIL or KOQSRVCD
- Data set name: MS SQL Statistics Detail
 - Table name: KOQSTATD
 - Historical table name: KOQ_MS_SQL_STATISTICS_DETAIL or KOQSTATD
- Data set name: MS SQL Statistics Summary
 - Table name: KOQSTATS
 - Historical table name: KOQ_MS_SQL_STATISTICS_SUMMARY or KOQSTATS
- Data set name: MS SQL Table Detail
 - Table name: KOQTBLD
 - Historical table name: KOQ_MS_SQL_TABLE_DETAIL or KOQTBLD
- Data set name: MS SQL Table Summary
 - Table name: KOQTBL
 - Historical table name: KOQ_MS_SQL_TABLE_SUMMARY or KOQTBL
- Data set name: MS SQL Text
 - Table name: KOQSQL
- Data set name: MS SQL Workload Group Stats

- Table name: KOQWLGS
- Historical table name: KOQ_MS_SQL_WORKLOAD_GROUP_STATS or KOQWLGS

Attribute descriptions

Attributes in each Microsoft SQL Server agent data set collect data that the agent uses for monitoring.

The descriptions of the data sets contain information such as description, type, and names for each attribute in the data set. Some attributes are designated as key attributes, which are identifier attributes for the data set. An attribute in *italic* indicates that it is available for display in the Cloud APM console dashboard pages.

MS SQL Additional Availability Database Details data set

MS SQL Additional Availability Database Details data set provides additional information about the availability databases that are hosted by the SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Database Name

The name of the availability database. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), Database Name (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Filestream Send Rate (KB Per Sec)

The rate (in KB per second) at which filestream data is sent to the secondary replica. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTREAM_SEND_RATE or FSTRSNRT (historical name), Filestream Send Rate (KB Per Sec) (caption), Filestream_Send_Rate (attribute name), and FSTRSNRT (column name).

Group Database ID

The ID of the availability group database. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_DATABASE_ID or GRDBID (historical name), Group Database ID (caption), Group_Database_ID (attribute name), and GRDBID (column name).

Group ID

The ID of the availability group to which the availability database belongs. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), Group ID (caption), Group_ID (attribute name), and GRPID (column name).

Log Bytes Sent Rate (KB Per Sec)

The rate (in KB per second) at which log records are currently being sent to the secondary replica. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected

(-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_SENT_RATE or LGBTSNRT (historical name), Log Bytes Sent Rate (KB Per Sec) (caption), Log_Bytes_Sent_Rate (attribute name), and LGBTSNRT (column name).

Log Send Queue Size (KB)

The amount of log bytes (in KB) in the log files of the primary database that has not been sent to the secondary replica. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SEND_QUEUE_SIZE or LGSNDQSZ (historical name), Log Send Queue Size (KB) (caption), Log_Send_Queue_Size (attribute name), and LGSNDQSZ (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Redo Queue Size (KB)

The amount of log bytes (in KB) that is not yet redone in the log files of the secondary replica. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDO_QUEUE_SIZE or REDOQSZ (historical name), Redo Queue Size (KB) (caption), Redo_Queue_Size (attribute name), and REDOQSZ (column name).

Redo Rate (KB Per Sec)

The rate (in KB per second) at which log records are being redone. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDO_RATE or REDORT (historical name), Redo Rate (KB Per Sec) (caption), Redo_Rate (attribute name), and REDORT (column name).

Role

Indicates whether a replica is currently a primary or secondary replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Resolving (1), Primary (2), Secondary (3), Invalid (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROLE (historical name), Role (caption), Role (attribute name), and ROLE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: **TIMESTAMP** (historical name), **Timestamp** (caption), **Timestamp** (attribute name), and **TIMESTAMP** (column name).

MS SQL Additional Database Detail data set

The MS SQL Additional Database Details data set contains attributes that provide information about the database details for SQL Server 2012 Enterprise Edition, or later. This data set provides information about database details in addition to the MS SQL Database Details data set.

This data set contains the following attributes:

Containment

Indicates whether containment is enabled for the database. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Containment (1), Partial Containment (2), Full Containment (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **CONTAINMENT** or **CONTAINED** (historical name), **Containment** (caption), **Containment** (attribute name), and **CONTAINED** (column name).

Database ID

The database ID. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **DATABASE_ID** or **DBID** (historical name), **Database ID** (caption), **Database_ID** (attribute name), and **DBID** (column name).

Database Name

The name of the database. The value format is an alphanumeric string, for example, KOQ3. Each database name is unique. The SQL Server also assigns each database a unique identification number. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_Applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **DATABASE_NAME_U** or **DBNAME** (historical name), **Database Name** (caption), **Database_Name_U** (attribute name), and **DBNAME** (column name).

Filestream Directory Name

The name of the filestream directory. This attribute is not supported in MSSQL Linux agent. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_Applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **FILETABLE_DIRECTORY_NAME** or **FILETBDIR** (historical name), **Filestream Directory Name** (caption), **FileTable_Directory_Name** (attribute name), and **FILETBDIR** (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: **HOST_NAME** or **HOSTNAME** (historical name), **Host Name** (caption), **Host_Name** (attribute name), and **HOSTNAME** (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Non Transactional Access Level

The level of non-transactional access for the database. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), Not Enabled (1), ReadOnly (2), Full Access (3), In Transition To ReadOnly (4), In Transition To Off (5). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NON_TRANSACTIONAL_ACCESS_LEVEL or NONTRACLVL (historical name), Non Transactional Access Level (caption), Non_Transactional_Access_Level (attribute name), and NONTRACLVL (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transaction Isolation Level of Memory Optimized Tables

Indicates whether the TRANSACTION ISOLATION LEVEL is set to SNAPSHOT. The SNAPSHOT isolation level is used to access the memory-optimized tables. The TRANSACTION ISOLATION LEVEL is set to a lower value, such as READ COMMITTED or READ UNCOMMITTED to access these tables. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Isolation level is not elevated (1), Minimum isolation level is SNAPSHOT (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TRANSACTION_ISOLATION_LEVEL_OF_MEMORY_OPTIMIZED_TABLES or ISMOESNPON (historical name), Transaction Isolation Level of Memory Optimized Tables (caption), Transaction_Isolation_Level_of_Memory_Optimized_Tables (attribute name), and ISMOESNPON (column name).

MS SQL Analysis Services Data Mining data set

The MS SQL Analysis Services Data Mining data set provides detailed information about the data mining of the SQL Server Analysis Services. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Cases Per sec

The rate at which cases are currently processed. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CASES_PER_SEC or CASESPERS (historical name), Cases Per sec (caption), Cases_Per_sec (attribute name), and CASESPERS (column name).

Concurrent Data Mining Queries

The number of data mining queries that are currently being processed. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONCURRENT_DATA_MINING_QUERIES or CONDMQUE (historical name), Concurrent Data Mining Queries (caption), Concurrent_Data_Mining_Questions (attribute name), and CONDMQUE (column name).

Current Models Processing

The number of SQL Server Analysis Services Data Mining models that are currently being processed. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_MODELS_PROCESSING or CRMODPROC (historical name), Current Models Processing (caption), Current_Models_Processing (attribute name), and CRMODPROC (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Predictions Per Sec

The rate at which predictions that are generated in data mining queries. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PREDICTIONS_PER_SEC or PREDCPERS (historical name), Predictions Per Sec (caption), Predictions_Per_Sec (attribute name), and PREDCPERS (column name).

Queries Per Sec

The rate at which data mining queries are generated. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERIES_PER_SEC or QUEPERS (historical name), Queries Per Sec (caption), Queries_Per_Sec (attribute name), and QUEPERS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Rows Per Sec

The number of rows that are currently processed per second for a data mining prediction query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROWS_PER_SEC or ROWPERS (historical name), Rows Per Sec (caption), Rows_Per_Sec (attribute name), and ROWPERS (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value format is an alphanumeric string with a maximum of 30 characters. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Predictions

The number of data mining prediction queries that are currently received by the server. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PREDICTIONS or TOTPRED (historical name), Total Predictions (caption), Total_Predictions (attribute name), and TOTPRED (column name).

Total Queries

The number of data mining queries that are currently received by the server. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_QUERIES or TOTQUE (historical name), Total Queries (caption), Total_Queries (attribute name), and TOTQUE (column name).

Total Rows

The number of rows that are currently returned by data mining queries. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS or TOTROWS (historical name), Total Rows (caption), Total_Rows (attribute name), and TOTROWS (column name).

MS SQL Analysis Services MDX Processing data set

The MS SQL Analysis Services MDX Processing data set provides a summary of data that is processed by Multidimensional Expressions (MDX). This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Current Cached Evaluation Nodes

The approximate number of cached evaluation nodes that are currently built by the MDX execution plans. This attribute is not supported on SQL Server 2005. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_CACHED_EVALUATION_NODES or CURCEVLNDS (historical name), Current Cached Evaluation Nodes (caption), Current_Cached_Evaluation_Nodes (attribute name), and CURCEVLNDS (column name).

Current Evaluation Nodes

The approximate number of evaluation nodes that are currently built by the MDX execution plans. This number includes active evaluation nodes and cached evaluation nodes. This attribute is not supported on SQL Server 2005. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_EVALUATION_NODES or CUREVLNDS (historical name), Current Evaluation Nodes (caption), Current_Evaluation_Nodes (attribute name), and CUREVLNDS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Autoexist

The number of times that the Autoexist operation is currently performed. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_AUTOEXIST or TOTAUTOEX (historical name), Total_Autoexist (caption), Total_Autoexist (attribute name), and TOTAUTOEX (column name).

Total Bulk Mode Evaluation Nodes

The number of bulk-mode evaluation nodes that are currently built by the MDX execution plans. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_BULK_MODE_EVALUATION_NODES or BMEVALNDS (historical name), Total Bulk Mode Evaluation Nodes (caption), Total_Bulk_Mode_Evaluation_Nodes (attribute name), and BMEVALNDS (column name).

Total Cached Bulk Mode Evaluation Nodes

The number of cached bulk-mode evaluation nodes that are currently built by the MDX execution plans. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_CACHED_BULK_MODE_EVALUATION_NODES or TOTCBMEVND (historical name), Total Cached Bulk Mode Evaluation Nodes (caption), Total_Cached_Bulk_Mode_Evaluation_Nodes (attribute name), and TOTCBMEVND (column name).

Total Cached Other Evaluation Nodes

The number of cached evaluation nodes that are currently built by the MDX execution plans, and that are not storage engine evaluation plans or bulk-mode evaluation plans. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CACHED_OTHER_EVALUATION_NODES or TOTCOTHREN (historical name), Total Cached Other Evaluation Nodes (caption), Total_Cached_Other_Evaluation_Nodes (attribute name), and TOTCOTHREN (column name).

Total Cached Storage Engine Evaluation Nodes

The number of cached storage engine evaluation nodes that are currently built by the execution plans. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_CACHED_STORAGE_ENGINE_EVALUATION_NODES or TOTCSENGEN (historical name), Total Cached Storage Engine Evaluation Nodes (caption), Total_Cached_Storage_Engine_Evaluation_Nodes (attribute name), and TOTCSENGEN (column name).

Total Calculation Cache Registered

Total number of calculation cache registered. This attribute is supported only in SQL server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CALCULATION_CACHE_REGISTERED or TOTCALCHRG (historical name), Total Calculation Cache Registered (caption), Total_Calculation_Cache_Registered (attribute name), and TOTCALCHRG (column name).

Total Calculations Covered

The number of evaluation nodes that are currently built by the MDX execution plans. This number includes active evaluation nodes and cached evaluation nodes. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CALCULATIONS_COVERED or TOTCALCVRD (historical name), Total Calculations Covered (caption), Total_Calculations_Covered (attribute name), and TOTCALCVRD (column name).

Total Cell By Cell Evaluation nodes

The number of cell-by-cell evaluation nodes that are currently built by the MDX execution plans. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CELL_BY_CELL_EVALUATION_NODES or TOTCBYCEND (historical name), Total Cell By Cell Evaluation nodes (caption), Total_Cell_By_Cell_Evaluation_nodes (attribute name), and TOTCBYCEND (column name).

Total Cell By Cell Hits In Cache of Evaluation nodes

The number of cell-by-cell hits that currently occurred in the evaluation nodes cache. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_CELL_BY_CELL_HITS_IN_CACHE_OF_EVALUATION_NODES or TCBCHICOEN (historical name), Total Cell By Cell Hits In Cache of Evaluation nodes (caption), Total_Cell_By_Cell_Hits_In_Cache_of_Evaluation_nodes (attribute name), and TCBCHICOEN (column name).

Total Cell By Cell Misses In Cache of Evaluation nodes

The number of cell-by-cell misses that currently occurred in the evaluation nodes cache. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_CELL_BY_CELL_MISSES_IN_CACHE_OF_EVALUATION_NODES or TCBCMICOEN (historical name), Total Cell By Cell Misses In Cache of Evaluation nodes (caption), Total_Cell_By_Cell_Misses_In_Cache_of_Evaluation_nodes (attribute name), and TCBCMICOEN (column name).

Total Cells Calculated

The number of cell properties that are currently calculated. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CELLS_CALCULATED or TOTCELCALC (historical name), Total Cells Calculated (caption), Total_Cells_Calculated (attribute name), and TOTCELCALC (column name).

Total Evaluation Nodes Calculations at Same Granularity

The number of evaluation nodes with calculations that have identical granularity level as that of the evaluation nodes. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_EVALUATION_NODES_CALCULATIONS_AT_SAME_GRANULARITY or TOTENCASG (historical name), Total Evaluation Nodes Calculations at Same Granularity (caption), Total_Evaluation_Nodes_Calculations_at_Same_Granularity (attribute name), and TOTENCASG (column name).

Total Evaluation Nodes Covered Single Cell

The number of evaluation nodes that are currently built by the MDX execution plans, and that covered only one cell. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_EVALUATION_NODES_COVERED_SINGLE_CELL or TOTENCSC (historical name), Total Evaluation Nodes Covered Single Cell (caption), Total_Evaluation_Nodes_Covered_Single_Cell (attribute name), and TOTENCSC (column name).

Total Evictions Of Evaluation Nodes

The number of evaluation nodes that are currently evicted because of collisions. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_EVICTIONS_OF_EVALUATION_NODES or TOTEOEN (historical name), Total Evictions Of Evaluation Nodes (caption), Total_Evictions_Of_Evaluation_Nodes (attribute name), and TOTEOEN (column name).

Total EXISTING Operators

Total number of times EXISTING set operator performed. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_EXISTING_OPERATORS or TOTEXSISOP (historical name), Total EXISTING Operators (caption), Total_EXISTING_Operators (attribute name), and TOTEXSISOP (column name).

Total Flat Cache Inserts

The number of cell values that are currently inserted into the flat calculation cache. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_FLAT_CACHE_INSERTS or TOTFLTCHIN (historical name), Total Flat Cache Inserts (caption), Total_Flat_Cache_Inserts (attribute name), and TOTFLTCHIN (column name).

Total Hash Index Hits In Cache Of Evaluation Nodes

The number of hits in the evaluation nodes cache that are currently satisfied by the hash index. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_HASH_INDEX_HITS_IN_CACHE_OF_EVALUATION_NODES or THIXHICOEN (historical name), Total Hash Index Hits In Cache Of Evaluation Nodes (caption), Total_Hash_Index_Hits_In_Cache_Of_Evaluation_Nodes (attribute name), and THIXHICOEN (column name).

Total Nonempty Algorithms

The number of times a NonEmpty algorithm is currently used. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NONEMPTY_ALGORITHMS or TOTNEMPALG (historical name), Total Nonempty Algorithms (caption), Total_Nonempty_Algorithms (attribute name), and TOTNEMPALG (column name).

Total Nonempty Algorithms For Calculated Members

The number of times that the NonEmpty algorithm currently looped through the calculated members. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_NONEMPTY_ALGORITHMS_FOR_CALCULATED_MEMBERS or TOTNEAFCM (historical name), Total Nonempty Algorithms For Calculated Members (caption), Total_Nonempty_Algorithms_For_Calculated_Members (attribute name), and TOTNEAFCM (column name).

Total Nonempty Unoptimized Algorithms

The number of times that a NonEmpty algorithm that is not optimized is currently used. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NONEMPTY_UNOPTIMIZED_ALGORITHMS or TOTNUNOPAG (historical name), Total Nonempty Unoptimized Algorithms (caption), Total_Nonempty_Unoptimized_Algorithms (attribute name), and TOTNUNOPAG (column name).

Total Recomputes

The number of cells that are currently computed again because of an error. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_RECOMPUTES or TOTRECOMP (historical name), Total Recomputes (caption), Total_Recomputes (attribute name), and TOTRECOMP (column name).

Total Sonar Subcubes

The number of subcubes that are currently generated by the query optimizer. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL SONAR SUBCUBES or TOTSONSCUB (historical name), Total Sonar Subcubes (caption), Total_Sonar_Subcubes (attribute name), and TOTSONSCUB (column name).

Total Storage Engine Evaluation Nodes

Total number of Storage Engine evaluation nodes built by MDX execution plans. This attribute is not supported in SQL server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_STORAGE_ENGINE_EVALUATION_NODES or TOTSEEN (historical name), Total Storage Engine Evaluation Nodes (caption), Total_Storage_Engine_Evaluation_Nodes (attribute name), and TOTSEEN (column name).

Total Subcube Hits In Cache Of Evaluation Nodes

The number of subcube hits that are currently present in the evaluation nodes cache. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_SUBCUBE_HITS_IN_CACHE_OF_EVALUATION_NODES or TOTSCHICEN (historical name), Total Subcube Hits In Cache Of Evaluation Nodes (caption), Total_Subcube_Hits_In_Cache_Of_Evaluation_Nodes (attribute name), and TOTSCHICEN (column name).

Total Subcube Misses In Cache Of Evaluation Nodes

The number of subcube misses that are currently present in the evaluation nodes cache. This attribute is not supported on SQL Server 2005. The type is integer (64-bit counter) with enumerated values.

The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_SUBCUBE_MISSES_IN_CACHE_OF_EVALUATION_NODES or TOTSCMICEN (historical name), Total Subcube Misses In Cache Of Evaluation Nodes (caption), Total_Subcube_Misses_In_Cache_Of_Evaluation_Nodes (attribute name), and TOTSCMICEN (column name).

MS SQL Analysis Services Memory Statistics data set

The MS SQL Analysis Services Memory Statistics data set provides information about memory that is used by the SQL Server Analysis Services. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Aggregate Cache (KB)

The amount of memory that is currently allocated to the file cache. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_CACHE or AVGCACHKB (historical name), Aggregate Cache (KB) (caption), Average_Cache (attribute name), and AVGCACHKB (column name).

Aggregation Map Files

The current number of aggregation map files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGGREGATION_MAP_FILES or AGGRMPFIL (historical name), Aggregation Map Files (caption), Aggregation_Map_Files (attribute name), and AGGRMPFIL (column name).

Cleaner Balance Per Sec

The rate at which balance and shrink operations are performed. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLEANER_BALANCE_PER_SEC or CLRBALPERS (historical name), Cleaner Balance Per Sec (caption), Cleaner_Balance_Per_Sec (attribute name), and CLRBALPERS (column name).

Cleaner Current Price

The current price of memory in dollars per byte and per unit of time. This price is normalized to 1000. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLEANER_CURRENT_PRICE or CLNRCURPRC (historical name), Cleaner Current Price (caption), Cleaner_Current_Price (attribute name), and CLNRCURPRC (column name).

Cleaner Memory (KB)

The sum of shrinkable memory and the amount of non-shrinkable memory that is currently known to the background cleaner. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLEANER_MEMORY or CLNRMEMKB (historical name), Cleaner Memory (KB) (caption), Cleaner_Memory (attribute name), and CLNRMEMKB (column name).

Cleaner Memory nonshrinkable (KB)

The amount of memory that cannot be purged by the background cleaner. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLEANER_MEMORY_NONSHRINKABLE or CMNONSHRKB (historical name), Cleaner Memory nonshrinkable (KB) (caption), Cleaner_Memory_nonshrinkable (attribute name), and CMNONSHRKB (column name).

Cleaner Memory shrinkable (KB)

The amount of memory that can be purged by the background cleaner. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLEANER_MEMORY_SHRINKABLE or CMEMSHRKBS (historical name), Cleaner Memory shrinkable (KB) (caption), Cleaner_Memory_shrinkable (attribute name), and CMEMSHRKBS (column name).

Cleaner Memory Shrunk KB Per Sec

The rate at which the cleaner memory is shrunk. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLEANER_MEMORY_SHRUNK_KB_SEC or CMEMSHRKBS (historical name), Cleaner Memory Shrunk KB Per Sec (caption), Cleaner_Memory_Shru... (attribute name), and CMEMSHRKBS (column name).

Dimension Index(hash) Files

The current number of dimension index (hash) files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DIMENSION_INDEX_FILES or DIMINDFIL (historical name), Dimension Index(hash) Files (caption), Dimension_Index_Files (attribute name), and DIMINDFIL (column name).

Dimension Property Files

The current number of dimension property files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DIMENSION_PROPERTY_FILES or DIMPROPFIL (historical name), Dimension Property Files (caption), Dimension_Property_Files (attribute name), and DIMPROPFIL (column name).

Dimension String Files

The current number of dimension string files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DIMENSION_STRING_FILES or DIMSTRFIL (historical name), Dimension String Files (caption), Dimension_String_Files (attribute name), and DIMSTRFIL (column name).

Fact Aggregation Files

The current number of fact aggregation files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FACT_AGGREGATION_FILES or FCTAGGRFIL (historical name), Fact Aggregation Files (caption), Fact_Aggregation_Files (attribute name), and FCTAGGRFIL (column name).

Fact Data Files

The current number of fact data files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FACT_DATA_FILES or FACTDATFIL (historical name), Fact Data Files (caption), Fact_Data_Files (attribute name), and FACTDATFIL (column name).

Fact String Files

The current number of fact string files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FACT_STRING_FILES or FACTSTRFIL (historical name), Fact String Files (caption), Fact_String_Files (attribute name), and FACTSTRFIL (column name).

Filestore (KB)

The amount of memory that is currently allocated to the filestore. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE or FILSTKB (historical name), Filestore (KB) (caption), Filestore (attribute name), and FILSTKB (column name).

Filestore Clock Pages Examined Per Sec

The rate at which the background cleaner examines the filestore clock pages and considers the pages for eviction. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_CLOCK_PAGES_EXAMINED_PER_SEC or FSCLKPGEPS (historical name), Filestore Clock Pages Examined Per Sec (caption), Filestore_Clock_Pages_Examined_Per_Sec (attribute name), and FSCLKPGEPS (column name).

Filestore Clock Pages HaveRef Per Sec

The rate at which the background cleaner examines the filestore clock pages that have a reference number because they are currently used. The type is real number (64-bit gauge) with two decimal

places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_CLOCK_PAGES_HAVEREF_PER_SEC or FSCPGSHRPS (historical name), Filestore Clock Pages HaveRef Per Sec (caption), Filestore_Clock_Pages_HaveRef_Per_Sec (attribute name), and FSCPGSHRPS (column name).

Filestore Clock Pages Valid Per Sec

The rate at which the background cleaner examines the filestore clock pages that are valid for eviction. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_CLOCK_PAGES_VALID_PER_SEC or FSCPGSVPS (historical name), Filestore Clock Pages Valid Per Sec (caption), Filestore_Clock_Pages_Valid_Per_Sec (attribute name), and FSCPGSVPS (column name).

Filestore KB Reads Per Sec

The rate at which filestore operations are read. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_KB_READS_PER_SEC or FSTRKBRPS (historical name), Filestore KB Reads Per Sec (caption), Filestore_KB_Reads_Per_Sec (attribute name), and FSTRKBRPS (column name).

Filestore KB Writes Per Sec

The rate at which filestore operations are written. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_KB_WRITES_PER_SEC or FSTRKBWPS (historical name), Filestore KB Writes Per Sec (caption), Filestore_KB_Writes_Per_Sec (attribute name), and FSTRKBWPS (column name).

Filestore Memory Pinned (KB)

The amount of filestore memory that is currently pinned. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_MEMORY_PINNED or FSTRMPINKB (historical name), Filestore Memory Pinned (KB) (caption), Filestore_Memory_Pinned (attribute name), and FSTRMPINKB (column name).

Filestore Page Faults Per Sec

The rate at which filestore page faults are generated. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_PAGEFAULTS_PER_SEC or FILSTPGFLS (historical name), Filestore Page Faults Per Sec (caption), Filestore_Page_Faults_Per_Sec (attribute name), and FILSTPGFLS (column name).

Filestore Reads Per Sec

The rate at which filestore pages are read. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_READS_PER_SEC or FLSTREADPS (historical name), Filestore Reads Per Sec (caption), Filestore_Reads_Per_Sec (attribute name), and FLSTREADPS (column name).

Filestore Writes Per Sec

The rate at which filestore pages are written. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTORE_WRITEPS (historical name), Filestore Writes Per Sec (caption), Filestore_Writes_Per_Sec (attribute name), and FILESTORE_WRITEPS (column name).

In Memory Aggregation Map File (KB)

The current size of the in-memory aggregation map file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_AGGRGATION_MAP_FILE or IMAGRMPFL (historical name), In Memory Aggregation Map File (KB) (caption), In_Memory_Aggregation_Map_File (attribute name), and IMAGRMPFL (column name).

In Memory Aggregation Map File KB Per Sec

The rate at which write operations occur in the in-memory aggregation map file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

IN_MEMORY_AGGRGATION_MAP_FILE_KB_PER_SEC or IMAGRMPFLS (historical name), In Memory Aggregation Map File KB Per Sec (caption), In_Memory_Aggregation_Map_File_KB_Per_Sec (attribute name), and IMAGRMPFLS (column name).

In Memory Dimension Index(hash) File (KB)

The current size of the in-memory dimension index (hash) file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_DIMENSION_INDEX_FILE or IMDMINDFL (historical name), In Memory Dimension Index(hash) File (KB) (caption), In_Memory_Dimension_Index_File (attribute name), and IMDMINDFL (column name).

In Memory Dimension Index(hash) File KB Per Sec

The rate at which write operations occur in the in-memory dimension index (hash) file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

IN_MEMORY_DIMENSION_INDEX_FILE_KB_PER_SEC or IMDMINDFLS (historical name), In Memory Dimension Index(hash) File KB Per Sec (caption), In_Memory_Dimension_Index_File_KB_Per_Sec (attribute name), and IMDMINDFLS (column name).

In Memory Fact Aggregation File (KB)

The current size of the in-memory fact aggregation file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_FACT_AGGRGATION_FILE or IMFTAGRFL (historical name), In Memory Fact Aggregation File (KB) (caption), In_Memory_Fact_Aggregation_File (attribute name), and IMFTAGRFL (column name).

In Memory Fact Aggregation File KB Per Sec

The rate at which write operations occur in the in-memory fact aggregation file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

IN_MEMORY_FACT_AGGREGATION_FILE_KB_PER_SEC or IMFTAGRFLS (historical name), In Memory Fact Aggregation File KB Per Sec (caption), In_Memory_Fact_Aggregation_File_KB_Per_Sec (attribute name), and IMFTAGRFLS (column name).

In Memory Fact Data File (KB)

The current size of the in-memory fact data file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_FACT_DATA_FILE or IMFTDATFL (historical name), In Memory Fact Data File (KB) (caption), In_Memory_Fact_Data_File (attribute name), and IMFTDATFL (column name).

In Memory Fact Data File KB Per Sec

The rate at which write operations occur in the in-memory fact data file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_FACT_DATA_FILE_KB_PER_SEC or IMFTDATFLS (historical name), In Memory Fact Data File KB Per Sec (caption), In_Memory_Fact_Data_File_KB_Per_Sec (attribute name), and IMFTDATFLS (column name).

In Memory Fact String File (KB)

The current size of the in-memory fact string file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_FACT_STRING_FILE or IMFTSTRFL (historical name), In Memory Fact String File (KB) (caption), In_Memory_Fact_String_File (attribute name), and IMFTSTRFL (column name).

In Memory Fact String File (KB) Per Sec

The rate at which write operations occur in the in-memory fact string file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_FACT_STRING_FILE_PER_SEC or IMFTSTRFLS (historical name), In Memory Fact String File (KB) Per Sec (caption), In_Memory_Fact_String_File_Per_Sec (attribute name), and IMFTSTRFLS (column name).

In Memory Map File (KB)

The current size of the in-memory map file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_MAP_FILE or IMMAPFL (historical name), In Memory Map File (KB) (caption), In_Memory_Map_File (attribute name), and IMMAPFL (column name).

In Memory Map File (KB) Per Sec

The rate at which write operations occur in the in-memory map file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_MAP_FILE_PER_SEC or IMMAPFLS (historical name), In Memory Map File (KB) Per Sec (caption), In_Memory_Map_File_Per_Sec (attribute name), and IMMAPFLS (column name).

In Memory Other File (KB)

The current size of the in-memory other file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_OTHER_FILE or IMOTHRFL (historical name), In Memory Other File (KB) (caption), In_Memory_Other_File (attribute name), and IMOTHRFL (column name).

In Memory Other File (KB) Per Sec

The rate at which write operations occur in the in-memory other file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_OTHER_FILE_PER_SEC or IMOTHRFLS (historical name), In Memory Other File (KB) Per Sec (caption), In_Memory_Other_File_Per_Sec (attribute name), and IMOTHRFLS (column name).

In Memory Property File (KB)

The current size of the in-memory property file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_PROPERTY_FILE or IMPROPFIL (historical name), In Memory Property File (KB) (caption), In_Memory_Property_File (attribute name), and IMPROPFIL (column name).

In Memory Property File KB Per Sec

The rate at which write operations occur in the in-memory property file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_PROPERTY_FILE_KB_PER_SEC or IMPROPFILS (historical name), In Memory Property File KB Per Sec (caption), In_Memory_Property_File_KB_Per_Sec (attribute name), and IMPROPFILS (column name).

In Memory String File (KB)

The current size of the in-memory string file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_STRING_FILE or IMSTRFIL (historical name), In Memory String File (KB) (caption), In_Memory_String_File (attribute name), and IMSTRFIL (column name).

In Memory String File KB Per Sec

The rate at which write operations occur in the in-memory string file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IN_MEMORY_STRING_FILE_KB_PER_SEC or IMSTRFILS (historical name), In Memory String File KB Per Sec (caption), In_Memory_String_File_KB_Per_Sec (attribute name), and IMSTRFILS (column name).

Map Files

The current number of map files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAP_FILES or MAPFL (historical name), Map Files (caption), Map_Files (attribute name), and MAPFL (column name).

Memory Limit High (KB)

The highest memory limit that is specified in the configuration file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_LIMIT_HIGH or MEMLMHIKB (historical name), Memory Limit High (KB) (caption), Memory_Limit_High (attribute name), and MEMLMHIKB (column name).

Memory Limit Low (KB)

The lowest memory limit that is specified in the configuration file. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_LIMIT_LOW or MEMLMLOKB (historical name), Memory Limit Low (KB) (caption), Memory_Limit_Low (attribute name), and MEMLMLOKB (column name).

Memory Usage (KB)

The amount of memory that is used by the server to calculate the cleaner memory price. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_USAGE or MEMUSGKB (historical name), Memory Usage (KB) (caption), Memory_Usage (attribute name), and MEMUSGKB (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Other Files

The current number of other files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OTHER_FILES or OTHRFILES (historical name), Other Files (caption), Other_Files (attribute name), and OTHRFILES (column name).

Page Pool 1 Alloc (KB)

The amount of memory that is borrowed from the 64 KB page pool. This memory is distributed to other parts of the server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_POOL_1_ALLOC or PGPL1ALCKB (historical name), Page Pool 1 Alloc (KB) (caption), Page_Pool_1_Alloc (attribute name), and PGPL1ALCKB (column name).

Page Pool 1 Lookaside (KB)

The amount of memory that is currently available in the 1 KB lookaside list. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_POOL_1_LOOKASIDE or PGPL1LKASD (historical name), Page Pool 1 Lookaside (KB) (caption), Page_Pool_1_Lookaside (attribute name), and PGPL1LKASD (column name).

Page Pool 64 Alloc (KB)

The amount of memory that is borrowed from the system. This memory is distributed to other parts of the server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_POOL_64_ALLOC or PGPL64ALC (historical name), Page Pool 64 Alloc (KB) (caption), Page_Pool_64_Alloc (attribute name), and PGPL64ALC (column name).

Page Pool 64 Lookaside (KB)

The amount of memory that is currently available in the 64 KB lookaside list. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_POOL_64_LOOKASIDE or PGPL64LASD (historical name), Page Pool 64 Lookaside (KB) (caption), Page_Pool_64_Lookaside (attribute name), and PGPL64LASD (column name).

Page Pool 8 Alloc (KB)

The amount of memory that is currently borrowed from the 64 KB page pool. The memory is distributed to other parts of the server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_POOL_8_ALLOC or PGPL8ALC (historical name), Page Pool 8 Alloc (KB) (caption), Page_Pool_8_Alloc (attribute name), and PGPL8ALC (column name).

Page Pool 8 Lookaside (KB)

The amount of memory that is currently available in the 8 KB lookaside list. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_POOL_8_LOOKASIDE or PGPL8LASD (historical name), Page Pool 8 Lookaside (KB) (caption), Page_Pool_8_Lookaside (attribute name), and PGPL8LASD (column name).

Potential In-memory Aggregation Map File (KB)

The potential size of the in-memory aggregation map files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_AGGREGATION_MAP_FILE or PIMAGRMFKB (historical name), Potential In-memory Aggregation Map File (KB) (caption), Potential_In-memory_Aggregation_Map_File (attribute name), and PIMAGRMFKB (column name).

Potential In-memory Dimension Index (Hash) File (KB)

The potential size of the in-memory dimension index (hash) files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_DIMENSION_INDEX_FILE or PIMDIHFKB (historical name), Potential In-memory Dimension Index (Hash) File (KB) (caption), Potential_In-memory_Dimension_Index_File (attribute name), and PIMDIHFKB (column name).

Potential In-memory Dimension Property File (KB)

The potential size of the in-memory dimension property files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_DIMENSION_PROPERTY_FILE or PIMDPRFKB (historical name), Potential In-memory Dimension Property File (KB) (caption), Potential_In-memory_Dimension_Property_File (attribute name), and PIMDPRFKB (column name).

Potential In-memory Dimension String File (KB)

The potential size of the in-memory dimension string files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_DIMENSION_STRING_FILE or PIMDSTRFKB (historical name), Potential In-memory Dimension String File (KB) (caption), Potential_In-memory_Dimension_String_File (attribute name), and PIMDSTRFKB (column name).

Potential In-memory Fact Aggregation File (KB)

The potential size of the in-memory fact aggregation files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_FACT_AGGREGATION_FILE or PIMFAGRFKB (historical name), Potential In-memory Fact Aggregation File (KB) (caption), Potential_In-memory_Fact_Aggregation_File (attribute name), and PIMFAGRFKB (column name).

Potential In-memory Fact Data File (KB)

The potential size of the in-memory fact data files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_FACT_DATA_FILE or PIMFDATFKB (historical name), Potential In-memory Fact Data File (KB) (caption), Potential_In-memory_Fact_Data_File (attribute name), and PIMFDATFKB (column name).

Potential In-memory Fact String File (KB)

The potential size of the in-memory fact string files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_FACT_STRING_FILE or PIMFSTRFKB (historical name), Potential In-memory Fact String File (KB) (caption), Potential_In-memory_Fact_String_File (attribute name), and PIMFSTRFKB (column name).

Potential In-memory Map File (KB)

The potential size of the in-memory map files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_MAP_FILE or PIMFMAPFKB (historical name), Potential In-memory Map File (KB) (caption), Potential_In-memory_Map_File (attribute name), and PIMFMAPFKB (column name).

Potential In-memory Other File (KB)

The potential size of the in-memory other files. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: POTENTIAL_IN-MEMORY_OTHER_FILE or PIMFOTRFKB (historical name), Potential In-memory Other File (KB) (caption), Potential_In-memory_Other_File (attribute name), and PIMFOTRFKB (column name).

Quota (KB)

The current amount of memory quota. A memory quota is also referred to as a memory grant or memory reservation. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUOTA or QUOTAKB (historical name), Quota (KB) (caption), Quota (attribute name), and QUOTAKB (column name).

Quota Blocked

The number of requests for memory quota that are currently blocked until some memory quotas are freed. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUOTA_BLOCKED or QUOTABLKD (historical name), Quota Blocked (caption), Quota_Blocked (attribute name), and QUOTABLKD (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Filestore IO Errors

The current number of filestore I/O errors. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_FILESTORE_IO_ERRORS or TOTFLSTER (historical name), Total Filestore IO Errors (caption), Total_Filestore_IO_Errors (attribute name), and TOTFLSTER (column name).

Total Filestore IO Errors Per Sec

The rate at which filestore I/O errors occur. The type is real number (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_FILESTORE_IO_ERRORS_PER_SEC or TOTFLSTERS (historical name), Total Filestore IO Errors Per Sec (caption), Total_Filestore_IO_Errors_Per_Sec (attribute name), and TOTFLSTERS (column name).

MS SQL Analysis Services Rows Processing data set

The MS SQL Analysis Services Rows Processing data set provides a summary of the rows that are processed by the SQL Server Analysis Services. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Rows Converted Per Sec

The rate at which rows are converted during processing per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROWS_CONVERTED_PER_SEC or ROWCONVPS (historical name), Rows Converted Per Sec (caption), Rows_Converted_Per_Sec (attribute name), and ROWCONVPS (column name).

Rows Read Per Sec

The rate at which rows are read from all the relational databases per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROWS_READ_PER_SEC or ROWRDPS (historical name), Rows Read Per Sec (caption), Rows_Read_Per_Sec (attribute name), and ROWRDPS (column name).

Rows Written Per Sec

The rate at which rows are written to the database during processing. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROWS_WRITTEN_PER_SEC or ROWWRPS (historical name), Rows Written Per Sec (caption), Rows_Written_Per_Sec (attribute name), and ROWWRPS (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Rows Converted

The number of rows that are currently converted during processing. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS_CONVERTED or TOTROWCONV (historical name), Total Rows Converted (caption), Total_Rows_Converted (attribute name), and TOTROWCONV (column name).

Total Rows Read

The number of rows that are currently read from all the relational databases. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS_READ or TOTROWRD (historical name), Total Rows Read (caption), Total_Rows_Read (attribute name), and TOTROWRD (column name).

Total Rows Written

The number of rows that are currently written to the database during processing. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS_WRITTEN or TOTROWWR (historical name), Total Rows Written (caption), Total_Rows_Written (attribute name), and TOTROWWR (column name).

MS SQL Analysis Services Storage Engine Query data set

The MS SQL Analysis Services Storage Engine Query data set provides the summary of statistics that are related to the storage engine query processing feature of the SQL Server Analysis Services. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Aggregation Hits Per Sec

The rate of aggregation hits. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGGREGATION_HITS_PER_SEC or AGRHITSPS (historical name), Aggregation Hits Per Sec (caption), Aggregation_Hits_Per_Sec (attribute name), and AGRHITSPS (column name).

Aggregation Lookups Per Sec

The rate of aggregation lookups. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGGREGATION_LOOKUPS_PER_SEC or AGRLOOKUPS (historical name), Aggregation Lookups Per Sec (caption), Aggregation_Lookups_Per_Sec (attribute name), and AGRLOOKUPS (column name).

Average Time Per Query

The average time in milliseconds that is taken to respond to each query that was answered since the last counter measurement. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_TIME_PER_QUERY or AVGTPQUERY (historical name), Average Time Per Query (caption), Average_Time_Per_Query (attribute name), and AVGTPQUERY (column name).

Bytes Sent Per Sec

The number of bytes that are currently sent by the server to the clients per second in response to queries. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BYTES_SENT_PER_SEC or BYTESENTPS (historical name), Bytes Sent Per Sec (caption), Bytes_Sent_Per_Sec (attribute name), and BYTESENTPS (column name).

Calculation Cache Hits Per Sec

The number of calculation hits per second. The calculation cache includes global, session, and query scope calculation caches. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CALCULATION_CACHE_HITS_PER_SEC or CLCACHEHPS (historical name), Calculation Cache Hits Per Sec (caption), Calculation_Cache_Hits_Per_Sec (attribute name), and CLCACHEHPS (column name).

Calculation Cache Lookups Per Sec

The number of calculation cache lookups per second. The calculation cache includes global, session, and query scope calculation caches. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CALCULATION_CACHE_LOOKUPS_PER_SEC or CCACHELKPS (historical name), Calculation Cache Lookups Per Sec (caption), Calculation_Cache_Lookups_Per_Sec (attribute name), and CCACHELKPS (column name).

Current Dimension Queries

The number of dimension queries that are currently being processed. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_DIMENSION_QUERIES or CURDMQUE (historical name), Current Dimension Queries (caption), Current_Dimension_Queries (attribute name), and CURDMQUE (column name).

Current Measure Group Queries

The number of measure group queries that are currently being worked on. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_MEASURE_GROUP_QUERIES or CURMGRQUE (historical name), Current Measure Group Queries (caption), Current_Measure_Group_Questions (attribute name), and CURMGRQUE (column name).

Current Pyramid Operations

The number of pyramid operations that are currently in progress. This attribute is supported only on SQL Server 2005. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_PYRAMID_OPERATIONS or CURPYRMOP (historical name), Current Pyramid Operations (caption), Current_Pyramid_Operations (attribute name), and CURPYRMOP (column name).

Data Bytes Per Sec

The number of bytes that are currently read per second from the data file. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_BYTES_PER_SEC or DATABYTSPS (historical name), Data Bytes Per Sec (caption), Data_Bits_Per_Sec (attribute name), and DATABYTSPS (column name).

Data Reads Per Sec

The number of logical read operations that are currently using the data file per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_READS_PER_SEC or DATAREDSPS (historical name), Data Reads Per Sec (caption), Data_Reads_Per_Sec (attribute name), and DATAREDSPS (column name).

Dimension Cache Hits Per Sec

The number of dimension cache hits per second. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DIMENSION_CACHE_HITS_PER_SEC or DIMCHITSPS (historical name), Dimension Cache Hits Per Sec (caption), Dimension_Cache_Hits_Per_Sec (attribute name), and DIMCHITSPS (column name).

Dimension Cache Lookups Per Sec

The rate of dimension cache lookups. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DIMENSION_CACHE_LOOKUPS_PER_SEC or DIMCLKPSPS (historical name), Dimension Cache Lookups Per Sec (caption), Dimension_Cache_Lookups_Per_Sec (attribute name), and DIMCLKPSPS (column name).

Dimension Queries Per Sec

The rate at which dimension queries are processed. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DIMENSION_QUERIES_PER_SEC or DMQUERYPS (historical name), Dimension Queries Per Sec (caption), Dimension_Queries_Per_Sec (attribute name), and DMQUERYPS (column name).

Filter Rows Excluded Per Sec

The number of rows that are excluded per second in the last query that was processed. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILTER_ROWS_EXCLUDED_PER_SEC or FILTROWEXS (historical name), Filter Rows Excluded Per Sec (caption), Filter_Rows_Excluded_Per_Sec (attribute name), and FILTROWEXS (column name).

Filter Rows Included Per Sec

The number of rows that are included per second in the last query that was processed. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILTER_ROWS_INCLUDED_PER_SEC or FILTROWINS (historical name), Filter Rows Included Per Sec (caption), Filter_Rows_Included_Per_Sec (attribute name), and FILTROWINS (column name).

Filtered Rows Per Sec

The rate at which rows are currently filtered. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILTERED_ROWS_PER_SEC or FILTROWS (historical name), Filtered Rows Per Sec (caption), Filtered_Rows_Per_Sec (attribute name), and FILTROWS (column name).

Flat Cache Hits Per Sec

The rate of flat cache hits. The flat cache includes global, session, and query scope flat caches. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FLAT_CACHE_HITS_PER_SEC or FCHITSPS (historical name), Flat Cache Hits Per Sec (caption), Flat_Cache_Hits_Per_Sec (attribute name), and FCHITSPS (column name).

Flat Cache Lookups Per Sec

The rate of flat cache lookups. The flat cache includes global, session, and query scope flat caches. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FLAT_CACHE_LOOKUPS_PER_SEC or FCLOKUPSPS (historical name), Flat Cache Lookups Per Sec (caption), Flat_Cache_Lookups_Per_Sec (attribute name), and FCLOKUPSPS (column name).

Index Bytes Per Sec

The number of bytes that are currently read from the index files per second. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INDEX_BYTES_PER_SEC or INDXBYTS (historical name), Index Bytes Per Sec (caption), Index_Bytes_Per_Sec (attribute name), and INDXBYTS (column name).

Index Reads Per Sec

The number of logical read operations that are currently using the index files per second. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INDEX_READS_PER_SEC or INDXRDS (historical name), Index Reads Per Sec (caption), Index_Reads_Per_Sec (attribute name), and INDXRDS (column name).

Map Bytes Per Sec

The number of bytes that are currently read from the map file per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAP_BYTES_PER_SEC or MAPBYTESPS (historical name), Map Bytes Per Sec (caption), Map_Bytes_Per_Sec (attribute name), and MAPBYTESPS (column name).

Map Reads Per Sec

The number of logical read operations that are currently using the map file per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAP_READS_PER_SEC or MAPREADSPS (historical name), Map Reads Per Sec (caption), Map_Reads_Per_Sec (attribute name), and MAPREADSPS (column name).

Measure Group Cache Hits Per Sec

The rate of group cache hits. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEASURE_GROUP_CACHE_HITS_PER_SEC or MSGRPCHT (historical name), Measure Group Cache Hits Per Sec (caption), Measure_Group_Cache_Hits_Per_Sec (attribute name), and MSGRPCHT (column name).

Measure Group Lookups Per Sec

The rate of measure group cache lookups. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEASURE_GROUP_LOOKUPS_PER_SEC or MSGRPCHLP (historical name), Measure Group Lookups Per Sec (caption), Measure_Group_Lookups_Per_Sec (attribute name), and MSGRPCHLP (column name).

Measure Group Queries Per Sec

The rate of measure group queries that are sent to the server. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEASURE_GROUP_QUERIES_PER_SEC or MGRPQUERYPS (historical name), Measure Group Queries Per Sec (caption), Measure_Group_Queries_Per_Sec (attribute name), and MGRPQUERYPS (column name).

Network Round Trips Per Sec

The rate of network round trips. The network round trips include all communication between the client and the server. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NETWORK_ROUND_TRIPS_PER_SEC or NETRTRIPPS (historical name), Network Round Trips Per Sec (caption), Network_Round_Trips_Per_Sec (attribute name), and NETRTRIPPS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Persisted Cache Hits Per Sec

The rate of persisted cache hits. Persisted caches are created by the MDX script cache statement. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERSISTED_CACHE_HITS_PER_SEC or PCHITSPS (historical name), Persisted Cache Hits Per Sec (caption), Persisted_Cache_Hits_Per_Sec (attribute name), and PCHITSPS (column name).

Persisted Cache Lookups Per Sec

The rate of persisted cache lookups. Persisted caches are created by the MDX script cache statement. This attribute is not supported on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected

(-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERSISTED_CACHE_LOOKUPS_PER_SEC or PCKLKPSPS (historical name), Persisted Cache Lookups Per Sec (caption), Persisted_Cache_Lookups_Per_Sec (attribute name), and PCKLKPSPS (column name).

Pyramid Operations Per Sec

The rate at which pyramid operations are currently started. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PYRAMID_OPERATIONS_PER_SEC or PYRMDOPS (historical name), Pyramid Operations Per Sec (caption), Pyramid_Operations_Per_Sec (attribute name), and PYRMDOPS (column name).

Queries Answered Per Sec

The rate at which queries are currently answered. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERIES_ANSWERED_PER_SEC or QUERYANSPS (historical name), Queries Answered Per Sec (caption), Queries_Answered_Per_Sec (attribute name), and QUERYANSPS (column name).

Queries From Cache Direct Per Sec

The rate at which queries are answered directly from the cache. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERIES_FROM_CACHE_DIRECT_PER_SEC or QFCDIRCTPS (historical name), Queries From Cache Direct Per Sec (caption), Queries_From_Cache_Direct_Per_Sec (attribute name), and QFCDIRCTPS (column name).

Queries From Cache Filtered Per Sec

The rate at which queries are answered by filtering the existing cache entry. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERIES_FROM_CACHE_FILTERED_PER_SEC or QFCFILTRPS (historical name), Queries From Cache Filtered Per Sec (caption), Queries_From_Cache_Filtered_Per_Sec (attribute name), and QFCFILTRPS (column name).

Queries From Files Per Sec

The rate at which queries are answered from files. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERIES_FROM_FILES_PER_SEC or QFCFILESPS (historical name), Queries From Files Per Sec (caption), Queries_From_Files_Per_Sec (attribute name), and QFCFILESPS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Rows Read Per Sec

The number of rows that are currently read by the server per second. This attribute is supported only on SQL Server 2005. The type is real number (64-bit gauge) with two decimal places of precision with

enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROWS_READ_PER_SEC or ROWSREADPS (historical name), Rows Read Per Sec (caption), Rows_Read_Per_Sec (attribute name), and ROWSREADPS (column name).

Rows Sent Per Sec

The rate at which rows are sent by the server to the clients. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROWS_SENT_PER_SEC or ROWSSENTPS (historical name), Rows Sent Per Sec (caption), Rows_Sent_Per_Sec (attribute name), and ROWSSENTPS (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Bytes Sent

The number of bytes that are currently sent by the server to the clients in response to queries. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_BYTES_SENT or TOTBYTSENT (historical name), Total Bytes Sent (caption), Total_Bytes_Sent (attribute name), and TOTBYTSENT (column name).

Total Dimension Queries

The current number of dimension queries. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DIMENSION_QUERIES or TOTDMQUERY (historical name), Total Dimension Queries (caption), Total_Dimension_Queries (attribute name), and TOTDMQUERY (column name).

Total Measure Group Queries

The current number of measure group queries. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_MEASURE_GROUP_QUERIES or TOTMEGRQUE (historical name), Total Measure Group Queries (caption), Total_Measure_Group_Queries (attribute name), and TOTMEGRQUE (column name).

Total Network Round Trips

The current number of network round trips. The network round trips include all communication between the client and the server. The type is integer (64-bit gauge) with enumerated values. The

following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NETWORK_ROUND_TRIPS or TOTNETRTRP (historical name), Total Network Round Trips (caption), Total_Network_Round_Trips (attribute name), and TOTNETRTRP (column name).

Total Pyramid Operations

The current number of pyramid operations. This attribute is supported only on SQL Server 2005. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PYRAMID_OPERATIONS or TOTPYRMOP (historical name), Total Pyramid Operations (caption), Total_Pyramid_Operations (attribute name), and TOTPYRMOP (column name).

Total Queries Answered

The current number of queries that are answered. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_QUERIES_ANSWERED or TOTQUEANS (historical name), Total Queries Answered (caption), Total_Qualities_Answered (attribute name), and TOTQUEANS (column name).

Total Queries From Cache Direct

The number of queries that are currently derived from the cache for every partition. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_QUERIES_FROM_CACHE_DIRECT or TOTQCCHEDRT (historical name), Total Queries From Cache Direct (caption), Total_Qualities_From_Cache_Direct (attribute name), and TOTQCCHEDRT (column name).

Total Queries From Cache Filtered

The number of queries that are currently answered by filtering the existing cache entries. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_QUERIES_FROM_CACHE_FILTERED or TOTQCHEFTD (historical name), Total Queries From Cache Filtered (caption), Total_Qualities_From_Cache_Filtered (attribute name), and TOTQCHEFTD (column name).

Total Queries From File

The number of queries that are currently answered from files. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_QUERIES_FROM_FILE or TOTQFRMFIL (historical name), Total Queries From File (caption), Total_Qualities_From_File (attribute name), and TOTQFRMFIL (column name).

Total Rows Sent

The number of rows that are currently sent by the server to the clients. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS_SENT or TOTROWSNT (historical name), Total Rows Sent (caption), Total_Rows_Sent (attribute name), and TOTROWSNT (column name).

MS SQL Audit Details data set

The MS SQL Audit Details data set provides details about the audit events that are written to the application log, the security logs, and the audit files. The data collection for this data set is limited to 7000 rows. This data set is supported for SQL Server 2008, or later.

This data set contains the following attributes:

Action ID

The ID of the audit action. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTION_ID or ACTIONID (historical name), Action ID (caption), Action_ID (attribute name), and ACTIONID (column name).

Audit Action Name

The name of the audit action. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUDIT_ACTION_NAME or ACTIONNM (historical name), Audit Action Name (caption), Audit_Action_Name (attribute name), and ACTIONNM (column name).

Audit Name

The name of the audit. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUDIT_NAME or AUDITNAME (historical name), Audit Name (caption), Audit_Name (attribute name), and AUDITNAME (column name).

Audit Specification Name

The name of the audit specification. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUDIT_SPECIFICATION_NAME or SPECNAME (historical name), Audit Specification Name (caption), Audit_Specification_Name (attribute name), and SPECNAME (column name).

Audit Specification Type

The type of the audit specification. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUDIT_SPECIFICATION_TYPE or SPECTYPE (historical name), Audit Specification Type (caption), Audit_Specification_Type (attribute name), and SPECTYPE (column name).

Class Type

The type of the entity that is being audited. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLASS_TYPE or CLASSTYPE (historical name), Class Type (caption), Class_Type (attribute name), and CLASSTYPE (column name).

Database Name

The database context in which the action occurred. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or ADBNAME (historical name), Database Name (caption), Database_Name (attribute name), and ADBNAME (column name).

Database Principal Name

The name of the owner of the audit that is registered on the server. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_PRINCIPAL_NAME or DBPNAME (historical name), Database Principal Name (caption), Database_Principal_Name (attribute name), and DBPNAME (column name).

Event Time

The local date and time when the audit action is triggered. This attribute is a key attribute. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EVENT_TIME or EVENTTIME (historical name), Event Time (caption), Event_Time (attribute name), and EVENTTIME (column name).

File Name

The name and location of the audit log file. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILE_NAME or FILENM (historical name), File Name (caption), File_Name (attribute name), and FILENM (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Object Name

The name of the entity on which the audit action occurred. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OBJECT_NAME or OBJECTNM (historical name), Object Name (caption), Object_Name (attribute name), and OBJECTNM (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Schema Name

The schema context in which the audit action occurred. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SCHEMA_NAME or SCHEMANM (historical name), Schema Name (caption), Schema_Name (attribute name), and SCHEMANM (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Server Principal Name

The current login to the SQL Server where the audit action occurred. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_PRINCIPAL_NAME or SPNAME (historical name), Server Principal Name (caption), Server_Principal_Name (attribute name), and SPNAME (column name).

Statement

The Transact-SQL statement that is specified by the user. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STATEMENT (historical name), Statement (caption), Statement (attribute name), and STATEMENT (column name).

Succeeded

Indicates whether the audit action succeeded. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), True (1), False (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SUCCEEDED (historical name), Succeeded (caption), Succeeded (attribute name), and SUCCEEDED (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

User Defined Event ID

The event ID that is specified by the user. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: USER_DEFINED_EVENT_ID or UEVENTID (historical name), User Defined Event ID (caption), User_Defined_Event_ID (attribute name), and UEVENTID (column name).

User Defined Information

Additional information that the user wants to record in the audit log file. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: USER_DEFINED_INFORMATION or UDINFO (historical name), User Defined Information (caption), User_Defined_Information (attribute name), and UDINFO (column name).

MS SQL Availability Database Details data set

The MS SQL Availability Database Details data set provides information about the availability databases that are hosted by the SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 5 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Database Name

The name of the availability database that is hosted by the SQL Server instance. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), *Database Name* (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Database State

Indicates the current state of the availability database. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Online (1), Restoring (2), Recovering (3), Recovery Pending (4), Suspect (5), Emergency (6), Offline (7). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_STATE or DBSTAT (historical name), *Database State* (caption), Database_State (attribute name), and DBSTAT (column name).

Group Database ID

The ID of the availability group database that is hosted by the SQL Server instance. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_DATABASE_ID or GRDBID (historical name), *Group Database ID* (caption), Group_Database_ID (attribute name), and GRDBID (column name).

Group ID

The ID of the availability group to which the availability database belongs. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), *Group ID* (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The name of the availability group to which the availability database belongs. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME_U or GRPNAME (historical name), *Group Name* (caption), Group_Name_U (attribute name), and GRPNAME (column name).

Local

Indicates whether the availability replica is hosted by the local SQL Server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), No (1), Yes (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCAL (historical name), *Local* (caption), Local (attribute name), and LOCAL (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Replica ID

The ID of the availability replica. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_ID or REPID (historical name), *Replica ID* (caption), Replica_ID (attribute name), and REPID (column name).

Role

Indicates whether the availability database is a primary or a secondary database. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Resolving (1), Primary (2), Secondary (3), Invalid (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROLE (historical name), *Role* (caption), Role (attribute name), and ROLE (column name).

Suspend Reason

The reason for a database to enter the suspended state. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Suspend From User (1), Suspend From Partner (2), Suspend From Redo (3), Suspend From Apply (4), Suspend From Capture (5), Suspend From Restart (6), Suspend From Undo (7), Suspend From Revalidation (8). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SUSPEND_REASON or SUSPRESN (historical name), *Suspend Reason* (caption), Suspend_Reason (attribute name), and SUSPRESN (column name).

Synchronization Health

Indicates the health of the availability database replica. The replica is healthy if the synchronized and the asynchronized replicas are in the Synchronized state. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Not Healthy (1), Partially Healthy (2), Healthy (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SYNCHRONIZATION_HEALTH or SYNCHLTH (historical name), *Synchronization Health* (caption), Synchronization_Health (attribute name), and SYNCHLTH (column name).

Synchronization State

The current synchronization state of the availability database replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Not Synchronizing (1), Synchronizing (2), Synchronized (3), Reverting (4), Initializing (5). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SYNCHRONIZATION_STATE or SYNSTATE (historical name), *Synchronization State* (caption), Synchronization_State (attribute name), and SYNSTATE (column name).

Synchronized Commit

Indicates whether the transaction commit is synchronized with the database replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Not Synchronized (1), Synchronized (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SYNCHRONIZED_COMMIT or SYNCCOM (historical name), *Synchronized Commit* (caption), Synchronized_Commit (attribute name), and SYNCCOM (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Database ID

The ID of the availability database that is hosted by the SQL Server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database ID (caption), Database_ID (attribute name), and DBID (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Availability Database Details In Cluster data set

The MS SQL Availability Database Details In Cluster data set provides detailed information about the health of the availability databases in each availability group on the Windows Server Failover Clustering (WFSC) cluster even if the availability database is not hosted by the local SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Database Name

The name of the availability database. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), Database Name (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Failover Ready

Indicates whether the secondary database is synchronized with the primary database. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Ready (1), Ready (2), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAILOVER_READY or FAILOVRED (historical name), Failover Ready (caption), Failover_Ready (attribute name), and FAILOVRED (column name).

Group Database ID

The availability group database ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_DATABASE_ID or GRDBID (historical name), Group Database ID (caption), Group_Database_ID (attribute name), and GRDBID (column name).

Group ID

The availability group ID. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), Group ID (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The availability group name. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME_U or GRPNAME (historical name), Group Name (caption), Group_Name_U (attribute name), and GRPNAME (column name).

Join State

Indicates whether the availability database on the availability replica is joined to the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Joined (1), Joined (2), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOIN_STATE or JOINSTAT (historical name), Join State (caption), Join_State (attribute name), and JOINSTAT (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Replica ID

The replica ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_ID or REPID (historical name), Replica ID (caption), Replica_ID (attribute name), and REPID (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Secondary Suspension Pending

Indicates whether the database is pending suspension after a forced failover. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: No (1), Yes (2), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SECONDARY_SUSPENSION_PENDING or SECSUPPEN (historical name), Secondary Suspension Pending (caption), Secondary_Suspension_Pending (attribute name), and SECSUPPEN (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Availability Database Statistics data set

MS SQL Availability Database Statistics data set provides information about the availability database statistics. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Database Name

The name of the availability database. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), Database Name (caption), Database_Name_U (attribute name), and UDBNAME (column name).

File Bytes Received Per Sec

The amount of filestream data that is currently being received per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILE_BYTES_RECEIVED_PER_SEC or FBTRECPs (historical name), File Bytes Received Per Sec (caption), File_Bytes_Received_Per_Sec (attribute name), and FBTRECPs (column name).

Log Bytes Received Per Sec

The amount of log bytes that is currently being received per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_RECEIVED_PER_SEC or LGBTRCPs (historical name), Log Bytes Received Per Sec (caption), Log_Bytes_Received_Per_Sec (attribute name), and LGBTRCPs (column name).

Log Send Queue (KB)

The amount of log bytes (in KB) that has been sent by the primary replica, but has not been received by the secondary replica. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SEND_QUEUE or LGSENQUE (historical name), Log Send Queue (KB) (caption), Log_Send_Queue (attribute name), and LGSENQUE (column name).

Mirrored Write Transaction Per Sec

The number of transactions that have been written to the mirrored database in the last second, and are waiting for the logs to be sent to the mirrored database to complete the commit operation. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRRORED_WRITE_TRANSACTION_PER_SEC or MRWRTRPS (historical name), Mirrored Write Transaction Per Sec (caption), Mirrored_Write_Transaction_Per_Sec (attribute name), and MRWRTRPS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Recovery Queue

The number of log records that are in the recovery queue because they are not redone. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RECOVERY_QUEUE or RECQUE (historical name), Recovery Queue (caption), Recovery_Queue (attribute name), and RECQUE (column name).

Redo Blocked Per Sec

The number of times that the redo operation is currently being blocked per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDO_BLOCKED_PER_SEC or REDBLKPS (historical name), Redo Blocked Per Sec (caption), Redo_Blocked_Per_Sec (attribute name), and REDBLKPS (column name).

Redo Bytes Remaining (KB)

The amount of log bytes (in KB) that has not completed the redo phase. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDO_BYTES_REMAINING or REDBTREM (historical name), Redo Bytes Remaining (KB) (caption), Redo_Bytes_Remaining (attribute name), and REDBTREM (column name).

Redone Bytes Per Sec

The amount of log bytes that was redone on the secondary database in the last second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDONE_BYTES_PER_SEC or REDBTPS (historical name), Redone Bytes Per Sec (caption), Redone_Bytes_Per_Sec (attribute name), and REDBTPS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: **TIMESTAMP** (historical name), **Timestamp** (caption), **Timestamp** (attribute name), and **TIMESTAMP** (column name).

Total Log Requiring Undo (KB)

The amount of log data (in KB) that must be undone. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **TOTAL_LOG_REQURING_UNDO** or **TLGREQUD** (historical name), **Total Log Requiring Undo (KB)** (caption), **Total_Log_Requring_Undo** (attribute name), and **TLGREQUD** (column name).

Undo Log Remaining (KB)

The amount of log data (in KB) that remains to be undone. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **UNDO_LOG_REMAINING** or **UNDLGREM** (historical name), **Undo Log Remaining (KB)** (caption), **Undo_Log_Remaining** (attribute name), and **UNDLGREM** (column name).

MS SQL Availability Database Summary data set

The MS SQL Availability Database Summary data set provides a summary of the information about availability databases that are hosted by the SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 5 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: **NODE** (historical name), **Node** (caption), **ORIGINNODE** (attribute name), and **ORIGINNODE** (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: **TIMESTAMP** (historical name), **Timestamp** (caption), **Timestamp** (attribute name), and **TIMESTAMP** (column name).

Total Non Healthy Databases

The total number of the availability databases that have been non-healthy since the agent startup. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: **TOTAL_NON_HEALTHY_DATABASES** or **TONHLDB** (historical name), **Total Non Healthy Databases** (caption), **Total_Non_Healthy_Databases** (attribute name), and **TONHLDB** (column name).

Max Non Healthy Databases

The maximum number of the availability databases that have been non-healthy since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_NON_HEALTHY_DATABASES or MXNHLDB (historical name), Max Non Healthy Databases (caption), Max_Non_Healthy_Databases (attribute name), and MXNHLDB (column name).

Max Non-Online Databases

The maximum number of databases that have not been online since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_NON-ONLINE_DATABASES or MXNONLDB (historical name), Max Non-Online Databases (caption), Max_Non-Online_Databases (attribute name), and MXNONLDB (column name).

Max Unsynchronized Commit DB Replicas

The maximum number of availability database replicas on which the transaction commit operation has not been synchronized since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_UNSYNCHRONIZED_COMMIT_DB_REPLICAS or MXNSCDBR (historical name), Max Unsynchronized Commit DB Replicas (caption), Max_Unsynchronized_Commit_DB_Replicas (attribute name), and MXNSCDBR (column name).

Max Unsynchronized Databases

The maximum number of the availability databases that have been unsynchronized since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_UNSYNCHRONIZED_DATABASES or MXUSYNDB (historical name), Max Unsynchronized Databases (caption), Max_Unsynchronized_Databases (attribute name), and MXUSYNDB (column name).

Min Non Healthy Databases

The minimum number of the availability databases that have been non-healthy since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_NON_HEALTHY_DATABASES or MNNHLDB (historical name), Min Non Healthy Databases (caption), Min_Non_Healthy_Databases (attribute name), and MNNHLDB (column name).

Min Non-Online Databases

The minimum number of databases that have not been online since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_NON-ONLINE_DATABASES or MNNONLDB (historical name), Min Non-Online Databases (caption), Min_Non-Online_Databases (attribute name), and MNNONLDB (column name).

Min Unsynchronized Commit DB Replicas

The minimum number of availability database replicas on which the transaction commit operation has not been synchronized since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_UNSYNCHRONIZED_COMMIT_DB_REPLICAS or MNNSCDBR (historical name), Min Unsynchronized Commit DB Replicas (caption), Min_Unsynchronized_Commit_DB_Replicas (attribute name), and MNNSCDBR (column name).

Min Unsynchronized Databases

The minimum number of the availability databases that have been unsynchronized since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_UNSYNCHRONIZED_DATABASES or MNUSYNDB (historical name), Min Unsynchronized Databases (caption), Min_Unsynchronized_Databases (attribute name), and MNUSYNDB (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Databases

The total number of availability databases that are hosted by the SQL Server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DATABASES or TOTDBS (historical name), Total Databases (caption), Total_Databases (attribute name), and TOTDBS (column name).

Total Non-Online Databases

The total number of databases that have not been online since the agent startup. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NON-ONLINE_DATABASES or TONONLDB (historical name), Total Non-Online Databases (caption), Total_Non-Online_Databases (attribute name), and TONONLDB (column name).

Total Primary Databases

The total number of primary availability databases that are hosted by the SQL Server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PRIMARY_DATABASES or TOTPMDB (historical name), Total Primary Databases (caption), Total_Primary_Databases (attribute name), and TOTPMDB (column name).

Total Secondary Databases

The total number of secondary availability databases that are hosted by the SQL Server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SECONDARY_DATABASES or TOTSCDB (historical name), Total Secondary Databases (caption), Total_Secondary_Databases (attribute name), and TOTSCDB (column name).

Total Unsynchronized Commit DB Replicas

The total number of availability database replicas on which the transaction commit operation has not been synchronized since the agent startup. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_UNSYNCHRONIZED_COMMIT_DB_REPLICAS or TONSCDBR (historical name), Total Unsynchronized Commit DB Replicas (caption), Total_Unsynchronized_Commit_DB_Replicas (attribute name), and TONSCDBR (column name).

Total Unsynchronized Databases

The total number of availability databases that have been unsynchronized since the agent startup. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_UNSYNCHRONIZED_DATABASES or TOUSYNDB (historical name), Total Unsynchronized Databases (caption), Total_Unsynchronized_Databases (attribute name), and TOUSYNDB (column name).

MS SQL Availability Group Listener Details data set

The MS SQL Availability Group Listener Details data set provides information about the availability group listeners of all the availability groups that are hosted by the local SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Group ID

The availability group ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), Group ID (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The availability group name. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME_U or GRPNAME (historical name), Group Name (caption), Group_Name_U (attribute name), and GRPNAME (column name).

Is Conformant

Indicates whether the IP address of the availability groups listener is conformant. An IP address is conformant if one of the following conditions exists:

- The IP configuration is created by the T-SQL statement.
- The IP configuration is created outside of the SQL Server, but the IP configuration can be modified by the T-SQL statement.

The type is integer (32-bit numeric property) with enumerated values. The following values are defined: No (1), Yes (2), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IS_CONFORMANT or ISCONF (historical name), Is Conformant (caption), Is_Conformant (attribute name), and ISCONF (column name).

Is DHCP

Indicates whether the IP address of the availability groups listener is configured by the Dynamic Host Configuration Protocol (DHCP). The type is integer (32-bit numeric property) with enumerated values. The following values are defined: No (1), Yes (2), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IS_DHCP or ISDHCP (historical name), Is DHCP (caption), Is_DHCP (attribute name), and ISDHCP (column name).

Listener ID

The resource ID of the availability group listener in the WSFC cluster. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LISTENER_ID or LISNID (historical name), Listener ID (caption), Listener_ID (attribute name), and LISNID (column name).

Listener IP Address

The IP address of the availability group listener. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LISTENER_IP_ADDRESS or LISNIPAD (historical name), Listener IP Address (caption), Listener_IP_Address (attribute name), and LISNIPAD (column name).

Listener Name

The name of the availability group listener. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LISTENER_NAME_U or LISNNM (historical name), Listener Name (caption), Listener_Name_U (attribute name), and LISNNM (column name).

Listener Subnet IP

The IP address of the subnet to which the availability group listener belongs. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LISTENER_SUBNET_IP or LISNSUBIP (historical name), Listener Subnet IP (caption), Listener_Subnet_IP (attribute name), and LISNSUBIP (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Port

The TCP port number of the availability group listener. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PORT (historical name), Port (caption), Port (attribute name), and PORT (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

State

The state of the availability group listener in the WSFC cluster. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Offline (1), Online (2), Online Pending (3), Failed (4), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STATE (historical name), State (caption), State (attribute name), and STATE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Availability Groups Details data set

The MS SQL Availability Groups data set provides details about the availability groups that are hosted by the local SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Group ID

The ID of the availability group to which the replica belongs. The value format is an alphanumeric string with a maximum of 128 characters. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), *Group ID* (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The name of the availability group to which the replica belongs. The value format is an alphanumeric string with a maximum of 384 characters. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME_U or GRPNAME (historical name), *Group Name* (caption), Group_Name_U (attribute name), and GRPNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Primary Replica ID

The ID of the primary replica. The format is an alphanumeric string with a maximum of 128 characters. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PRIMARY_REPLICA_ID or REPLEID (historical name), *Primary Replica ID* (caption), Primary_Replica_ID (attribute name), and REPLEID (column name).

Synchronization Health

A summary of the synchronization health of all the replicas in the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Healthy (1), Partially Healthy (2), Healthy (3), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SYNCHRONIZATION_HEALTH or SYNHEAL (historical name), *Synchronization Health* (caption), Synchronization_Health (attribute name), and SYNHEAL (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Backup Preference

The preferred location to perform backup operations on the availability databases that are available in the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Primary Replica (1), Secondary Replica (2), Primary Replica acceptable if no Secondary Replica Available (3), No Preference (4), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BACKUP_PREFERENCE or BACKPRE (historical name), Backup Preference (caption), Backup_Preference (attribute name), and BACKPRE (column name).

Database Count

The total number of databases that are connected to the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_COUNT or DBCOUNT (historical name), Database Count (caption), Database_Count (attribute name), and DBCOUNT (column name).

Failure Condition

The condition level that is specified for triggering an automatic failure of the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: SQL Service Down or Group Lease Expired (1), health check timeout Threshold Group Exceeded (2), Critical SQL Server Internal Errors (3), SQL Server Internal Errors (4), Unsolvable Deadlock or Exhaustion of SQL Engine (5), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAILURE_CONDITION or FAILCOND (historical name), Failure Condition (caption), Failure_Condition (attribute name), and FAILCOND (column name).

Group Role

Indicates whether the local server instance is a primary or a secondary replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Resolving (1), Primary (2), Secondary (3), Invalid (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUPS_ROLE or GRPROLE (historical name), Group Role (caption), Groups_Role (attribute name), and GRPROLE (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Last Group Role

Indicates whether the local server instance is a primary or a secondary replica for the availability group in the last sample. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Resolving (1), Primary (2), Secondary (3), Invalid (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_GROUPS_ROLE or LGRPROLE (historical name), Last Group Role (caption), Last_Groups_Role (attribute name), and LGRPROLE (column name).

Local Replica Recovery Health

The recovery health status of the replica that is hosted by the local SQL Server instance in the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: In Progress (1), Online (2), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCAL_REPLICA_RECOVERY_HEALTH or LOCREPL (historical name), Local Replica Recovery Health (caption), Local_Replica_Recovery_Health (attribute name), and LOCREPL (column name).

Primary Replica

The name of the primary replica. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PRIMARY_REPLICA_U or REPLENAM (historical name), Primary Replica (caption), Primary_Replica_U (attribute name), and REPLENAM (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Secondary Replicas Count

The number of secondary replicas in the availability group. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SECONDARY_REPLICA_COUNT or NUMSECRP (historical name), Secondary Replicas Count (caption), Secondary_Replica_Count (attribute name), and NUMSECRP (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Wait Timeout (ms)

The duration (in milliseconds) that the system waits for a response from the SQL Server instance before the instance is assumed to be slow or not responding. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: WAIT_TIMEOUT or WAITTME (historical name), Wait Timeout (ms) (caption), Wait_Timeout (attribute name), and WAITTME (column name).

MS SQL Availability Groups Details In Cluster data set

The MS SQL Availability Groups Details In Cluster data set provides information about each availability group in the Windows Server Failover Clustering (WFSC) cluster even if the local SQL Server instance is not a part of any availability group in the WFSC cluster. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Backup Preference

The preferred location that is specified for performing backup operations on the availability database. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Primary Replica (1), Secondary Replica (2), Primary Replica Acceptable If No Secondary Replica Available (3), No Preference (4), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BACKUP_PREFERENCE or BACKPRE (historical name), Backup Preference (caption), Backup_Preference (attribute name), and BACKPRE (column name).

Cluster Resource ID

The WSFC cluster resource ID. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLUSTER_RESOURCE_ID or CLURESID (historical name), Cluster Resource ID (caption), Cluster_Resource_ID (attribute name), and CLURESID (column name).

Failure Condition

The value that indicates when an automatic failover occurs. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: SQL Service Down or Group Lease Expired (1), Health Check Timeout Threshold Group Exceeded (2), Critical SQL Server Internal Errors (3), SQL Server Internal Errors (4), Unsolvable Deadlock or Exhaustion of SQL Engine (5), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAILURE_CONDITION or FAILCOND (historical name), Failure Condition (caption), Failure_Condition (attribute name), and FAILCOND (column name).

Group ID

The availability group ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), Group ID (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The availability group name. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME_U or GRPNAME (historical name), Group Name (caption), Group_Name_U (attribute name), and GRPNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Resource Group ID

The WSFC cluster resource group ID. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RESOURCE_GROUP_ID or RESGRPID (historical name), Resource Group ID (caption), Resource_Group_ID (attribute name), and RESGRPID (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Wait Timeout (ms)

The duration (in milliseconds) that the system waits for the sp_server_diagnostics stored procedure to return the health information of the SQL Server instance before the server instance is assumed to be slow or not responding. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: WAIT_TIMEOUT or WAITTME (historical name), Wait Timeout (ms) (caption), Wait_Timeout (attribute name), and WAITTME (column name).

MS SQL Availability Groups Summary data set

The MS SQL Availability Groups Summary data set provides a summary of the availability groups that are hosted by the local SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Groups

The total number of availability groups that are available on the SQL Server instance in the current interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_GROUPS or TOTGRP (historical name), *Total Groups* (caption), Total_Groups (attribute name), and TOTGRP (column name).

Total Non Healthy Groups

The total number of availability groups that are currently not healthy in the current interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NON_HEALTHY_GROUPS or TOTNHGRP (historical name), *Total Non Healthy Groups* (caption), Total_Non_Healthy_Groups (attribute name), and TOTNHGRP (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Max Non Healthy Groups

The maximum number of availability groups in which at least one availability replica that is hosted by the SQL Server instance has not been healthy since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_NON_HEALTHY_GROUPS or MAXNHGRP (historical name), Max Non Healthy Groups (caption), Max_Non_Healthy_Groups (attribute name), and MAXNHGRP (column name).

Max Non Online Local Replica Groups

The maximum number of availability groups in which at least one availability replica that is hosted by the SQL Server instance has not been online since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_NON_ONLINE_LOCAL_REPLICA_GROUPS or MAXOLRG (historical name), Max Non Online Local Replica Groups (caption), Max_Non_Online_Local_Replica_Groups (attribute name), and MAXOLRG (column name).

Min Non Healthy Groups

The minimum number of availability groups in which at least one availability replica that is hosted by the SQL Server instance has not been healthy since the agent startup. The type is integer (32-bit

numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_NON_HEALTHY_GROUPS or MINNHGRP (historical name), Min Non Healthy Groups (caption), Min_Non_Healthy_Groups (attribute name), and MINNHGRP (column name).

Min Non Online Local Replica Groups

The minimum number of availability groups in which at least one availability replica that is hosted by the SQL Server instance has not been online since the agent startup. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_NON_ONLINE_LOCAL_REPLICA_GROUPS or MINOLRG (historical name), Min Non Online Local Replica Groups (caption), Min_Non_Online_Local_Replica_Groups (attribute name), and MINOLRG (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Non Online Local Replica Groups

The total number of availability groups in which at least one availability replica that is hosted by the SQL Server instance is not online in the current interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NON_ONLINE_LOCAL_REPLICA_GROUPS or TNOGRP (historical name), Total Non Online Local Replica Groups (caption), Total_Non_Online_Local_Replica_Groups (attribute name), and TNOGRP (column name).

Total Primary Role Group

The total number of availability groups in which the SQL Server instance is the primary replica in the current interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PRIMARY_ROLE_GROUP or TOTPGRP (historical name), Total Primary Role Group (caption), Total_Primary_Role_Group (attribute name), and TOTPGRP (column name).

Total Secondary Role Group

The total number of availability groups in which the SQL Server instance is the secondary replica in the current interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SECONDARY_ROLE_GROUP or TOTSGRP (historical name), Total Secondary Role Group (caption), Total_Secondary_Role_Group (attribute name), and TOTSGRP (column name).

MS SQL Availability Replicas Details data set

The MS SQL Availability Replicas Details data set provides details about the availability replicas in each availability group. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Availability Mode

Indicates whether the availability replica is running in the synchronous-commit mode or the asynchronous-commit mode. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Asynchronous Commit (1), Synchronous Commit (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVAILABILITY_MODE or AVAMODE (historical name), Availability Mode (caption), Availability_Mode (attribute name), and AVAMODE (column name).

Backup Priority

The priority that is specified to perform backup operations on the replica. The value of this attribute is displayed as 0 if you select the Exclude Replica check box while configuring an availability replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BACKUP_PRIORITY or BACKPROR (historical name), Backup Priority (caption), Backup_Priority (attribute name), and BACKPROR (column name).

Create Date

The date when the replica was created. This attribute shows No Data if the SQL Server query returns the NULL value. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL Server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CREATE_DATE or CRDATE (historical name), Create Date (caption), Create_Date (attribute name), and CRDATE (column name).

Endpoint URL

The database mirroring endpoint that is specified by the user. The value format is an alphanumeric string. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENDPOINT_URL or ENDPTURL (historical name), Endpoint URL (caption), Endpoint_URL (attribute name), and ENDPTURL (column name).

Failover Mode

Indicates whether the failover mode of the availability replica is manual or automatic. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Manual (2), Automatic (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAILOVER_MODE or FAILOVMOD (historical name), Failover Mode (caption), Failover_Mode (attribute name), and FAILOVMOD (column name).

Group ID

The ID of the availability group to which the replica belongs. The value format is an alphanumeric string. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), Group ID (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The name of the availability group to which the replica belongs. The value format is an alphanumeric string. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME or GRPNAME (historical name), Group Name (caption), Group_Name (attribute name), and GRPNAME (column name).

Modify Date

The date when the replica was last modified. This attribute shows No Data if the SQL Server query returns the NULL value. The type is timestamp with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MODIFY_DATE or MODDATE (historical name), Modify Date (caption), Modify_Date (attribute name), and MODDATE (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Owner SID

The security identifier (SID) of the external owner of the replica. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OWNER_SID or OWNSID (historical name), Owner SID (caption), Owner_SID (attribute name), and OWNSID (column name).

Primary Connection Type

Indicates the type of connection that is allowed for a primary replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), All (3), Read Write (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PRIMARY_CONNECTION_TYPE or PRMCNTYP (historical name), Primary Connection Type (caption), Primary_Connection_Type (attribute name), and PRMCNTYP (column name).

Replica ID

The ID of the replica. The value format is an alphanumeric string. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_ID or REPID (historical name), Replica ID (caption), Replica_ID (attribute name), and REPID (column name).

Replica Server Name

The server and instance name of the location that hosts the replica. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_SERVER_NAME or RPSRNAME (historical name), Replica Server Name (caption), Replica_Server_Name (attribute name), and RPSRNAME (column name).

Role

Indicates whether a replica is currently a primary or secondary replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Resolving (1), Primary (2), Secondary (3), Invalid (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROLE (historical name), Role (caption), Role (attribute name), and ROLE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Secondary Connection Type

Indicates the type of connection that is allowed for a secondary replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), No (1), Read Only (2), All (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SECONDARY_CONNECTION_TYPE or SECCNTYP (historical name), Secondary Connection Type (caption), Secondary_Connection_Type (attribute name), and SECCNTYP (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Session Timeout

The duration (in seconds) that the replica waits to receive a message from another replica before assuming that the connection between the primary and the secondary replica has failed. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SESSION_TIMEOUT or SESTMOUT (historical name), Session Timeout (caption), Session_Timeout (attribute name), and SESTMOUT (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Availability Replicas Details In Cluster data set

The MS SQL Availability Replicas Details In Cluster data set provides detailed information about each availability replica in the Windows Server Failover Clustering (WFSC) cluster. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Cluster Node

The name of the cluster node where the availability replica is located. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLUSTER_NODE or CLUNODE (historical name), Cluster Node (caption), Cluster_Node (attribute name), and CLUNODE (column name).

Group ID

The availability group ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), Group ID (caption), Group_ID (attribute name), and GRPID (column name).

Group Name

The availability group name. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME_U or GRPNAME (historical name), Group Name (caption), Group_Name_U (attribute name), and GRPNAME (column name).

Join State

Indicates whether the availability replica is joined to the availability group in the WSFC cluster. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Joined (1), Joined As Stand Alone Instance (2), Joined As Failover Cluster Instance (3), Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOIN_STATE or JOINSTAT (historical name), Join State (caption), Join_State (attribute name), and JOINSTAT (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Replica ID

The availability replica ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_ID or REPID (historical name), Replica ID (caption), Replica_ID (attribute name), and REPID (column name).

Replica Server Name

The server and instance name of the location where the availability replica is hosted. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_SERVER_NAME or REPSVRNM (historical name), Replica Server Name (caption), Replica_Server_Name (attribute name), and REPSVRNM (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Availability Replicas Statistics data set

The MS SQL Availability Replicas Statistics data set provides information about the statistics of the availability replicas that are available in the AlwaysOn availability group, which is currently hosted on the SQL Server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later.

This data set contains the following attributes:

Bytes Received Per Sec

The number of bytes that are currently being received from the availability replica per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BYTES_RECEIVED_PER_SEC or BTRECPSC (historical name), Bytes Received Per Sec (caption), Bytes_Received_Per_Sec (attribute name), and BTRECPSC (column name).

Bytes Sent Per Sec

The number of bytes that are currently being sent per second from the primary replica to the secondary replica, or from the secondary replica to the primary replica. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BYTES_SENT_PER_SEC or BTSENPSC (historical name), Bytes Sent Per Sec (caption), Bytes_Sent_Per_Sec (attribute name), and BTSENPSC (column name).

Bytes Sent To Transport Per Sec

The number of bytes that are currently being sent per second through the network from the primary replica to the secondary replica, or from the secondary replica to the primary replica. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BYTES_SENT_TO_TRANSPORT_PER_SEC or BTSNTRPS (historical name), Bytes Sent To Transport Per Sec (caption), Bytes_Sent_To_Transport_Per_Sec (attribute name), and BTSNTRPS (column name).

Flow Control Per Sec

The number of flow control operations that have been currently initiated per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FLOW_CONTROL_PER_SEC or FLWCNTPS (historical name), Flow Control Per Sec (caption), Flow_Control_Per_Sec (attribute name), and FLWCNTPS (column name).

Flow Control Time (ms)

The time (in milliseconds) that the log stream messages waited to send the flow control information to the availability database. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FLOW_CONTROL_TIME or FLWCNTTM (historical name), Flow Control Time (ms) (caption), Flow_Control_Time (attribute name), and FLWCNTTM (column name).

Instance Name

The name of the local SQL Server instance. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INSTANCE_NAME or AVINSNAM (historical name), Instance Name (caption), Instance_Name (attribute name), and AVINSNAM (column name).

Messages Sent To Transport Per Sec

The number of AlwaysOn messages that are currently being sent per second through the network from the primary replica to the secondary replica, or from the secondary replica to the primary replica. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGES_SENT_TO_TRANSPORT_PER_SEC or MSGTTRPS (historical name), Messages Sent To Transport Per Sec (caption), Messages_Sent_To_Transport_Per_Sec (attribute name), and MSGTTRPS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Received Messages Per Sec

The number of AlwaysOn messages that are currently being received from the primary replica per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RECEIVED_MESSAGES_PER_SEC or RECMGPSC (historical name), Received Messages Per Sec (caption), Received_Messages_Per_Sec (attribute name), and RECMGPSC (column name).

Resent Messages Per Sec

The number of AlwaysOn messages that are currently being resent to the availability replica per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RESENT_MESSAGES_PER_SEC or RSNMSGPS (historical name), Resent Messages Per Sec (caption), Resent_Messages_Per_Sec (attribute name), and RSNMSGPS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Sent Messages Per Sec

The number of AlwaysOn messages that are currently being sent to the availability replica per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SENT_MESSAGES_PER_SEC or SNDMSGPS (historical name), Sent Messages Per Sec (caption), Sent_Messages_Per_Sec (attribute name), and SNDMSGPS (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Availability Replicas Status data set

The MS SQL Availability Replicas Status data set provides details about the status of the replicas in each availability group. This data set is supported for SQL Server 2012 Enterprise Edition, or later. If the SQL Server instance is a primary replica in an availability group, the MS SQL Availability Replicas Status Summary data set displays the details of all the replicas that are present in all the availability groups. However, if the SQL Server instance is a secondary replica in an availability group, the MS SQL Availability Replicas Status Summary data set displays the details of only those replicas that are hosted by the local instance. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 5 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Connection State

The current connection state of the availability replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Disconnected (1), Connected (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONNECTION_STATE or CONSTATE (historical name), *Connection State* (caption), Connection_State (attribute name), and CONSTATE (column name).

Group ID

The ID of the availability group to which the replica belongs. The value format is an alphanumeric string with a maximum of 64 characters. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_ID or GRPID (historical name), *Group ID* (caption), Group_ID (attribute name), and GRPID (column name).

Local

Indicates whether the replica is hosted by the local instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), No (1), Yes (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCAL (historical name), *Local* (caption), Local (attribute name), and LOCAL (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Operational State

The current operational state of the availability replica. This attribute shows the value No Data on the replica that is not local. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Pending Failover (1), Pending (2), Online (3), Offline (4), Failed (5), Failed No Quorum (6), Invalid (7). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OPERATIONAL_STATE or OPSTATE (historical name), *Operational State* (caption), Operational_State (attribute name), and OPSTATE (column name).

Recovery Health

Indicates whether the databases that are connected to the availability group are online or are being recovered after a failover. This attribute shows the value No Data on the replica that is not local. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), In Progress (1), Online (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RECOVERY_HEALTH or RECHLTH (historical name), *Recovery Health* (caption), Recovery_Health (attribute name), and RECHLTH (column name).

Replica ID

The ID of the replica. The value format is an alphanumeric string with a maximum of 64 characters. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICA_ID or REPID (historical name), *Replica ID* (caption), Replica_ID (attribute name), and REPID (column name).

Role

Indicates whether the replica is a primary or a secondary replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Resolving (1), Primary (2), Secondary (3), Invalid (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROLE (historical name), *Role* (caption), Role (attribute name), and ROLE (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), *Server* (caption), Server (attribute name), and SERVERID (column name).

Synchronization Health

The synchronization state of all the databases that are connected to the availability group on the availability replica. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Not Healthy (1), Partially Healthy (2), Healthy (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SYNCHRONIZATION_HEALTH or SYNCHLTH (historical name), *Synchronization Health* (caption), Synchronization_Health (attribute name), and SYNCHLTH (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Databases

The total number of databases that are hosted by the availability replica. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DATABASES or TOTDBS (historical name), *Total Databases* (caption), Total_Databases (attribute name), and TOTDBS (column name).

Group Name

The name of the availability group to which the replica belongs. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GROUP_NAME or GRPNAME (historical name), Group Name (caption), Group_Name (attribute name), and GRPNAME (column name).

Last Connect Error Number

The last connection error number. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_CONNECT_ERROR_NUMBER or LTSYCERNO (historical name), Last Connect Error Number (caption), Last_Connect_Error_Number (attribute name), and LTSYCERNO (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

MS SQL Availability Replicas Status Summary data set

The MS SQL Availability Replicas Status Summary data set provides a summary of the status of availability replicas in all the availability groups that are hosted on the local server instance. This data set is supported for SQL Server 2012 Enterprise Edition, or later. If the SQL Server instance is a primary replica in an availability group, the MS SQL Availability Replicas Status Summary data set displays the details of all the replicas that are present in all the availability groups. However, if the SQL Server instance is a secondary replica in an availability group, the MS SQL Availability Replicas Status Summary data set displays the details of only those replicas that are hosted by the local instance. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 5 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Non Healthy Replicas

The total number of non-healthy replicas in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NON_HEALTHY_REPLICAS or TOTNHLRP (historical name), *Total Non Healthy Replicas* (caption), Total_Non_Healthy_Replicas (attribute name), and TOTNHLRP (column name).

Max Disconnected Replicas

The maximum number of replicas that have been disconnected since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_DISCONNECTED_REPLICAS or MXDSCNRP (historical name), Max Disconnected Replicas (caption), Max_Disconnected_Replicas (attribute name), and MXDSCNRP (column name).

Max Local Failed Replicas

The maximum number of local replicas that have failed since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_LOCAL_FAILED_REPLICAS or MAXLCFRP (historical name), Max Local Failed Replicas (caption), Max_Local_Failed_Replicas (attribute name), and MAXLCFRP (column name).

Max Non Healthy Replicas

The maximum number of replicas that have not been healthy since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_NON_HEALTHY_REPLICAS or MAXNHLRP (historical name), Max Non Healthy Replicas (caption), Max_Non_Healthy_Replicas (attribute name), and MAXNHLRP (column name).

Max Secondary Replicas

The maximum number of secondary replicas that have been connected since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_SECONDARY_REPLICAS or MAXSCREP (historical name), Max Secondary Replicas (caption), Max_Secondary_Replicas (attribute name), and MAXSCREP (column name).

Min Disconnected Replicas

The minimum number of replicas that have been disconnected since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric

property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_DISCONNECTED_REPLICAS or MNDSCNRP (historical name), Min Disconnected Replicas (caption), Min_Disconnected_Replicas (attribute name), and MNDSCNRP (column name).

Min Local Failed Replicas

The minimum number of local replicas that have failed since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_LOCAL_FAILED_REPLICAS or MINLCFRP (historical name), Min Local Failed Replicas (caption), Min_Local_Failed_Replicas (attribute name), and MINLCFRP (column name).

Min Non Healthy Replicas

The minimum number of replicas that have not been healthy since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_NON_HEALTHY_REPLICAS or MINNHLRP (historical name), Min Non Healthy Replicas (caption), Min_Non_Healthy_Replicas (attribute name), and MINNHLRP (column name).

Min Secondary Replicas

The minimum number of secondary replicas that have been connected since the agent startup in all the availability groups that are hosted on the local server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIN_SECONDARY_REPLICAS or MINSCREP (historical name), Min Secondary Replicas (caption), Min_Secondary_Replicas (attribute name), and MINSCREP (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Connected Replicas

The total number of connected replicas in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CONNECTED_REPLICAS or TOTCNREP (historical name), Total Connected Replicas (caption), Total_Connected_Replicas (attribute name), and TOTCNREP (column name).

Total Disconnected Replicas

The total number of disconnected replicas in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DISCONNECTED_REPLICAS or TOTDCNRP (historical name), Total Disconnected Replicas (caption), Total_Disconnected_Replicas (attribute name), and TOTDCNRP (column name).

Total Local Failed Replicas

The total number of local replicas that failed in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_LOCAL_FAILED_REPLICAS or TOTLCFRP (historical name), Total Local Failed Replicas (caption), Total_Local_Failed_Replicas (attribute name), and TOTLCFRP (column name).

Total Local Replicas

The total number of local replicas in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_LOCAL_REPLICAS or TOTLCREP (historical name), Total Local Replicas (caption), Total_Local_Replicas (attribute name), and TOTLCREP (column name).

Total Primary Replicas

The total number of primary replicas in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PRIMARY_REPLICAS or TOTPMREP (historical name), Total Primary Replicas (caption), Total_Primary_Replicas (attribute name), and TOTPMREP (column name).

Total Secondary Replicas

The total number of secondary replicas in all the availability groups that are hosted on the local server instance. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SECONDARY_REPLICAS or TOTSCREP (historical name), Total Secondary Replicas (caption), Total_Secondary_Replicas (attribute name), and TOTSCREP (column name).

MS SQL Batch Stats data set

The MS SQL Batch Stats data set contains attributes that you can use to monitor detailed information about batch requests.

This data set contains the following attributes:

Auto-Param Attempts Per Sec

The number of auto-parameterization attempts per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUTO_PARAM_ATTEMPTS or APARAMSEC (historical name), Auto-Param Attempts Per Sec (caption), Auto_Param_Attempts (attribute name), and APARAMSEC (column name).

Batch Requests Per Sec

The number of Transact-SQL command batch requests received per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BATCH_REQUESTS or BATCHSEC (historical name), Batch Requests Per Sec (caption), Batch_Requests (attribute name), and BATCHSEC (column name).

Failed Auto-Params Per Sec

The number of failed auto-parameterization attempts per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAILED_AUTO_PARAMS or FAPARAMSEC (historical name), Failed Auto-Params Per Sec (caption), Failed_Auto_Params (attribute name), and FAPARAMSEC (column name).

Forced Parameterization Per Sec

The number of successful forced parameterization attempts per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FORCED_PARAMS or FPARAMSEC (historical name), Forced Parameterization Per Sec (caption), Forced_Params (attribute name), and FPARAMSEC (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Safe Auto-Params Per Sec

The number of safe auto-parameterization attempts per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SAFE_AUTO-PARAMS or SAPARAMSEC (historical name), Safe Auto-Params Per Sec (caption), Safe_Auto-Params (attribute name), and SAPARAMSEC (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Server Attention Rate

The number of attentions per second. An attention is a request by the client to end the currently running request. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_ATTENTION_RATE or ATTNRATE (historical name), Server Attention Rate (caption), Server_Attention_Rate (attribute name), and ATTNRATE (column name).

Server Compilations Per Sec

The number of server compilations per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_COMPILATIONS or CMPLSEC (historical name), Server Compilations Per Sec (caption), Server_Compilations (attribute name), and CMPLSEC (column name).

Server Re-Compilations Per Sec

The number of times the server recompiles per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_RE_COMPILATIONS or RCMPLSEC (historical name), Server Re-Compilations Per Sec (caption), Server_Re_Compilations (attribute name), and RCMPLSEC (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Unsafe Auto-Params Per Sec

The number of unsafe auto-parameterization attempts per second. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: UNSAFE_AUTO_PARAMS or UAPARAMSEC (historical name), Unsafe Auto-Params Per Sec (caption), Unsafe_Auto_Params (attribute name), and UAPARAMSEC (column name).

MS SQL Blocked Queries Details data set

The MS SQL Blocked Queries Details data set contains attributes that provide information about the queries that are currently blocked on the SQL Server.

This data set contains the following attributes:

Blocked Database Name (Unicode)

The name of the database against which the blocked query is currently being run. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKED_DATABASE_NAME_U or BLKEDDBNM (historical name), Blocked Database Name (Unicode) (caption), Blocked_Database_Name_U (attribute name), and BLKEDDBNM (column name).

Blocked Parent Query Text

The SQL text of the parent of the blocking query. The value format is an alphanumeric string with a maximum of 2048 characters. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKED_PARENT_QUERY_TEXT or BLKEDPQXT (historical name), Blocked Parent Query Text (caption), Blocked_Parent_Query_Text (attribute name), and BLKEDPQXT (column name).

Blocked Query Text

The SQL text of the query that is currently blocked on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKED_QUERY_TEXT or BLKEDQTEXT (historical name), Blocked Query Text (caption), Blocked_Query_Text (attribute name), and BLKEDQTEXT (column name).

Blocked Session ID

The session ID that is blocking the query that is currently being run on the SQL Server. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKED_SESSION_ID or BLKEDSESID (historical name), Blocked Session ID (caption), Blocked_Session_ID (attribute name), and BLKEDSESID (column name).

Blocked User Name (Unicode)

The name of the user who has run the blocked SQL query. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKED_USER_NAME_U or BLKEDUSRNM (historical name), Blocked User Name (Unicode) (caption), Blocked_User_Name_U (attribute name), and BLKEDUSRNM (column name).

Blocking Query Text

The SQL text of the query blocks one or more than one queries on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_QUERY_TEXT or BLKINQTEXT (historical name), Blocking Query Text (caption), Blocking_Query_Text (attribute name), and BLKINQTEXT (column name).

Blocking Reason

The reason for blocking the query. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_REASON or BLKINGRSON (historical name), Blocking Reason (caption), Blocking_Reason (attribute name), and BLKINGRSON (column name).

Blocking Resource

The resource for which the query is currently blocked on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_RESOURCE or BLKINGSRC (historical name), Blocking Resource (caption), Blocking_Resource (attribute name), and BLKINGSRC (column name).

Blocking Session ID

The ID of the session that is currently blocking the execution of a query on the SQL Server. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_SESSION_ID or BLKINSESID (historical name), Blocking Session ID (caption), Blocking_Session_ID (attribute name), and BLKINSESID (column name).

Blocking User Name

The login name of the user who ran the blocking query. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_USER_NAME or BLKINUSRNM (historical name), Blocking User Name (caption), Blocking_User_Name (attribute name), and BLKINUSRNM (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Wait Time (Sec.)

The total wait time (in seconds) since the query has been blocked on the SQL Server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_WAIT_TIME_SEC or TOTWETTIMS (historical name), Total Wait Time (Sec.) (caption), Total_Wait_Time_Sec (attribute name), and TOTWETTIMS (column name).

MS SQL Configuration data set

The MS SQL Configuration data set contains attributes that you can use to monitor the configuration of a SQL server.

This data set contains the following attributes:

Config Parameter

The name of the configuration parameter. The value format is an alphanumeric string with a maximum of 30 characters; for example, Partition group. The sysconfigures and syscurconfigs system tables store the configuration parameters. Use the parameter name to track the performance of a particular configuration parameter. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: CONFIG_PARAMETER or CONFIGNAME (historical name), Config Parameter (caption), Config_Parameter (attribute name), and CONFIGNAME (column name).

Config Value

The value for the configuration parameter. The value format is an alphanumeric string with a maximum of 12 characters; for example, 1000. The value was initially specified in the sysconfigures system table. Monitor this value to track its effect on performance, and to compare the initial setting to the maximum, minimum, and run values. The type is string.

The following names are defined for this attribute: CONFIG_VALUE or CONFIGVAL (historical name), Config Value (caption), Config_Value (attribute name), and CONFIGVAL (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Maximum Value

The maximum value that can be specified for the configuration parameter. The value format is an alphanumeric string with a maximum of 12 characters; for example, 2147483647. The type is string.

The following names are defined for this attribute: MAXIMUM_VALUE or MAXVAL (historical name), Maximum Value (caption), Maximum_Value (attribute name), and MAXVAL (column name).

Minimum Value

The minimum value that can be specified for the configuration parameter. The value format is an alphanumeric string with a maximum of 12 characters; for example, 1. The type is string.

The following names are defined for this attribute: MINIMUM_VALUE or MINVAL (historical name), Minimum Value (caption), Minimum_Value (attribute name), and MINVAL (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Parm Type

The type of configuration parameter. If the configuration parameter is dynamic, a server restart is not required. The type is string with enumerated values. The following values are defined: Static (0), Dynamic (1), Advanced (2), Dynamic and Advanced (3), Unknown (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PARM_TYPE or PARMTYPE (historical name), Parm Type (caption), Parm_Type (attribute name), and PARMTYPE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Run Value

The value the SQL server is using for the configuration parameter. The value format is an alphanumeric string with a maximum of 12 characters; for example, 1000. The type is string.

The following names are defined for this attribute: RUN_VALUE or RUNVALUE (historical name), Run Value (caption), Run_Value (attribute name), and RUNVALUE (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Customized SQL Query Details data set

The MS SQL Customized SQL Query Details data set provides a summary of the customized SQL queries that are defined in the properties file. This data set is supported for SQL Server 2012 Enterprise Edition, or later. This data set is not supported in MSSQL APM agent.

This data set contains the following attributes:

Customized SQL Definition File

The location where the file that contains customized SQL queries is saved. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CUSTOMIZED_SQL_DEFINITION_FILE or CUSQLDFFL (historical name), Customized SQL Definition File (caption), Customized_SQL_Definition_File (attribute name), and CUSQLDFFL (column name).

Database Name

The maximum number of the availability databases that have been unsynchronized since the agent startup. The type is string with enumerated values. The following values are defined: Not Collected

(Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), Database Name (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Definition File Last Modified Time

The date and time when the customized SQL queries definition file was last modified. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEFINITION_FILE_LAST_MODIFIED_TIME or DFFLLSTMDT (historical name), Definition File Last Modified Time (caption), Definition_File_Last_Modified_Time (attribute name), and DFFLLSTMDT (column name).

Last Execution Error Message

The reason why the SQL query failed when it was last executed. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_EXECUTION_ERROR_MESSAGE or LSTEXERMSG (historical name), Last Execution Error Message (caption), Last_Execution_Error_Message (attribute name), and LSTEXERMSG (column name).

Last Execution Status

The status of the last execution of the SQL query. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), Success (1), Fail (2), Rejected (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_EXECUTION_STATUS or LSTEXSTUS (historical name), Last Execution Status (caption), Last_Execution_Status (attribute name), and LSTEXSTUS (column name).

Last Execution Time

The time when the SQL query was last executed. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_EXECUTION_TIME or LSTEXTIM (historical name), Last Execution Time (caption), Last_Execution_Time (attribute name), and LSTEXTIM (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

SQL ID

The customized SQL query ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_ID or SQLID (historical name), SQL ID (caption), SQL_ID (attribute name), and SQLID (column name).

SQL Query

The content of the customized SQL query that is defined in the properties file. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_QUERY or SQLQUE (historical name), SQL Query (caption), SQL_Query (attribute name), and SQLQUE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Customized SQL Query Result data set

This data set displays the result of customized SQL queries. The data collection for this data set is limited to 1000 rows. This data set is not supported in MSSQL APM agent.

This data set contains the following attributes:

Integer Column Name 1

The name of the first column of integer data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_NAME_1 or INTCOLNM1 (historical name), Integer Column Name 1 (caption), Integer_Column_Name_1 (attribute name), and INTCOLNM1 (column name).

Integer Column Name 2

The name of the second column of integer data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_NAME_2 or INTCOLNM2 (historical name), Integer Column Name 2 (caption), Integer_Column_Name_2 (attribute name), and INTCOLNM2 (column name).

Integer Column Name 3

The name of the third column of integer data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_NAME_3 or INTCOLNM3 (historical name), Integer Column Name 3 (caption), Integer_Column_Name_3 (attribute name), and INTCOLNM3 (column name).

Integer Column Name 4

The name of the fourth column of integer data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_NAME_4 or INTCOLNM4 (historical name), Integer Column Name 4 (caption), Integer_Column_Name_4 (attribute name), and INTCOLNM4 (column name).

Integer Column Name 5

The name of the fifth column of integer data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_NAME_5 or INTCOLNM5 (historical name), Integer Column Name 5 (caption), Integer_Column_Name_5 (attribute name), and INTCOLNM5 (column name).

Integer Column Value 1

The value of the first column of integer data type in an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_VALUE_1 or INTCOLVL1 (historical name), Integer Column Value 1 (caption), Integer_Column_Value_1 (attribute name), and INTCOLVL1 (column name).

Integer Column Value 2

The value of the second column of integer data type in an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_VALUE_2 or INTCOLVL2 (historical name), Integer Column Value 2 (caption), Integer_Column_Value_2 (attribute name), and INTCOLVL2 (column name).

Integer Column Value 3

The value of the third column of integer data type in an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_VALUE_3 or INTCOLVL3 (historical name), Integer Column Value 3 (caption), Integer_Column_Value_3 (attribute name), and INTCOLVL3 (column name).

Integer Column Value 4

The value of the fourth column of integer data type in an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_VALUE_4 or INTCOLVL4 (historical name), Integer Column Value 4 (caption), Integer_Column_Value_4 (attribute name), and INTCOLVL4 (column name).

Integer Column Value 5

The value of the fifth column of integer data type in an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INTEGER_COLUMN_VALUE_5 or INTCOLVL5 (historical name), Integer Column Value 5 (caption), Integer_Column_Value_5 (attribute name), and INTCOLVL5 (column name).

More Columns

Indicates whether the number of columns in an SQL query result is more than five for the string and integer data types or more than four for the time data type. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Yes (1), No (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MORE_COLUMNS or MORECOL (historical name), More Columns (caption), More_Columns (attribute name), and MORECOL (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

SQL ID

The customized SQL query ID. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_ID or SQLID (historical name), SQL ID (caption), SQL_ID (attribute name), and SQLID (column name).

String Column Name 1

The name of the first column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_NAME_1 or STRCOLNM1 (historical name), String Column Name 1 (caption), String_Column_Name_1 (attribute name), and STRCOLNM1 (column name).

String Column Name 2

The name of the second column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_NAME_2 or STRCOLNM2 (historical name), String Column Name 2 (caption), String_Column_Name_2 (attribute name), and STRCOLNM2 (column name).

String Column Name 3

The name of the third column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_NAME_3 or STRCOLNM3 (historical name), String Column Name 3 (caption), String_Column_Name_3 (attribute name), and STRCOLNM3 (column name).

String Column Name 4

The name of the fourth column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_NAME_4 or STRCOLNM4 (historical name), String Column Name 4 (caption), String_Column_Name_4 (attribute name), and STRCOLNM4 (column name).

String Column Name 5

The name of the fifth column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_NAME_5 or STRCOLNM5 (historical name), String Column Name 5 (caption), String_Column_Name_5 (attribute name), and STRCOLNM5 (column name).

String Column Value 1

The value of the first column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_VALUE_1 or STRCOLVL1 (historical name), String Column Value 1 (caption), String_Column_Value_1 (attribute name), and STRCOLVL1 (column name).

String Column Value 2

The value of the second column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_VALUE_2 or STRCOLVL2 (historical name), String Column Value 2 (caption), String_Column_Value_2 (attribute name), and STRCOLVL2 (column name).

String Column Value 3

The value of the third column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_VALUE_3 or STRCOLVL3 (historical name), String Column Value 3 (caption), String_Column_Value_3 (attribute name), and STRCOLVL3 (column name).

String Column Value 4

The value of the fourth column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_VALUE_4 or STRCOLVL4 (historical name), String Column Value 4 (caption), String_Column_Value_4 (attribute name), and STRCOLVL4 (column name).

String Column Value 5

The value of the fifth column of string data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STRING_COLUMN_VALUE_5 or STRCOLVL5 (historical name), String Column Value 5 (caption), String_Column_Value_5 (attribute name), and STRCOLVL5 (column name).

Time Column Name 1

The name of the first column of time data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_NAME_1 or TIMCOLNM1 (historical name), Time Column Name 1 (caption), Time_Column_Name_1 (attribute name), and TIMCOLNM1 (column name).

Time Column Name 2

The name of the second column of time data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_NAME_2 or TIMCOLNM2 (historical name), Time Column Name 2 (caption), Time_Column_Name_2 (attribute name), and TIMCOLNM2 (column name).

Time Column Name 3

The name of the third column of time data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_NAME_3 or TIMCOLNM3 (historical name), Time Column Name 3 (caption), Time_Column_Name_3 (attribute name), and TIMCOLNM3 (column name).

Time Column Name 4

The name of the fourth column of time data type in an SQL query. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_NAME_4 or TIMCOLNM4 (historical name), Time Column Name 4 (caption), Time_Column_Name_4 (attribute name), and TIMCOLNM4 (column name).

Time Column Value 1

The value of the first column of time data type in an SQL query. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_VALUE_1 or TIMCOLVL1 (historical name), Time Column Value 1 (caption), Time_Column_Value_1 (attribute name), and TIMCOLVL1 (column name).

Time Column Value 2

The value of the second column of time data type in an SQL query. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_VALUE_2 or TIMCOLVL2 (historical name), Time Column Value 2 (caption), Time_Column_Value_2 (attribute name), and TIMCOLVL2 (column name).

Time Column Value 3

The value of the third column of time data type in an SQL query. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_VALUE_3 or TIMCOLVL3 (historical name), Time Column Value 3 (caption), Time_Column_Value_3 (attribute name), and TIMCOLVL3 (column name).

Time Column Value 4

The value of the fourth column of time data type in an SQL query. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_COLUMN_VALUE_4 or TIMCOLVL4 (historical name), Time Column Value 4 (caption), Time_Column_Value_4 (attribute name), and TIMCOLVL4 (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Database Detail data set

The MS SQL Database Detail data set contains attributes that you can use to monitor the performance and usage of MS SQL database. The data collection for this data set is limited to 6000 rows. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Accessible

Access validation of the database. When the database is in offline, restoring, or transition state, the value of this attribute is displayed as No. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), YES (0), NO (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACCESS_VALIDATION or ACSVAL (historical name), Accessible (caption), Access_Validation (attribute name), and ACSVAL (column name).

Active Transactions

The number of active transactions for the database. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTIVE_TRANSACTIONS or ACTIVTRANS (historical name), Active Transactions (caption), Active_Transactions (attribute name), and ACTIVTRANS (column name).

Aggregate Database Freespace (MB)

The total amount of freespace (in MB) that is available in all the data files and on the storage device. The value format is a decimal number with two decimal places allowed. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGGREGATE_DATABASE_FREESPACE or ADBFREE (historical name), *Aggregate Database Freespace (MB)* (caption), Aggregate_Database_Freespace (attribute name), and ADBFREE (column name).

Data File Names

The names of all physical files for data that make up this database. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_FILE_NAMES or DATAFILENM (historical name), *Data File Names* (caption), Data_File_Names (attribute name), and DATAFILENM (column name).

Data Size (MB)

The number of megabytes (MB) allocated for the data only segments of the database. The value format is a decimal number with two decimal places allowed; for example, 50.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_SIZE or DBSIZE (historical name), *Data Size (MB)* (caption), Data_Size (attribute name), and DBSIZE (column name).

Database Growth Percent

The percentage of growth for the database from the last sample to the current sample. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_GROWTH_PERCENT or DBGRWTHPCT (historical name), *Database Growth Percent* (caption), Database_Growth_Percent (attribute name), and DBGRWTHPCT (column name).

Database Max Growth Size(MB)

The maximum size to which the database can grow in MB. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_MAXIMUM_GROWTH_SIZE or DBMAXGRSZ (historical name), *Database Max Growth Size(MB)* (caption), Database_Maximum_Growth_Size (attribute name), and DBMAXGRSZ (column name).

Database Name (Unicode)

The name of the database. Each database name is unique. The SQL server also assigns each database its own identification number. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), *Database Name (Unicode)* (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Database Name

The name of the database. The value format is an alphanumeric string with a maximum of 30 characters; for example, KOQ3. Each database name is unique. The SQL server also assigns each database its own identification number. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), *Database Name* (caption), Database_Name (attribute name), and DBNAME (column name).

Database Space Pct Used

The amount of aggregate space (in megabytes) used in the database as a percentage of total space allowed. Use this attribute to warn you when you need to extend the database. If you run out of space, you can no longer use the database. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_SPACE_PERCENT_USED or PCTDBUSE (historical name), *Database Space Pct Used* (caption), Database_Space_Percent_Used (attribute name), and PCTDBUSE (column name).

Database State

Reports the database state. The type is string with enumerated values. The following values are defined: ONLINE (ONLINE), RESTORING (RESTORING), RECOVERING (RECOVERING), RECOVERY PENDING (RECOVERY_PENDING), SUSPECT (SUSPECT), EMERGENCY (EMERGENCY), OFFLINE (OFFLINE), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_STATE or DBSTATE (historical name), *Database State* (caption), Database_State (attribute name), and DBSTATE (column name).

Database Status

Reports when the database is unavailable. If a database is offline, you cannot access it. Use this attribute to warn that the database has become unavailable. When the database is in transition state, the database status is displayed as Available. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Available (0), Not Available (1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_STATUS or DBSTAT (historical name), *Database Status* (caption), Database_Status (attribute name), and DBSTAT (column name).

DB ID

The ID for the database. The value format is an alphanumeric string with a maximum of 12 characters; for example, 2156. This value is stored in the sysdatabases table. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DB_ID or DBID (historical name), *DB ID* (caption), DB_ID (attribute name), and DBID (column name).

DB Owner (Unicode)

The SQL server-assigned user ID for the owner of the database. Use the create database command to establish this identifier. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DB_OWNER_U or UOWNER (historical name), *DB Owner (Unicode)* (caption), DB_Owner_U (attribute name), and UOWNER (column name).

DBO Only Access

Indicates whether the database has a status of DBO only. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DBO_ONLY_ACCESS or DBOONLY (historical name), *DBO Only Access* (caption), DBO_Only_Access (attribute name), and DBOONLY (column name).

Error Status

Indicates whether the database has an error status. A database with an error status has a status of suspect, crashed, or recovery. Check the status bits in the sysdatabases table to determine the cause of the error. Use the database consistency checker (dbcc) to verify the database integrity. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_STATUS or ERRSTATUS (historical name), *Error Status* (caption), Error_Status (attribute name), and ERRSTATUS (column name).

Free Space Accounting Suppressed

Indicates whether the free space accounting option is disabled for the database. The no-free-space-actg option turns off free space accounting on non-log segments only. Information about free space is inaccurate when free space accounting is turned off. Use the no-free-space-actg option and the checkpoint command to speed recovery. No time is needed to count free space for non-log segments. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FREE_SPACE_ACCOUNTING_SUPPRESSED or NOFREESPA (historical name), *Free Space Accounting Suppressed* (caption), Free_Space_Accounting_Suppressed (attribute name), and NOFREESPA (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), Host_Name (attribute name), and HOSTNAME (column name).

Log File Names

The names of all physical files for logs that make up this database. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_FILE_NAMES or LOGFILENM (historical name), *Log File Names* (caption), Log_File_Names (attribute name), and LOGFILENM (column name).

Log Freespace (MB)

The number of megabytes (MB) of free space in the transaction log for the database. The value includes the number of megabytes of free space on the log only partitions. The value format is a decimal number with two decimal places, for example, 8. 00. Various types of transactions, such as mass updates and bulk copying can involve extensive logging. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_FREESPACE or LOGFREE (historical name), *Log Freespace (MB)* (caption), Log_Freespace (attribute name), and LOGFREE (column name).

Log Growths

The total number of times the database transaction log is expanded since the SQL Server was started. The format is an integer. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_GROWTHS or LOGGRTH (historical name), *Log Growths* (caption), Log_Growths (attribute name), and LOGGRTH (column name).

Log Max Growth Size(MB)

The maximum size to which the log can grow in MB. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_MAXIMUM_GROWTH_SIZE or LOGMAXGRSZ (historical name), *Log Max Growth Size(MB)* (caption), Log_Maximum_Growth_Size (attribute name), and LOGMAXGRSZ (column name).

Log Size (MB)

The number of megabytes (MB) allocated for the transaction log for the database. The value includes the number of megabytes allocated for the transaction log on the log only partitions. The value format is a decimal number with two decimal places allowed; for example, 500. The appropriate size for a

transaction log depends on how the database is used. Several factors, such as the number of transactions, type of transactions, and number of users, effect sizing. Evaluate the need to truncate the transaction log to prevent it from filling up. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SIZE or LOGSIZE (historical name), *Log Size (MB)* (caption), Log_Size (attribute name), and LOGSIZE (column name).

Log Space Pct Used

The percentage of the transaction log that is full. When the log has filled up all of its allocated space, transactions that require logging are rejected. Use this attribute to set an alert based on a percent full threshold, then spawn a task to dump or truncate the transaction log or to do both. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SPACE_PERCENT_USED or PCTLOGUSED (historical name), *Log Space Pct Used* (caption), Log_Space_Percent_Used (attribute name), and PCTLOGUSED (column name).

Long Running Transaction Name

The long running transaction name. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LONG_RUNNING_TRANSACTION_NAME or LRTNAME (historical name), *Long Running Transaction Name* (caption), Long_Running_Transaction_Name (attribute name), and LRTNAME (column name).

Long Running Transaction Process ID

The longest running transaction. The Long running transactions are transactions that are open for longer than the LongRunningTransColl standard collector parameter. The format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LONG_RUNNING_TRANSACTION_PROCESS_ID or LRTPID (historical name), *Long Running Transaction Process ID* (caption), Long_Running_Transaction_Process_ID (attribute name), and LRTPID (column name).

Long Running Transaction Time (Sec.)

The time for which the longest running transaction is active (in seconds). The format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LONG_RUNNING_TRANSACTION_TIME_PER_SEC or LRTTPS (historical name), *Long Running Transaction Time (Sec.)* (caption), Long_Running_Transaction_Time_Per_Sec (attribute name), and LRTTPS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Oldest Open Transaction(min)

The age (in minutes) of the oldest open transaction in the database transaction log. Use this attribute to determine up to which point in time the transaction log can be truncated and backed up. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OLDEST_OPEN_TRANSACTION or OLDOPTRAN (historical name), *Oldest Open Transaction(min)* (caption), Oldest_Open_Transaction (attribute name), and OLDOPTRAN (column name).

Page Verify

The option that identifies and notifies incomplete I/O transactions that have occurred because of disk I/O errors. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: NONE (0), TORN PAGE DETECTION (1), CHECKSUM (2), Not Collected (-1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_VERIFY or PAGEVERIFY (historical name), *Page Verify* (caption), Page_Verify (attribute name), and PAGEVERIFY (column name).

Read Only Access

Indicates whether the database has a status of read only. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: READ_ONLY_ACCESS or READONLY (historical name), *Read Only Access* (caption), Read_Only_Access (attribute name), and READONLY (column name).

Recovery Model

The method used to maintain the transaction log. The three types of recovery models include Simple, Full, and Bulk-logged. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: FULL (1), BULK LOGGED (2), SIMPLE (3), Not Collected (-1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RECOVERY_MODEL or RECOVMODEL (historical name), *Recovery Model* (caption), Recovery_Model (attribute name), and RECOVMODEL (column name).

Replicated Transactions

The number of transactions in the publication database transaction log that are marked for replication but have not yet been delivered to the distribution database. Use this attribute to discover if there are any bottlenecks in the replication process. Bottlenecks cause delays in the published transactions reaching the subscriber database. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICATED_TRANSACTIONS or REPTRANS (historical name), *Replicated Transactions* (caption), Replicated_Transactions (attribute name), and REPTRANS (column name).

Replication Latency (ms)

The number of milliseconds between the time a transaction marked for replication is entered into the publication database transaction log and the time it is read out of the log and delivered to the distribution database. Use this attribute to discover if there are any bottlenecks in the replication process. Bottlenecks cause delays in the published transactions reaching the subscriber database. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICATION_LATENCY or REPTLAT (historical name), *Replication Latency (ms)* (caption), Replication_Latency (attribute name), and REPTLAT (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters. For example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), *Server* (caption), Server (attribute name), and SERVERID (column name).

Single User Access

Indicates whether the database has a status of single user. A database with a status of single user can only be accessed by one user at a time. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SINGLE_USER_ACCESS or SINGUSER (historical name), *Single User Access* (caption), Single_User_Access (attribute name), and SINGUSER (column name).

Table Count

The number of tables that exist in the database. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_COUNT or TABLECNT (historical name), *Table Count* (caption), Table_Count (attribute name), and TABLECNT (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Devices

The number of devices allocated for the database. The value format is an integer in the range 1-128; for example, 4. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DEVICES or NUMDEVICES (historical name), *Total Devices* (caption), Total_Devices (attribute name), and NUMDEVICES (column name).

Transactions per Sec.

The number of transactions started for the database per second for the current interval. When the number of transactions collected by the agent in the last interval is greater than or equal to the number of transactions collected in the current interval, then 0 (zero) is displayed as the attribute value on the console. When the number of transactions collected in the last interval is less than the number of transactions collected in the current interval, then the value is calculated as follows: The number of transactions collected in the last interval is subtracted from the number of transactions collected in the current interval, and the value after subtraction is divided by the time interval between these two collections. When the value of transactions collected in the last or the current interval, is negative, then Not_Collected is displayed on the console. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TRANSACTIONS_PER_SECOND or TRANSSEC (historical name), *Transactions per Sec.* (caption), Transactions_per_Second (attribute name), and TRANSSEC (column name).

Write Transactions per sec

The number of transactions written to the database and committed in the last second. The format is an integer. When the number of written transactions collected by the agent in the last interval is greater than or equal to the number of written transactions collected in the current interval, then 0 (zero) is displayed as the attribute value on the console. When the number of written transactions collected in the last interval is less than the number of written transactions collected in the current interval, then the value is calculated as follows: The number of written transactions collected in the last interval is subtracted from the number of written transactions collected in the current interval, and the value after subtraction is divided by the time interval between these two collections. When the value of written transactions, which are collected in the last or the current interval, is negative, then Not_Collected is displayed on the console. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: WRITE_TRANSACTIONS_PER_SEC or WRTRANSEC (historical name), *Write Transactions per sec* (caption), Write_Transactions_per_sec (attribute name), and WRTRANSEC (column name).

Aggregate Database Freespace Pct

The percentage of the total amount of freespace that is available in all the data files and on the storage device out of the total space allowed. The value format is a decimal number with two decimal places allowed. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGGREGATE_DATABASE_FREESPACE_PCT or ADBFREEPCT (historical name), Aggregate Database Freespace Pct (caption), Aggregate_Database_Freespace_Pct (attribute name), and ADBFREEPCT (column name).

Data Files Freespace (MB)

The amount of freespace (in MB) that is available in all the data files. The value format is a decimal number with two decimal places allowed. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_FILES_FREESPACE or DTFFREE (historical name), Data Files Freespace (MB) (caption), Data_Files_Freespace (attribute name), and DTFFREE (column name).

Data Freespace (MB)

The number of megabytes (MB) of free space for the database. The value includes the megabytes of free data only. The value format is a decimal number with two decimal places allowed; for example, 10.00. This value will be same as that of Aggregate Database Freespace. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_FREESPACE or DBFREE (historical name), Data Freespace (MB) (caption), Data_Freespace (attribute name), and DBFREE (column name).

Data Freespace Percent

The percentage of maximum available free space for the database. The value format is percentage with two decimal places allowed; for example, 20.00. The product calculates this percentage from the maximum growth size for the database and the amount of space that is used by the database. Consider using the alter database command to expand the database size. This value will be same as that of Aggregate Database Freespace pct. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_FREESPACE_PERCENT or PCTDBFREE (historical name), Data Freespace Percent (caption), Data_Freespace_Percent (attribute name), and PCTDBFREE (column name).

DB Owner

The SQL server-assigned user ID for the owner of the database. The value format is an alphanumeric string with a maximum of 8 characters; for example, DBO. Use the create database command to establish this identifier. The type is string.

The following names are defined for this attribute: DB_OWNER or OWNER (historical name), DB Owner (caption), DB_Owner (attribute name), and OWNER (column name).

Dump Tran Date

The timestamp that indicates the date on which the dump transaction command was last executed for the database. The dump transaction command copies the uncommitted transactions in the transaction log. Refer to this date to determine when the latest backup of the transaction log was made. This attribute value may not be available for the SQL Server to capture. The format is YYYYMMDD, where, YYYY indicates year, MM indicates month, and DD indicates day. For example, 20020501 indicates May 1, 2002. The type is string with enumerated values. The following values are defined: Not Collected (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DUMP_TRAN_DATE or DUMPTRDATE (historical name), Dump Tran Date (caption), Dump_Tran_Date (attribute name), and DUMPTRDATE (column name).

Last Database Size(MB)

The database size in MB from the previous sample. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_DATABASE_SIZE or LASTDBSIZE (historical name), Last Database Size(MB) (caption), Last_Database_Size (attribute name), and LASTDBSIZE (column name).

Log Bytes Flushed Per Sec

The number of bytes of the log file that are flushed in the current interval. When the value of flushed bytes of the log file collected by the agent in the last interval is greater than or equal to the value of flushed bytes of the log file collected in the current interval, then 0 (zero) is displayed as the attribute value on the console. When the value of flushed bytes of the log file collected in the last interval is less than the value of flushed bytes of the log file collected in the current interval, then the value is calculated as follows: The value of flushed bytes of the log file collected in the last interval is subtracted from the value of flushed bytes of the log file collected in the current interval, and the value after subtraction is divided by the time interval between these two collections. When the value of flushed bytes of the log file collected in the last or the current interval, is negative, then Not_Collected is displayed on the console. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_FLUSHED_PER_SEC or LOGBFLSHPS (historical name), Log Bytes Flushed Per Sec (caption), Log_Bytes_Flushed_Per_Sec (attribute name), and LOGBFLSHPS (column name).

Log Freespace Percent

The percentage of free space in the transaction log for the database. The value format is a percentage with two decimal places, for example, 20.00. Various types of transactions, such as mass updates and bulk copying in, can involve extensive logging. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_FREESPACE_PERCENT or PCTLOGFREE (historical name), Log Freespace Percent (caption), Log_Freespace_Percent (attribute name), and PCTLOGFREE (column name).

Mirroring Role

Indicates whether the database is mirror or principal, or if mirroring is not applied. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Mirror (2), Principal (1), No Data (-3), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRRORING_ROLE or MIRROLE (historical name), Mirroring Role (caption), Mirroring_Role (attribute name), and MIRROLE (column name).

No CKPT After Recovery

Indicates whether a record for the checkpoint is added to the transaction log when the database is recovered. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NO_CKPT_AFTER_RECOVERY or NOCKPTREC (historical name), No CKPT After Recovery (caption), No_CKPT_After_Recovery (attribute name), and NOCKPTREC (column name).

Replicated Transaction Rate

The rate per second at which transactions have been read out of the publication database transaction log and delivered to the distribution database. Use this attribute to discover if there are any

bottlenecks in the replication process. Bottlenecks cause delays in the published transactions reaching the subscriber database. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPLICATED_TRANSACTION_RATE or REPRATE (historical name), Replicated Transaction Rate (caption), Replicated_Transaction_Rate (attribute name), and REPRATE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Select Into Bulkcopy Enabled

Indicates whether the select into/bulkcopy option is enabled for the database. The type is string with enumerated values. The following values are defined: No (No), Yes (Yes), No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SELECT_INTO_BULKCOPY_ENABLED or SELINTOB (historical name), Select Into Bulkcopy Enabled (caption), Select_Into_Bulkcopy_Enabled (attribute name), and SELINTOB (column name).

Suspect Database

Reports whether a database is marked as suspect. Use this attribute to warn you that a database cannot be loaded. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: True (0), False (1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SUSPECT_DATABASE or SUSPECTDB (historical name), Suspect Database (caption), Suspect_Database (attribute name), and SUSPECTDB (column name).

MS SQL Database Mirroring data set

The MS SQL Database Mirroring data set provides information about the database mirroring objects. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Bytes Received Per Sec

The number of bytes received per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are

defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BYTES_RECEIVED_PER_SEC or BRPERSEC (historical name), Bytes Received Per Sec (caption), Bytes_Received_Per_Sec (attribute name), and BRPERSEC (column name).

Bytes Sent Per Sec

The number of bytes received per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BYTES_SENT_PER_SEC or BSPERSEC (historical name), Bytes Sent Per Sec (caption), Bytes_Sent_Per_Sec (attribute name), and BSPERSEC (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Log Bytes Received Per Sec

The number of bytes of log received per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_RECEIVED_PER_SEC or LBRPSEC (historical name), Log Bytes Received Per Sec (caption), Log_Bytes_Received_Per_Sec (attribute name), and LBRPSEC (column name).

Log Bytes Redone From Cache Per Sec

The number of log bytes redone from the Database Mirroring log cache per second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_REDONE_FROM_CACHE_PER_SEC or LBRFCPS (historical name), Log Bytes Redone From Cache Per Sec (caption), Log_Bytes_Redone_From_Cache_Per_Sec (attribute name), and LBRFCPS (column name).

Log Bytes Sent From Cache Per Sec

The number of log bytes sent from the Database Mirroring log cache in the last second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_SENT_FROM_CACHE_PER_SEC or LBSFCPS (historical name), Log Bytes Sent From Cache Per Sec (caption), Log_Bytes_Sent_From_Cache_Per_Sec (attribute name), and LBSFCPS (column name).

Log Bytes Sent Per Sec

The number of log bytes sent per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_BYTES_SENT_PER_SEC or LBSPSEC (historical name), Log Bytes Sent Per Sec (caption), Log_Bytes_Sent_Per_Sec (attribute name), and LBSPSEC (column name).

Log Compressed Bytes Rcvd Per Sec

The number of compressed log bytes received in the last second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_COMPRESSED_BYTES_RCVD_PER_SEC or LCBRPS (historical name), Log Compressed Bytes Rcvd Per Sec (caption), Log_Compressed_Bytes_Rcvd_Per_Sec (attribute name), and LCBRPS (column name).

Log Compressed Bytes Sent Per Sec

The number of compressed bytes of log sent in the last second. The number of compressed log bytes sent in the last second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_COMPRESSED_BYTES_SENT_PER_SEC or LCBSPS (historical name), Log Compressed Bytes Sent Per Sec (caption), Log_Compressed_Bytes_Sent_Per_Sec (attribute name), and LCBSPS (column name).

Log Harden Time (ms)

The number of milliseconds the log blocks waited to be hardened to disk in the last second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_HARDEN_TIME or LHARDTIME (historical name), Log Harden Time (ms) (caption), Log_Harden_Time (attribute name), and LHARDTIME (column name).

Log Remaining For Undo(MB)

The total number of log bytes (in KB) that remain to be scanned by the new mirror server after failover. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_REMAINING_FOR_UNDO or LRFUKB (historical name), Log Remaining For Undo(MB) (caption), Log_Remaining_For_Undo (attribute name), and LRFUKB (column name).

Log Scanned For Undo(MB)

The total number of log bytes (in KB) that are scanned by the new mirror server after failover. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SCANNED_FOR_UNDO or LSFUKB (historical name), Log Scanned For Undo(MB) (caption), Log_Scanned_For_Undo (attribute name), and LSFUKB (column name).

Log Send Flow Control Time (ms)

The time (in milliseconds) awaited by the log stream messages for the send flow control in the last second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated

values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SEND_FLOW_CONTROL_TIME or LOGSFCT (historical name), Log Send Flow Control Time (ms) (caption), Log_Send_Flow_Control_Time (attribute name), and LOGSFCT (column name).

Log Send Queue(MB)

The total number of megabytes of log that are not sent to the mirror server. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOG_SEND_QUEUE or LOGSQKB (historical name), Log Send Queue(MB) (caption), Log_Send_Queue (attribute name), and LOGSQKB (column name).

Mirror Name (Unicode)

The name of the database. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRROR_NAME or MIRNAME (historical name), Mirror Name (Unicode) (caption), Mirror_Name (attribute name), and MIRNAME (column name).

Mirrored Write Transactions Per Sec

The number of transactions which wrote to the mirrored database in the last second, that waited for log to be sent to the mirror. The number of transactions (that waited for log to be sent to the mirror) written to the mirrored database in the last second. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRRORED_WRITE_TRANSACTIONS_PER_SEC or MWTPS (historical name), Mirrored Write Transactions Per Sec (caption), Mirrored_Write_Transactions_Per_Sec (attribute name), and MWTPS (column name).

Mirroring Role

The current role of the local database in the database mirroring session. The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Mirror (0), Principal (1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRRORING_ROLE or MIRROLE (historical name), Mirroring Role (caption), Mirroring_Role (attribute name), and MIRROLE (column name).

Mirroring State

The state of the mirror database. The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Suspended (0), Disconnected (1), Synchronizing (2), Pending Failover (3), Synchronized (4), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRRORING_STATE or MIRSTATE (historical name), Mirroring State (caption), Mirroring_State (attribute name), and MIRSTATE (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Pages Sent Per Sec

The number of pages sent per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value

Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGES_SENT_PER_SEC or PSPERSEC (historical name), Pages Sent Per Sec (caption), Pages_Sent_Per_Sec (attribute name), and PSPERSEC (column name).

Receives Per Sec

The number of mirroring messages received per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RECEIVES_PER_SEC or RECPSEC (historical name), Receives Per Sec (caption), Receives_Per_Sec (attribute name), and RECPSEC (column name).

Redo Bytes Per Sec

The number of bytes of log redone by the mirror database per second. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDO_BYTES_PER_SEC or REDOPSEC (historical name), Redo Bytes Per Sec (caption), Redo_Bytes_Per_Sec (attribute name), and REDOPSEC (column name).

Redo Queue(MB)

The number of megabytes that redo on the mirror database is behind the hardened log. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDO_QUEUE or REDOQKB (historical name), Redo Queue(MB) (caption), Redo_Queue (attribute name), and REDOQKB (column name).

Sample Timestamp

The time when data is collected. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Send Receive Ack Time

The time (in milliseconds) awaited by the messages to receive an acknowledgment from the partner. The format is an integer. This attribute is not supported on the SQL Server 2005 Enterprise Edition. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SEND_RECEIVE_ACK_TIME or SRATIME (historical name), Send Receive Ack Time (caption), Send_Receive_Ack_Time (attribute name), and SRATIME (column name).

Sends Per Sec

The number of sends initiated per second at the time of monitoring interval. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SENDS_PER_SEC or SPERSEC (historical name), Sends Per Sec (caption), Sends_Per_Sec (attribute name), and SPERSEC (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Synchronization State

The state of synchronization (is failover possible). The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Failover not possible (0), Failover possible (1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SYNCHRONIZATION_STATE or SYNCSTATE (historical name), Synchronization State (caption), Synchronization_State (attribute name), and SYNCSTATE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transaction Delay (ms)

The time (in milliseconds) awaited by the transaction termination to receive an acknowledgment. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TRANSACTION_DELAY or TRANDEL (historical name), Transaction Delay (ms) (caption), Transaction_Delay (attribute name), and TRANDEL (column name).

MS SQL Database Summary data set

The MS SQL Database Summary data set contains attributes that you can use to monitor summary information about SQL server databases. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 3 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Databases Critical

The total number of databases that are in the critical state. A database is in the critical state when the (log space percent that is used exceeds 90), or (database space percent that is used exceeds 90), or (database status is equal to 1), or (database state is equal to offline or emergency), or (database error status is equal to yes). The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DATABASES_CRITICAL or TDBCRT (historical name), *Total Databases Critical* (caption), Total_Databases_Critical (attribute name), and TDBCRT (column name).

Total Databases Warning

The total number of databases that are in the warning state. A database is in the warning state when the (log space percent that is used is greater than 70 and less than equal to 90), or (database space percent that is used is greater than 70 and less than equal to 90), or (database state is not equal to offline, online, and emergency). The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_DATABASES_WARNING or TDBWARN (historical name), *Total Databases Warning* (caption), Total_Databases_Warning (attribute name), and TDBWARN (column name).

Total Databases

The number of databases for the SQL server. The value format is an integer; for example, 10. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_DATABASES or NUMDBS (historical name), *Total Databases* (caption), Total_Databases (attribute name), and NUMDBS (column name).

Total DBs DBO Only

The number of databases with a status of DBO only. A database with a status of database owner can be accessed only by users with DBO authority. The value format is an integer; for example, 5. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_DBs_DBO_ONLY or NUMDBO (historical name), *Total DBs DBO Only* (caption), Total_DBs_DBO_Only (attribute name), and NUMDBO (column name).

Total DBs in Error

The number of databases with an error status. A database with an error status is a database with a status of suspect, crashed, or recovery. The value format is an integer; for example, 2. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_DBs_IN_ERROR or NUMERRST (historical name), *Total DBs in Error* (caption), Total_DBs_in_Error (attribute name), and NUMERRST (column name).

Total DBs No Free Space Accounting

The number of databases that have the free space accounting option disabled. The value format is an integer; for example, 5. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_DBs_NO_FREE_SPACE_ACCOUNTING or NUMNOFRSP (historical name), *Total DBs No Free Space Accounting* (caption), Total_DBs_No_Free_Space_Accounting (attribute name), and NUMNOFRSP (column name).

Total DBs Read Only

The number of databases with a status of read only. The value format is an integer; for example, 5. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_DBs_READ_ONLY or NUMRO (historical name), *Total DBs Read Only* (caption), Total_DBs_Read_Only (attribute name), and NUMRO (column name).

Total DBs Single User

The number of databases with a status of single user. A database with a status of single user can be accessed by only one user at a time. The value format is an integer; for example, 2. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_DBs_SINGLE_USER or NUMSING (historical name), *Total DBs Single User* (caption), Total_DBs_Single_User (attribute name), and NUMSING (column name).

Current Interval (Sec.)

The number of seconds that have elapsed between the previous sample and the current sample. The value format is an integer; for example, 90. A sample contains the data that IBM Tivoli Monitoring for

Microsoft SQL Server collects about each SQL server. New data becomes available if a new interval has occurred and data has been refreshed at the CMS hub. The type is integer (32-bit numeric property).

The following names are defined for this attribute: CURRENT_INTERVAL or CURRINTL (historical name), Current Interval (Sec.) (caption), Current_Interval (attribute name), and CURRINTL (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Minimum Pct Data Freespace

The lowest percentage of aggregate free space in data only and data and log segments on a device allocated to a database. The value format is a percentage with two decimal places allowed; for example, 5.00. The amount of space needed by the database depends on its anticipated activity. Set alerts for abnormal conditions. Evaluate the possibility of adding more space to the database. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MINIMUM_PCT_DATA_FREESPACE or MINDBFREE (historical name), Minimum Pct Data Freespace (caption), Minimum_Pct_Data_Freespace (attribute name), and MINDBFREE (column name).

Minimum Pct Log Freespace

The lowest percentage of free space in log only and data and log segments on a device allocated for the database transaction log. The value format is a percentage with two decimal places allowed; for example, 5.00. The amount of space needed by the transaction log depends on the type and quantity of the transactions and the frequency of backups. Set alerts for abnormal conditions. Evaluate the possibility of expanding the log. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MINIMUM_PCT_LOG_FREESPACE or MINLOGFREE (historical name), Minimum Pct Log Freespace (caption), Minimum_Pct_Log_Freespace (attribute name), and MINLOGFREE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters. For example: CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Device Detail data set

The MS SQL Device Detail data set contains attributes that you can use to monitor the usage and performance of devices. The data collection for this data set is limited to 10000 rows. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Autogrowth Unit

Indicates the unit of file growth. If the autogrowth feature is enabled for a device, the file growth value is displayed either in a percentage or MB. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), Percentage (1), MB (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUTOGROWTH_UNIT or ATOGRUNT (historical name), *Autogrowth Unit* (caption), Autogrowth_Unit (attribute name), and ATOGRUNT (column name).

Autogrowth Value

Indicates the growth value of a file. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUTOGROWTH_VALUE or ATOGRVAL (historical name), *Autogrowth Value* (caption), Autogrowth_Value (attribute name), and ATOGRVAL (column name).

Autogrowth

Indicates whether the autogrowth feature is enabled for a device. If the autogrowth feature is enabled, the maximum file size is displayed as restricted or unrestricted. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Disabled (1), Restricted (2), Unrestricted (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AUTOGROWTH or ATOGROWT (historical name), *Autogrowth* (caption), Autogrowth (attribute name), and ATOGROWT (column name).

Database Name (Unicode)

The Database name. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), *Database Name (Unicode)* (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Device Free Space (MB)

The number of megabytes (MB) of free space on a device. The value format is a decimal with two decimal places allowed; for example, 10.00. Adequate space is required to support database and system administration activities. Set alerts for abnormal conditions. Refer to this value when estimating space needed for tables, indexes, logs, and system administration. The type is real number

(32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEVICE_FREE_SPACE or DVFREE (historical name), *Device Free Space (MB)* (caption), *Device_Free_Space* (attribute name), and DVFREE (column name).

Device Free Space Percent

The percentage of free space on the device. The value format is a percentage with two decimal places allowed; for example, 15.00. Adequate space is required to support database and system administration activities. Set alerts for abnormal conditions. Refer to this value when estimating space needed for tables, indexes, logs, and system administration. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEVICE_FREE_SPACE_PERCENT or PCTDVFREE (historical name), *Device Free Space Percent* (caption), *Device_Free_Space_Percent* (attribute name), and PCTDVFREE (column name).

Device Name (Unicode)

The Device name. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: DEVICE_NAME_U or UDEVICE (historical name), *Device Name (Unicode)* (caption), *Device_Name_U* (attribute name), and UDEVICE (column name).

Device Size (MB)

The number of megabytes (MB) on the device. The value format is an integer; for example, 20. Use this value when analyzing the amount of free space. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEVICE_SIZE or DVSIZE (historical name), *Device Size (MB)* (caption), *Device_Size* (attribute name), and DVSIZE (column name).

Device Status

Indicates the status of a device. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), ONLINE (1), RESTORING (2), RECOVERING (3), RECOVERY PENDING (4), SUSPECT (5), INFORMATIONAL (6), OFFLINE (7), DEFUNCT (8), INACCESSIBLE (9), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEVICE_STATUS or DEVSTATUS (historical name), *Device Status* (caption), *Device_Status* (attribute name), and DEVSTATUS (column name).

Device Type

Indicates the type of device allocated for the database. The value indicates the type of data that is stored on the device. Databases are frequently spread across several devices due to size, performance, and recoverability issues. The type is string with enumerated values. The following values are defined: Not collected (-1), Database file (2), Log file (1), FILESTREAM (5), For future use (3), Full text (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEVICE_TYPE or DVTYPE (historical name), *Device Type* (caption), *Device_Type* (attribute name), and DVTYPE (column name).

Mirror Device Name

The name of the mirror device for the database. The value format is an alphanumeric string with a maximum of 64 characters; for example, /dev/rsd2g . The mirror device duplicates the contents of a primary device. Refer to the logical and physical names of the device that is being mirrored. The type is string with enumerated values. The following values are defined: Not Applicable (N/A). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIRROR_DEVICE_NAME or MIRROR (historical name), *Mirror Device Name* (caption), *Mirror_Device_Name* (attribute name), and MIRROR (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Physical Device Name (Unicode)

The Name of the physical device. Valid values include text strings with a maximum of 192 bytes. The type is string.

The following names are defined for this attribute: PHYSICAL_DEVICE_NAME_U or UPHYNAME (historical name), *Physical Device Name (Unicode)* (caption), PHYSICAL_DEVICE_NAME_U (attribute name), and UPHYNAME (column name).

Physical Device Name

The name of the physical device allocated for the database. The value format is an alphanumeric string with a maximum of 30 characters; for example, DATA_1. The type is string.

The following names are defined for this attribute: PHYSICAL_DEVICE_NAME or PHYNAME (historical name), *Physical Device Name* (caption), PHYSICAL_DEVICE_NAME (attribute name), and PHYNAME (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Database Name

The name of the database. The value format is alphanumeric string with a maximum 30 characters; for example, KOQ3. Each database name is unique. The SQL server also assigns each database its own identification number. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), Database Name (caption), DATABASE_NAME (attribute name), and DBNAME (column name).

Device Name

The name of the device allocated for the database. The value format is an alphanumeric string with a maximum of 64 characters; for example, DATA_1. The logical name of the device is stored in the name column of the sysdevices table. Use the device name in storage-management commands. The type is string.

The following names are defined for this attribute: DEVICE_NAME or DEVICE (historical name), Device Name (caption), DEVICE_NAME (attribute name), and DEVICE (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), HOST_NAME (attribute name), and HOSTNAME (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product

collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Server Version

The version of the SQL Server. The value format is the version in the format w. x. y. z; for example, 10. 0. 2531. 0. The type is string.

The following names are defined for this attribute: SQL_SERVER_VER or SQLVER (historical name), Server Version (caption), SQL_Server_Ver (attribute name), and SQLVER (column name).

Server Version (Superceded)

The version of the SQL Server. The value format is the version in the format x. y. z; for example, 10. 0. 253. For the complete version of the SQL Server, see the Server Version attribute. The type is string.

The following names are defined for this attribute: SERVER_VERSION or SRVVERSN (historical name), Server Version (Superceded) (caption), Server_Version (attribute name), and SRVVERSN (column name).

MS SQL Error Event Detail data set

The MS SQL Error Event Details data set provides detailed information about the selected errors in the SQL Server error log. The data is collected at the time that the error occurred. This data set is an event-type data set.

This data set contains the following attributes:

Error ID

The ID of the error message. The value format is an integer. For example, 2520. The errors with text that are not numbered are written to the SQL Server message logs. The type is string.

The following names are defined for this attribute: Error ID (caption), Error_ID (attribute name), and ERRORID (column name).

Error SPID

The ID of the session in which the event has occurred. This attribute is available for reports and queries, but not for eventing thresholds. The type is string.

The following names are defined for this attribute: Error SPID (caption), Error_SPID (attribute name), and ERRSPID (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Message Age (Min.)

The number of minutes that have elapsed since the error occurred. The format is an integer, for example, 2. Use this attribute to help identify the current messages. The type is integer (32-bit numeric property).

The following names are defined for this attribute: Message Age (Min.) (caption), Message_Age (attribute name), and MSGAGE (column name).

Message Issuer

The source of the error message. The value format is an alphanumeric string with a maximum of 8 characters, for example, JSMITH. The type is string.

The following names are defined for this attribute: Message Issuer (caption), Message_Issuer (attribute name), and MSGISSUER (column name).

Message Text (Unicode)

The message text. This attribute is available for reports and queries, but not for eventing thresholds. The type is string.

The following names are defined for this attribute: Message Text (Unicode) (caption), Message_Text_U (attribute name), and UMSGTEXT (column name).

Message Timestamp

The timestamp that indicates the date and time at which the error occurred. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the agent collected data from the SQL Server on January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Message Timestamp (caption), Message_Timestamp (attribute name), and MSGTIME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time at which the agent collected sample data for the SQL Server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the agent collected the data from the SQL Server on January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value format is an alphanumeric string with a maximum of 32 characters. For example, CFS_SVR. The type is string.

The following names are defined for this attribute: Server (caption), Server (attribute name), and SERVERID (column name).

Severity Level

The severity level of the error. The type is integer (32-bit numeric property).

The following names are defined for this attribute: Severity Level (caption), Severity_Level (attribute name), and SEVERITY (column name).

SQL State Code

The value that represents state of the SQL Server when the error occurred. The value format is an integer. For example, 37. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SQL State Code (caption), SQL_State_Code (attribute name), and SQLSTATE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Expensive Query Plans data set

The MS SQL Expensive Query Plans data set provides the top 10 cached query plans according to the performance statistics in the SQL server. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Average CPU Time(ms)

The average CPU time (in milliseconds) that is required to execute the plan since the plan was compiled. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_CPU_TIME or AVGCPUTIME (historical name), *Average CPU Time(ms)* (caption), Average_CPU_Time (attribute name), and AVGCPUTIME (column name).

Average Duration(ms)

The average elapsed time (in milliseconds) that is required to execute this plan. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_DURATION or AVGDUR (historical name), *Average Duration(ms)* (caption), Average_Duration (attribute name), and AVGDUR (column name).

Average Logical Reads

The average logical reads that are performed to execute this plan since the plan was compiled. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_LOGICAL_READS or AVGLOGREDS (historical name), *Average Logical Reads* (caption), Average_Logical_Reads (attribute name), and AVGLOGREDS (column name).

Average Logical Writes

The average logical writes that are performed to execute this plan since the plan was compiled. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_LOGICAL_WRITES or AVGLOGWRTS (historical name), *Average Logical Writes* (caption), Average_Logical_Writes (attribute name), and AVGLOGWRTS (column name).

Average Physical Reads

The average physical reads that are performed to execute this plan since the plan was compiled. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_PHYSICAL_READS or AVGPHYREDS (historical name), *Average Physical Reads* (caption), Average_Physical_Reads (attribute name), and AVGPHYREDS (column name).

Database Name (Unicode)

The name of the database against which the query is executed. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or UDBNAME (historical name), *Database Name (Unicode)* (caption), Database_Name (attribute name), and UDBNAME (column name).

Execution rate

The value in percentage that is calculated by dividing the execution count by the value that is obtained after subtracting the creation time of the query from the last execution time of the query. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EXECUTION_RATE or EXERATE (historical name), *Execution rate* (caption), Execution_rate (attribute name), and EXERATE (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Query Age in Minutes

The time interval (in minutes) between the last time that the query was executed and the time when the query was created. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGE or QUERYAGE (historical name), *Query Age in Minutes* (caption), Age (attribute name), and QUERYAGE (column name).

Query State

The state of the query that is based on the threshold (Normal/Critical) The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), Normal (1), Critical (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_STATE or QUERYSTATE (historical name), *Query State* (caption), Query_State (attribute name), and QUERYSTATE (column name).

Query Text (Unicode)

The SQL text of the expensive query plan. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_TEXT or USQLTEXT (historical name), *Query Text (Unicode)* (caption), *Query_Text* (attribute name), and USQLTEXT (column name).

Sorting Criteria

The sorting criteria for top N expensive queries. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Average Physical Reads (1), Average Logical Reads (2), Average Logical Writes (3), Average Duration (4), Average CPU Time (5), Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SORTING_CRITERIA or SORTBY (historical name), *Sorting Criteria* (caption), *Sorting_Criteria* (attribute name), and SORTBY (column name).

SQL Handle

The SQL handle of the SQL query that is executed on the SQL server. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected), Not Applicable (Not_applicable), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_HANDLE or SQLHANDLE (historical name), *SQL Handle* (caption), *SQL_Handle* (attribute name), and SQLHANDLE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

CPU Rate (percentage)

The value of CPU time in percentage (per interval) that is required to execute this plan since the last collection. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_RATE or CPURATE (historical name), CPU Rate (percentage) (caption), *CPU_Rate* (attribute name), and CPURATE (column name).

Executions Per Interval

The number of times that the query is executed in the sample interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EXECUTIONS_PER_INTERVAL or EXECPERINT (historical name), Executions Per Interval (caption), *Executions_Per_Interval* (attribute name), and EXECPERINT (column name).

ROWNO

The row number in the sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), *ROWNO* (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), *SAMPLENO* (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Elapsed Time(ms)

The total elapsed time (in milliseconds) that is required to execute this plan. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ELAPSED_TIME or TELPTIME (historical name), Total Elapsed Time(ms) (caption), Total_Elapsed_Time (attribute name), and TELPTIME (column name).

MS SQL Filegroup Detail data set

The MS SQL Filegroup Detail data set provides details about the filegroups for each database in the SQL Server instance. The data collection for this data set is limited to 10000 rows.

This data set contains the following attributes:

Database Name (Unicode)

The name of the database. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or UDBNAME (historical name), Database Name (Unicode) (caption), Database_Name (attribute name), and UDBNAME (column name).

Filegroup ID

The filegroup identification number. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILEGROUP_ID or FGID (historical name), Filegroup ID (caption), Filegroup_ID (attribute name), and FGID (column name).

Filegroup Max Growth Size

The maximum amount of space (in MB) that the filegroup can grow to. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILEGROUP_MAX_GROWTH_SIZE or FGMAXSIZE (historical name), Filegroup Max Growth Size (caption), Filegroup_Max_Growth_Size (attribute name), and FGMAXSIZE (column name).

Filegroup Name

The name of the filegroup. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILEGROUP_NAME or FGNAME (historical name), Filegroup Name (caption), Filegroup_Name (attribute name), and FGNAME (column name).

Filegroup Size

The size of the filegroup in MB. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILEGROUP_SIZE or FGSIZE (historical name), Filegroup Size (caption), Filegroup_Size (attribute name), and FGSIZE (column name).

Filegroup Space Pct Used

The amount of space (in MB) that is used in the filegroup as a percentage of total space allowed. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILEGROUP_SPACE_PERCENT_USED or FGPCTUSE (historical name), Filegroup Space Pct Used (caption), Filegroup_Space_Percent_Used (attribute name), and FGPCTUSE (column name).

Filegroup Status

Indicates the status of a filegroup. If the status of a file in a filegroup is offline, the status of the filegroup is displayed as offline. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), ONLINE (1), OFFLINE (2), INACCESSIBLE (3), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILEGROUP_STATUS or FGPSTATUS (historical name), Filegroup Status (caption), Filegroup_Status (attribute name), and FGPSTATUS (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:ss, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL FileTable Detail data set

The MS SQL FileTable Detail data set contains attributes that provide information about each FileTable in the SQL Server. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Database Id

The database ID. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database Id (caption), Database_Id (attribute name), and DBID (column name).

FileTable Directory

The name of the directory of the FileTable. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILETABLE_DIRECTORY or FILTBLDIR (historical name), FileTable Directory (caption), FileTable_Directory (attribute name), and FILTBLDIR (column name).

FileTable Enabled

Indicates whether the FileTable data is available for non-transactional access. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1), No (1), Yes (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILETABLE_ENABLED or FILTBLENB (historical name), FileTable Enabled (caption), FileTable_Enabled (attribute name), and FILTBLENB (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Table Id

The table ID. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_ID or TBLID (historical name), Table Id (caption), Table_Id (attribute name), and TBLID (column name).

Table Name

The table name. The value format is an alphanumeric string with a maximum of 30 characters. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_NAME or UTBLNAME (historical name), Table Name (caption), Table_Name (attribute name), and UTBLNAME (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Non Transactional Handles Opened

The total number of non-transactional file handles that are currently open and that are associated with the FileTable. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

TOTAL_NON_TRANSACTIONAL_HANDLES_OPENED or TNTRNHNOP (historical name), Total Non Transactional Handles Opened (caption), Total_Non_Transactional_Handles_Opened (attribute name), and TNTRNHNOP (column name).

MS SQL Individual Queries Details data set

The MS SQL Individual Queries Details data set contains attributes that provide information about each query in the query batch that is currently running on the SQL Server.

This data set contains the following attributes:

Avg CPU Time (Sec.)

The average CPU time (in seconds) that is taken to run an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVG_CPU_TIME_SEC or AVGCPUTM (historical name), Avg CPU Time (Sec.) (caption), Avg_CPU_Time_Sec (attribute name), and AVGCPUTM (column name).

Avg Logical Reads

The average number of logical read operations that are completed by an SQL query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVG_LOGICAL_READS or AVGLOGREDS (historical name), Avg Logical Reads (caption), Avg_Logical_Reads (attribute name), and AVGLOGREDS (column name).

Avg Logical Writes

The average number of logical write operations that are completed by an SQL query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVG_LOGICAL_WRITES or AVGLOGWRTS (historical name), Avg Logical Writes (caption), Avg_Logical_Writes (attribute name), and AVGLOGWRTS (column name).

Avg Physical Reads

The average number of physical read operations that are completed by an SQL query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVG_PHYSICAL_READS or AVGPHYREDS (historical name), Avg Physical Reads (caption), Avg_Physical_Reads (attribute name), and AVGPHYREDS (column name).

Database Name (Unicode)

The name of the database against which the SQL query is running on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or DBSNAME (historical name), Database Name (Unicode) (caption), Database_Name_U (attribute name), and DBSNAME (column name).

Last Execution Time

The time when an SQL query was last run on the SQL Server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL Server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_EXECUTION_TIME or LASTEXECTM (historical name), Last Execution Time (caption), Last_Execution_Time (attribute name), and LASTEXECTM (column name).

Max CLR Time(Sec.)

The maximum common language runtime (CLR, in seconds) that is used for a single execution of an SQL query inside the .NET framework CLR. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_CLR_TIME_SEC or MAXCLRTIME (historical name), Max CLR Time(Sec.) (caption), Max_CLR_Time_Sec (attribute name), and MAXCLRTIME (column name).

Max CPU Time(Sec.)

The maximum CPU time (in seconds) that is used for a single execution of an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_CPU_TIME_SEC or MAXCPUTIME (historical name), Max CPU Time(Sec.) (caption), Max_CPU_Time_Sec (attribute name), and MAXCPUTIME (column name).

Max logical Reads

The maximum number of logical read operations that are completed by a single execution of an SQL query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_LOGICAL_READS or MAXLOGREDS (historical name), Max logical Reads (caption), Max_logical_Reads (attribute name), and MAXLOGREDS (column name).

Max logical Writes

The maximum number of logical write operations that are completed by a single execution of an SQL query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_LOGICAL_WRITES or MAXLOGWRTS (historical name), Max logical Writes (caption), Max_logical_Writes (attribute name), and MAXLOGWRTS (column name).

Max Physical Reads

The maximum number of physical read operations that are completed by a single execution of an SQL query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_PHYSICAL_READS or MAXPHYREDS (historical name), Max Physical Reads (caption), Max_Physical_Reads (attribute name), and MAXPHYREDS (column name).

Max Waiting Time(Sec.)

The maximum waiting time (in seconds) that is taken by a single execution of an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_WAITING_TIME_SEC or MAXWAITTIM (historical name), Max Waiting Time(Sec.) (caption), Max_Waiting_Time_Sec (attribute name), and MAXWAITTIM (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Query Execution Count

The total number of times an SQL query has been run since the query was last compiled. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_EXECUTION_COUNT or QEXECCOUNT (historical name), Query Execution Count (caption), Query_Execution_Count (attribute name), and QEXECCOUNT (column name).

Query Hash

The query hash of an SQL query. The value format is an alphanumeric string with a maximum of 2048 characters. This attribute, with the Session ID attribute, is used to create the primary key. This attribute is a key attribute. The type is integer (64-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_HASH or QUERYHASH (historical name), Query Hash (caption), Query_Hash (attribute name), and QUERYHASH (column name).

Query Text

The SQL text of a query in an SQL query batch. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_TEXT or CHILDQTEXT (historical name), Query Text (caption), Query_Text (attribute name), and CHILDQTEXT (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Session ID

The ID of the session in which the query is run. This attribute is a key attribute. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SESSION_ID or SESSID (historical name), Session ID (caption), Session_ID (attribute name), and SESSID (column name).

SQL Handle

The SQL handle of the query. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_HANDLE or SQLHANDLE (historical name), SQL Handle (caption), SQL_Handle (attribute name), and SQLHANDLE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total CPU Time (Sec.)

The total amount of CPU time (in seconds) that is used by all the executions of an SQL query. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_CPU_TIME_SEC or TOTCPUUTM (historical name), Total CPU Time (Sec.) (caption), Total_CPU_Time_Sec (attribute name), and TOTCPUUTM (column name).

Total Percent CPU Used

The percentage of the available CPU memory that is used for executing an SQL query. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PERCENT_CPU_USED or TOTPCPUUSD (historical name), Total Percent CPU Used (caption), Total_Percent_CPU_Used (attribute name), and TOTPCPUUSD (column name).

Total Percent Waiting

The percentage of total wait operations that are completed when an SQL query is run. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PERCENT_WAITING or PERWAITING (historical name), Total Percent Waiting (caption), Total_Percent_Waiting (attribute name), and PERWAITING (column name).

MS SQL Integration Service Details data set

The MS SQL Integration Service Details data set provides details about the SQL Server integration service pipeline for an instance of the SQL Server. This data set is not supported in MSSQL Linux agent.

This data set contains the following attributes:

Buffer In Use

The number of buffers that are currently used for the pipeline. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BUFFER_IN_USE or BUFINUS (historical name), Buffer In Use (caption), Buffer_In_Use (attribute name), and BUFINUS (column name).

Buffer Memory (MB)

The amount of memory in MB that is currently allocated to the buffers that are in the pipeline. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BUFFER_MEMORY or BUFMEM (historical name), Buffer Memory (MB) (caption), Buffer_Memory (attribute name), and BUFMEM (column name).

Buffer Spooled

The number of buffers that are currently spooled to disk. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BUFFER_SPOOLED or BUFSPOOLD (historical name), Buffer Spooled (caption), Buffer_Spooled (attribute name), and BUFSPOOLD (column name).

Flat Buffer Memory (MB)

The amount of memory that is currently allocated to flat memory buffers. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FLAT_BUFFER_MEMORY or FLTBUFMEM (historical name), Flat Buffer Memory (MB) (caption), Flat_Buffer_Memory (attribute name), and FLTBUFMEM (column name).

Flat Buffers In Use

The number of flat memory buffers that are currently used for the pipeline. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FLAT_BUFFERS_IN_USE or FLTBUFINUS (historical name), Flat Buffers In Use (caption), Flat_Buffers_In_Use (attribute name), and FLTBUFINUS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Private Buffer Memory (MB)

The amount of memory that is currently allocated to private transformation buffers. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PRIVATE_BUFFER_MEMORY or PRBUFMEM (historical name), Private Buffer Memory (MB) (caption), Private_Buffer_Memory (attribute name), and PRBUFMEM (column name).

Private Buffers In Use

The number of private transformation buffers that are currently used for the pipeline. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PRIVATE_BUFFERS_IN_USE or PRBUFINUS (historical name), Private Buffers In Use (caption), Private_Buffers_In_Use (attribute name), and PRBUFINUS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Rows Read

The total number of rows that are currently being read from all data sources. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS_READ or TOTROWRD (historical name), Total Rows Read (caption), Total_Rows_Read (attribute name), and TOTROWRD (column name).

Total Rows Written

The total number of rows that are currently being written to all data destinations. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ROWS_WRITTEN or TOTROWWR (historical name), Total Rows Written (caption), Total_Rows_Written (attribute name), and TOTROWWR (column name).

MS SQL Job Detail data set

The MS SQL Job Detail data set contains attributes that you can use to monitor SQL Server jobs. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Current Execution Step

The current step that is being executed in the job. The value format is an alphanumeric string with a maximum of 128 characters. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_EXECUTION_STEP or CUEXESTP (historical name), *Current Execution Step* (caption), Current_Execution_Step (attribute name), and CUEXESTP (column name).

Current Retry Attempt

The current number of retry attempts that are done on a step of a running job. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_RETRY_ATTEMPT or CURRETAT (historical name), *Current Retry Attempt* (caption), Current_Retry_Attempt (attribute name), and CURRETAT (column name).

Current Status

The current job status. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Executing (1), Waiting for thread (2), Between retries (3), Idle (4), Suspended (5), Performing completion actions (7), Unknown (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_STATUS or CURSTATUS (historical name), *Current Status* (caption), Current_Status (attribute name), and CURSTATUS (column name).

Enabled

Whether or not the job is enabled to run. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: No (0), Yes (1), Unknown (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENABLED or JOBENABLED (historical name), *Enabled* (caption), Enabled (attribute name), and JOBENABLED (column name).

Job Duration

The amount of time it took for the job to complete (in seconds). The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_DURATION or JOBDUR (historical name), *Job Duration* (caption), Job_Duration (attribute name), and JOBDUR (column name).

Job Error Code

Error code for the last completion of the job. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_ERROR_CODE or JOBERRCODE (historical name), *Job Error Code* (caption), Job_Error_Code (attribute name), and JOBERRCODE (column name).

Job Execution Duration

The time (in seconds) that has elapsed since the job started. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_EXECUTION_DURATION or JOBEXEDUR (historical name), *Job Execution Duration* (caption), Job_Execution_Duration (attribute name), and JOBEXEDUR (column name).

Job Failure Count

The number of unsuccessful executions of the job. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_FAILURE_COUNT or JOBAFAILCNT (historical name), *Job Failure Count* (caption), *Job_Failure_Count* (attribute name), and JOBAFAILCNT (column name).

Job Id

The SQL Server job ID. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: JOB_ID or JOBID (historical name), *Job Id* (caption), *Job_Id* (attribute name), and JOBID (column name).

Job Name

The SQL Server job name. The type is string.

The following names are defined for this attribute: JOB_NAME or JOBNAME (historical name), *Job Name* (caption), *Job_Name* (attribute name), and JOBNAME (column name).

Job Owner

The name of the owner of the job. The value format is an alphanumeric string with a maximum of 128 characters. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_OWNER or JOBOWNER (historical name), *Job Owner* (caption), *Job_Owner* (attribute name), and JOBOWNER (column name).

Job Status

Current status of the job. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Failed (0), Succeeded (1), Retry (2), Canceled (3), In-progress (4), Unknown (5), Unknown (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_STATUS or JOBSTATNUM (historical name), *Job Status* (caption), *Job_Status* (attribute name), and JOBSTATNUM (column name).

Job Success Count

The number of successful executions of the job. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_SUCCESS_COUNT or JOBSUCCNT (historical name), *Job Success Count* (caption), *Job_Success_Count* (attribute name), and JOBSUCCNT (column name).

Last Run Outcome

The last job execution status. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Failed (0), Succeeded (1), Canceled (3), Unknown (5), Unknown (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_RUN_OUTCOME or LSTRUNSTAT (historical name), *Last Run Outcome* (caption), *Last_Run_Outcome* (attribute name), and LSTRUNSTAT (column name).

Last Run Timestamp

The timestamp of last job execution. If the job has not been run at all, the Last Run Timestamp value is 'N/P' (Not Present). If there is an error retrieving the job information, the Last Run Timestamp value is 'N/P'. The type is timestamp with enumerated values. The following values are defined: Invalid (0000000000000000), N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_RUN_TIMESTAMP or LSTRUNTIME (historical name), *Last Run Timestamp* (caption), *Last_Run_Timestamp* (attribute name), and LSTRUNTIME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Current Execution Status

The current execution status of the job. The type is integer (32-bit gauge) with enumerated values.

The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_EXECUTION_STATUS or CUEXISTA (historical name), Current Execution Status (caption), Current_Execution_Status (attribute name), and CUEXISTA (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Job Category ID

Job category ID for the job. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_CATEGORY_ID or JOBCATID (historical name), Job Category ID (caption), Job_Category_ID (attribute name), and JOBCATID (column name).

Job Category Name

Job category name for the job. The type is string.

The following names are defined for this attribute: JOB_CATEGORY_NAME or JOBCATNAME (historical name), Job Category Name (caption), Job_Category_Name (attribute name), and JOBCATNAME (column name).

Job Description

The description of the job. The value format is an alphanumeric string with a maximum of 512 characters. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_DESCRIPTION or JOBDESC (historical name), Job Description (caption), Job_Description (attribute name), and JOBDESC (column name).

Job Type

The type of the job. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Local Job (1), Multiserver Job (2), No Target Server (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_TYPE or JOBTYPE (historical name), Job Type (caption), Job_Type (attribute name), and JOBTYPE (column name).

Next Run Timestamp

The timestamp of the next scheduled job execution. If a job is not scheduled to run, its Next Run Timestamp is 'N/C' (Not Configured). If there is an error retrieving the job information, the Next Run Timestamp values is 'N/P'(Not Present). The type is timestamp with enumerated values. The following values are defined: N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NEXT_RUN_TIMESTAMP or NXTRUNTIME (historical name), Next Run Timestamp (caption), Next_Run_Timestamp (attribute name), and NXTRUNTIME (column name).

Notify Level Eventlog

The value that indicates when a notification event must be sent to the Microsoft Windows application log. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Never (0), When a job succeeds (1), When the job fails (2), Whenever the job completes (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NOTIFY_LEVEL_EVENTLOG or NFLVEVL (historical name), Notify Level Eventlog (caption), Notify_Level_Eventlog (attribute name), and NFLVEVL (column name).

Number of Steps

The current number of steps in the job. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NUMBER_OF_STEPS or NUMSTEPS (historical name), Number of Steps (caption), Number_of_Steps (attribute name), and NUMSTEPS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:ss, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Start Step ID

The step identifier from where the execution of the job starts. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: START_STEP_ID or STEPID (historical name), Start Step ID (caption), Start_Step_ID (attribute name), and STEPID (column name).

Version number

The version number of the job. The version number is automatically updated every time that the job is modified. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_NUMBER or VERNUM (historical name), Version number (caption), Version_number (attribute name), and VERNUM (column name).

MS SQL Job Summary data set

The MS SQL Job Summary data set contains attributes that you can use to monitor jobs. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Active Jobs

The number of active jobs since the agent was started. The format is an integer. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTIVE_JOBS or ACTJOBS (historical name), *Active Jobs* (caption), Active_Jobs (attribute name), and ACTJOBS (column name).

Failed Jobs

The number of failed jobs since the SQL Server agent was started. The format is an integer. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAILED_JOBS or FAILJOBS (historical name), *Failed Jobs* (caption), Failed_Jobs (attribute name), and FAILJOBS (column name).

Job Success Rate

The rate at which jobs are successful since the agent was started. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_SUCCESS_RATE or JOBSUCRATE (historical name), *Job Success Rate* (caption), Job_Success_Rate (attribute name), and JOBSUCRATE (column name).

Jobs Executed Current Interval

The number of processed jobs in the current sampling interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOBS_EXECUTED_CURRENT_INTERVAL or JOBEXECUR (historical name), *Jobs Executed Current Interval* (caption), Jobs_Executed_Current_Interval (attribute name), and JOBEXECUR (column name).

Jobs Failed Current Interval

The number of jobs that failed in the current sample interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOBS_FAILED_CURRENT_INTERVAL or NUMJBADCUR (historical name), *Jobs Failed Current Interval* (caption), Jobs_Failed_Current_Interval (attribute name), and NUMJBADCUR (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Queued Jobs

The number of jobs in a queue since the agent was started. The format is an integer. This attribute is not supported in MSSQL Linux agent. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUEUED_JOBS or QJOBS (historical name), *Queued Jobs* (caption), Queued_Jobs (attribute name), and QJOBS (column name).

Successful Jobs

The number of successful jobs since the agent was started. The format is an integer. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SUCCESSFUL_JOBS or SUCCJOBS (historical name), *Successful Jobs* (caption), Successful_Jobs (attribute name), and SUCCJOBS (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Warning Jobs Current Interval

The number of processed jobs in the current sampling interval with the status as unknown or canceled. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: WARNING_JOBS_CURRENT_INTERVAL or WARNJOBCUR (historical name), *Warning Jobs Current Interval* (caption), Warning_Jobs_Current_Interval (attribute name), and WARNJOBCUR (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), Host_Name (attribute name), and HOSTNAME (column name).

Jobs Failed Since Startup

The number of failed jobs since the agent was started. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOBS_FAILED_SINCE_STARTUP or NUMJBADTOT (historical name), *Jobs Failed Since Startup* (caption), Jobs_Failed_Since_Startup (attribute name), and NUMJBADTOT (column name).

Max Jobs Failed Interval

Maximum number of failed jobs across all samples since the agent was started. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_JOBS_FAILED_INTERVAL or MAXJBADCUR (historical name), *Max Jobs Failed Interval* (caption), Max_Jobs_Failed_Interval (attribute name), and MAXJBADCUR (column name).

Max Running Jobs Interval

Maximum number of running jobs across all samples since the agent was started. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_RUNNING_JOBS_INTERVAL or MAXJRUNCUR (historical name), Max Running Jobs Interval (caption), Max_Running_Jobs_Interval (attribute name), and MAXJRUNCUR (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Running Jobs Current Interval

The number of running jobs for the current sample interval. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RUNNING_JOBS_CURRENT_INTERVAL or NUMJRUNCUR (historical name), Running Jobs Current Interval (caption), Running_Jobs_Current_Interval (attribute name), and NUMJRUNCUR (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Lock Conflict Detail data set

The MS SQL Lock Conflict Detail data set contains attributes that you can use to monitor detailed information about a selected lock conflict. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Blocking Process ID

The identifier for the process that is blocking a request for a lock. The value format is an alphanumeric string with a maximum of 8 characters; for example, 12. Some blocking processes can become stranded. Investigate eventing thresholds in which a process is being blocked for an extended period

of time. For more information on a blocking process, query the sysprocesses and use the sp-lock procedure. The type is integer (32-bit numeric property).

The following names are defined for this attribute: BLOCKING_PROCESS_ID or BLOCKID (historical name), *Blocking Process ID* (caption), Blocking_Process_ID (attribute name), and BLOCKID (column name).

Client Group ID

The group ID of the user executing the process. This value does not apply to Microsoft SQL Server version 7. The value format is an alphanumeric string with a maximum of 8 characters; for example, ACCT_1. The type is string.

The following names are defined for this attribute: CLIENT_GROUP_ID or SRVGROUP (historical name), *Client Group ID* (caption), Client_Group_ID (attribute name), and SRVGROUP (column name).

Client User ID

ID of the user who executed the command. The value format is an alphanumeric string with a maximum of 8 characters. The type is string.

The following names are defined for this attribute: CLIENT_USER_ID or CLNTUSER (historical name), *Client User ID* (caption), Client_User_ID (attribute name), and CLNTUSER (column name).

Database Name (Unicode)

The name of the database. The type is string.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), *Database Name (Unicode)* (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Requestor Process ID

The ID of the blocked process. The value format is an alphanumeric string with a maximum of 16 characters; for example, 21. The type is integer (32-bit numeric property).

The following names are defined for this attribute: REQUESTOR_PROCESS_ID or PROCESSID (historical name), *Requestor Process ID* (caption), Requestor_Process_ID (attribute name), and PROCESSID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Database Id

ID of the database that is locked. The value format is an integer. The type is integer (32-bit numeric property).

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database_Id (caption), Database_Id (attribute name), and DBID (column name).

Database Name

The name of the database. The value format is an alphanumeric string with a maximum of 30 characters; for example, KOQ3. Each database name is unique. The SQL server also assigns each database its own identification number. The type is string.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), Database Name (caption), Database_Name (attribute name), and DBNAME (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The number of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Server User ID

The ID of the server user. The value format is an alphanumeric string with a maximum of 8 characters. The type is string.

The following names are defined for this attribute: SERVER_USER_ID or SRVUSER (historical name), Server User ID (caption), Server_User_ID (attribute name), and SRVUSER (column name).

MS SQL Lock Detail data set

The MS SQL Lock Detail data set contains attributes that you can use to monitor detailed information about lock contention by lock type. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Database Id

The ID of the database that is locked. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database Id (caption), Database_Id (attribute name), and DBID (column name).

Database Name (Unicode)

The name of the locked database. The type is string.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), *Database Name (Unicode)* (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Lock Request Status

Enumeration that identifies the Lock Request Status. The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Granted (1), Converting (2), Waiting (3), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_REQUEST_STATUS_E or LCKRQSTATE (historical name), *Lock Request Status* (caption), Lock_Request_Status_E (attribute name), and LCKRQSTATE (column name).

Lock Resource Type

Enumeration that identifies the Lock Resource Type. The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: NULL Resource (1), Database (2), File (3), Index (4), Table (5), Page (6), Key (7), Extent (8), RID (9), Application (10), Metadata (11), Heap or Btree (12), Allocation Unit (13), Object (14). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_RESOURCE_TYPE_E or LCKRSCTYPE (historical name), *Lock Resource Type* (caption), Lock_Resource_Type_E (attribute name), and LCKRSCTYPE (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Page Number

The page number of the table being locked. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_NUMBER or PAGE (historical name), *Page Number* (caption), Page_Number (attribute name), and PAGE (column name).

Process Holding Lock

The ID of the process holding the lock. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCESS_HOLDING_LOCK or PIDHOLD (historical name), *Process Holding Lock* (caption), Process_Holding_Lock (attribute name), and PIDHOLD (column name).

Table Name

The name of the table being locked. The value format is an alphanumeric string with a maximum of 16 characters. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_NAME or TBNAME (historical name), *Table Name* (caption), Table_Name (attribute name), and TBNAME (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Type

Indicates the type of lock on the resource that is being requested. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TYPE or LOCKTYPE (historical name), *Type* (caption), Type (attribute name), and LOCKTYPE (column name).

Database Name

The name of the database that is locked. The value format is an alphanumeric string with a maximum of 16 characters; for example, DB12. The type is string.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), Database Name (caption), Database_Name (attribute name), and DBNAME (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Table Name (Unicode)

Name of Table being locked. The type is string.

The following names are defined for this attribute: TABLE_NAME_U or UTBNAME (historical name), Table Name (Unicode) (caption), Table_Name_U (attribute name), and UTBNAME (column name).

MS SQL Lock Resource Type Summary data set

The MS SQL Lock Resource Type Summary data set provides information about locks on individual resource types. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute

and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Average Wait Time locks (ms)

The average amount of wait time (in milliseconds) for each lock request that resulted in a wait. The type is integer (32-bit counter).

The following names are defined for this attribute: AVERAGE_WAIT_TIME_LOCKS or AVGWAITTM (historical name), *Average Wait Time locks (ms)* (caption), Average_Wait_Time_locks (attribute name), and AVGWAITTM (column name).

Lock Requests per Sec.

The number of new locks and lock conversions per second requested from the lock manager. This value is calculated on an interval, and does not reflect a total value. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: LOCK_REQUESTS_PER_SECOND or LCKREQSEC (historical name), *Lock Requests per Sec.* (caption), Lock_Requests_per_Second (attribute name), and LCKREQSEC (column name).

Lock Resource Type

An enumeration of the resources that the SQL Server can lock. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: NULL Resource (1), Database (2), File (3), Index (4), Table (5), Page (6), Key (7), Extent (8), RID (9), Application (10), Metadata (11), Heap or Btree (12), Allocation Unit (13), Object (14), OibTrackTbl (15), RowGroup (16), UNKNOWN LOCK RESOURCE (17), UNUSED1 (18), LAST RESOURCE (19). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_RESOURCE_TYPE or LOCKRSCTYP (historical name), *Lock Resource Type* (caption), Lock_Resource_Type (attribute name), and LOCKRSCTYP (column name).

Lock Timeouts per Sec.

The number of lock requests per second that timed out, including requests for NOWAIT locks. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: LOCK_TIMEOUTS_PER_SECOND or LCKTOSEC (historical name), *Lock Timeouts per Sec.* (caption), Lock_Timeouts_per_Second (attribute name), and LCKTOSEC (column name).

Lock Wait Time (ms)

The total wait time (in milliseconds) for locks in the last second. The type is integer (32-bit gauge).

The following names are defined for this attribute: LOCK_WAIT_TIME or LCKWAITTM (historical name), *Lock Wait Time (ms)* (caption), Lock_Wait_Time (attribute name), and LCKWAITTM (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number Deadlocks per Second

The number of Deadlocks per second for the current sample interval. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: NUMBER_DEADLOCKS_PER_SECOND or LCKDDLKSEC (historical name), *Number Deadlocks per Second* (caption), Number_Deadlocks_per_Second (attribute name), and LCKDDLKSEC (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: **TIMESTAMP** (historical name), **Timestamp** (caption), **Timestamp** (attribute name), and **TIMESTAMP** (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: **HOST_NAME** or **HOSTNAME** (historical name), **Host Name** (caption), **Host_Name** (attribute name), and **HOSTNAME** (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: **SAMPLE_TIMESTAMP** or **SAMPTIME** (historical name), **Sample Timestamp** (caption), **Sample_Timestamp** (attribute name), and **SAMPTIME** (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: **SERVER** or **SERVERID** (historical name), **Server** (caption), **Server** (attribute name), and **SERVERID** (column name).

MS SQL Lock Summary data set

The MS SQL Lock Summary data set contains attributes that you can use to monitor the number of database locks that are currently open. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Exclusive Locks

The number of granted Exclusive locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: **EXCLUSIVE_LOCKS** or **EXCLOK** (historical name), **Exclusive Locks** (caption), **Exclusive_Locks** (attribute name), and **EXCLOK** (column name).

Intent Locks

The number of granted Intent locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: **INTENT_LOCKS** or **TOTINTLOK** (historical name), **Intent Locks** (caption), **Intent_Locks** (attribute name), and **TOTINTLOK** (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: **NODE** (historical name), **Node** (caption), **ORIGINNODE** (attribute name), and **ORIGINNODE** (column name).

Page Locks

The number of granted Page locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: PAGE_LOCKS or TOTPAGLOK (historical name), *Page Locks* (caption), *Page_Locks* (attribute name), and TOTPAGLOK (column name).

Shared Locks

The number of granted Shared locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: SHARED_LOCKS or SHRLOK (historical name), *Shared Locks* (caption), *Shared_Locks* (attribute name), and SHRLOK (column name).

Table Locks

The number of granted Table locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: TABLE_LOCKS or TOTBLLOK (historical name), *Table Locks* (caption), *Table_Locks* (attribute name), and TOTBLLOK (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Exclusive Extent Locks

The number of granted Exclusive Extent locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: EXCLUSIVE_EXTENT_LOCKS or EXCEXTLOK (historical name), *Exclusive Extent Locks* (caption), *Exclusive_Extent_Locks* (attribute name), and EXCEXTLOK (column name).

Exclusive Intent Locks

The number of granted Exclusive Intent locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: EXCLUSIVE_INTENT_LOCKS or EXCINTLOK (historical name), *Exclusive Intent Locks* (caption), *Exclusive_Intent_Locks* (attribute name), and EXCINTLOK (column name).

Exclusive Page Locks

The number of granted Exclusive Page locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: EXCLUSIVE_PAGE_LOCKS or EXCPAGLOK (historical name), *Exclusive Page Locks* (caption), *Exclusive_Page_Locks* (attribute name), and EXCPAGLOK (column name).

Exclusive Table Locks

The number of granted Exclusive Table locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: EXCLUSIVE_TABLE_LOCKS or EXCTBLLOK (historical name), *Exclusive Table Locks* (caption), *Exclusive_Table_Locks* (attribute name), and EXCTBLLOK (column name).

Extent Locks

The number of granted Extent locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: EXTENT_LOCKS or TOTEXTLOK (historical name), *Extent Locks* (caption), *Extent_Locks* (attribute name), and TOTEXTLOK (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), *Host_Name* (attribute name), and HOSTNAME (column name).

Pct Max Locks

The percentage of locks on resources of the maximum number of locks allowed by the SQL server. The value format is a percentage with two decimal places allowed; for example, 10.00. Setting a higher limit for the maximum number of locks does not impair performance. If your operations exceed the

number of available locks, you can increase this limit. This attribute is not available for use in eventing thresholds or for historical data collection. To create a eventing threshold, use the MS SQL Server Enterprise View data set. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERCENT_MAX_LOCKS or PCTLOCKS (historical name), Pct Max Locks (caption), Percent_Max_Locks (attribute name), and PCTLOCKS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Shared Intent Locks

The number of granted Shared Intent locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: SHARED_INTENT_LOCKS or SHRINTLOK (historical name), Shared Intent Locks (caption), Shared_Intent_Locks (attribute name), and SHRINTLOK (column name).

Shared Page Locks

The number of granted Shared Page locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: SHARED_PAGE_LOCKS or SHRPAGLOK (historical name), Shared Page Locks (caption), Shared_Page_Locks (attribute name), and SHRPAGLOK (column name).

Shared Table Locks

The number of granted Shared Table locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: SHARED_TABLE_LOCKS or SHRTBLLOK (historical name), Shared Table Locks (caption), Shared_Table_Locks (attribute name), and SHRTBLLOK (column name).

Table Lock Escalations per Sec

The number of times that the locks on a table were escalated. This attribute is not supported in MSSQL Linux agent. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_LOCK_ESCALATIONS_PER_SEC or TLCKESCSEC (historical name), Table Lock Escalations per Sec (caption), Table_Lock_Escalations_per_Sec (attribute name), and TLCKESCSEC (column name).

Total Lock Conflicts

The total number of processes involved in lock conflicts. The value format is an integer. This attribute is not available for use in eventing thresholds or for historical data collection. To create a eventing threshold, use the MS SQL Server Enterprise View data set. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_LOCK_CONFLICTS or NUMCONFL (historical name), Total Lock Conflicts (caption), Total_Lock_Conflicts (attribute name), and NUMCONFL (column name).

Update Extent Locks

The number of granted Update Extent locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: UPDATE_EXTENT_LOCKS or UPDEXTLOK (historical name), Update Extent Locks (caption), Update_Extent_Locks (attribute name), and UPDEXTLOK (column name).

Update Page Locks

The number of granted Update Page Locks. The type is integer (32-bit gauge).

The following names are defined for this attribute: UPDATE_PAGE_LOCKS or UPDPAGLOK (historical name), Update Page Locks (caption), Update_Page_Locks (attribute name), and UPDPAGLOK (column name).

MS SQL Log Shipping DB Details data set

The MS SQL Log Shipping DB Details data set provides detailed information about log shipping databases.

This data set contains the following attributes:

Database ID

Text: The ID of the primary or secondary database for the log shipping configuration. This attribute is a key attribute. The type is integer (32-bit numeric property).

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database ID (caption), Database_ID (attribute name), and DBID (column name).

Database Name

The name of the database. The type is string.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), Database Name (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Last File Processed

The absolute path of the most recent transaction log backup or restore. The type is string.

The following names are defined for this attribute: LAST_FILE_PROCESSED or LSTFILE (historical name), Last File Processed (caption), Last_File_Processed (attribute name), and LSTFILE (column name).

Last File Processed Time

The date and time at which the file is processed. The format is YY/MM/DD HH:mm, where, YY indicates year, MM indicates month, DD indicates day, HH indicates hour, and mm indicates minute. The type is string.

The following names are defined for this attribute: LAST_FILE_PROCESSED_TIME or LSTFILTM (historical name), Last File Processed Time (caption), Last_File_Processed_Time (attribute name), and LSTFILTM (column name).

Message

The job completion message text. The type is string.

The following names are defined for this attribute: MESSAGE or ERRMSG (historical name), Message (caption), Message (attribute name), and ERRMSG (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Sample Timestamp

The time when data is collected. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Server Type

Indicates whether the server is primary or secondary. The format is an integer (32-bit numeric property) with enumerated values. The following values are defined: Primary (0), Secondary (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_TYPE or SRVTYPEELS (historical name), Server Type (caption), Server_Type (attribute name), and SRVTYPEELS (column name).

Source

The source of the error message or event. The type is string.

The following names are defined for this attribute: SOURCE or LSSOURCE (historical name), Source (caption), Source (attribute name), and LSSOURCE (column name).

Threshold

The number of minutes allowed to elapse between backup or restore operations before an alert is generated. The type is integer (32-bit numeric property).

The following names are defined for this attribute: THRESHOLD (historical name), Threshold (caption), Threshold (attribute name), and THRESHOLD (column name).

Threshold Exceeded

Indicates whether the threshold is exceeded. The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: No (0), Yes (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THRESHOLD_EXCEEDED or THRESEXCD (historical name), Threshold Exceeded (caption), Threshold_Exceeded (attribute name), and THRESEXCD (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Log Shipping Errors data set

The MS SQL Log Shipping Errors data set provides detailed information about log shipping errors. This data set is not supported on the SQL Server 2005 and 2008 Express Edition.

This data set contains the following attributes:

Agent Type

The type of log shipping job. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Backup (0), Copy (1), Restore (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGENT_TYPE or AGTTYPE (historical name), Agent Type (caption), Agent_Type (attribute name), and AGTTYPE (column name).

Database ID

The ID of the primary or secondary database for the log shipping configuration. This attribute is a key attribute. The type is integer (32-bit numeric property).

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database ID (caption), Database_ID (attribute name), and DBID (column name).

Database Name

The name of the database. The type is string.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), Database Name (caption), Database_Name (attribute name), and DBNAME (column name).

Error Source

The source of the error message or event. The type is string.

The following names are defined for this attribute: ERROR_SOURCE or ERSOURCE (historical name), Error Source (caption), Error_Source (attribute name), and ERSOURCE (column name).

Log Time

The date and time at which the record is created. The format is YY/MM/DD HH:mm, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, and mm indicates minute.

The type is string.

The following names are defined for this attribute: LOG_TIME or LOGTIME (historical name), Log Time (caption), Log_Time (attribute name), and LOGTIME (column name).

Message

The job completion message text. The type is string.

The following names are defined for this attribute: MESSAGE or ERRMSG (historical name), Message (caption), Message (attribute name), and ERRMSG (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Sample Timestamp

The time when data is collected. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Sequence Number

An incremental number that indicates the correct order of information for errors that span multiple records. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SEQUENCE_NUMBER or SEQNUM (historical name), Sequence Number (caption), Sequence_Number (attribute name), and SEQNUM (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Log Shipping Jobs Detail data set

The MS SQL Log Shipping Jobs Detail data set provides detailed information about log shipping jobs. This data set is not supported on the SQL Server 2005 and 2008 Express Edition.

This data set contains the following attributes:

Agent Type

The type of log shipping job. The format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Backup (0), Copy (1), Restore (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGENT_TYPE or AGTTYPE (historical name), Agent Type (caption), Agent_Type (attribute name), and AGTTYPE (column name).

Database ID

The ID of the primary or secondary database for the log shipping configuration. For jobs that do not have an associated database, the value for the Database ID attribute is displayed as -1. The type is integer (32-bit numeric property).

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), Database ID (caption), Database_ID (attribute name), and DBID (column name).

Database Name

The name of the database. The type is string.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), Database Name (caption), Database_Name (attribute name), and DBNAME (column name).

Job Name

The log shipping job name. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: JOB_NAME or JOBNAME (historical name), Job Name (caption), Job_Name (attribute name), and JOBNAME (column name).

Job Status

The status of the session. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Starting (0), Running (1), Success (2), Error (3), Warning (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: JOB_STATUS or JOBSTATUS (historical name), Job Status (caption), Job_Status (attribute name), and JOBSTATUS (column name).

Log Time

The date and time at which the record is created. The format is YY/MM/DD HH:mm, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, and mm indicates minute. The type is string.

The following names are defined for this attribute: LOG_TIME or LOGTIME (historical name), Log Time (caption), Log_Time (attribute name), and LOGTIME (column name).

Message

The job completion message text. The type is string.

The following names are defined for this attribute: MESSAGE or ERRMSG (historical name), Message (caption), Message (attribute name), and ERRMSG (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Sample Timestamp

The time when data is collected. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Session ID

The primary ID for backup or the secondary ID for copy or restore. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SESSION_ID or SESSID (historical name), Session ID (caption), Session_ID (attribute name), and SESSID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Memory Manager data set

The MS SQL Memory Manager data set provides information about the memory usage. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 2 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Connection Memory(MB)

The total amount of dynamic memory the server is using for maintaining connections in megabytes. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONNECTION_MEMORY or CONMEM (historical name), *Connection Memory(MB)* (caption), Connection_Memory (attribute name), and CONMEM (column name).

Lock Memory(MB)

The total amount of dynamic memory the server is using for locks in megabytes. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_MEMORY or LKMEMKB (historical name), *Lock Memory(MB)* (caption), Lock_Memory (attribute name), and LKMEMKB (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Optimizer Memory(MB)

The total amount of dynamic memory the server is using for query optimization in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OPTIMIZER_MEMORY or OMEMKB (historical name), *Optimizer Memory(MB)* (caption), Optimizer_Memory (attribute name), and OMEMKB (column name).

SQL Cache Memory(MB)

The total amount of dynamic memory the server is using for the dynamic SQL cache in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_CACHE_MEMORY or SQLCAMEM (historical name), *SQL Cache Memory(MB)* (caption), SQL_Cache_Memory (attribute name), and SQLCAMEM (column name).

Target Server Memory(MB)

The total amount of dynamic memory the server is willing to consume in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TARGET_SERVER_MEMORY or TARGMEM (historical name), *Target Server Memory(MB)* (caption), Target_Server_Memory (attribute name), and TARGMEM (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Server Memory(MB)

The total amount of dynamic memory the server is currently consuming in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SERVER_MEMORY or TOTMEM (historical name), *Total Server Memory(MB)* (caption), Total_Server_Memory (attribute name), and TOTMEM (column name).

Granted Workspace Memory(MB)

The total amount of memory granted to running processes. This memory is used for the hash, sort, and create index operations in megabytes. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GRANTED_WORKSPACE_MEMORY or GRNTMEM (historical name), *Granted Workspace Memory(MB)* (caption), Granted_Workspace_Memory (attribute name), and GRNTMEM (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Lock Blocks

The current number of lock blocks that are in use on the server. Refreshed periodically. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_BLOCKS or LKBLKS (historical name), Lock Blocks (caption), Lock_Blocks (attribute name), and LKBLKS (column name).

Lock Blocks Allocated

The current number of allocated lock blocks. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_BLOCKS_ALLOCATED or LKBLKAL (historical name), Lock Blocks Allocated (caption), Lock_Blocks_Allocated (attribute name), and LKBLKAL (column name).

Lock Owner Blocks

The number of lock owner blocks that are in use on the server. The blocks are refreshed periodically. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_OWNER_BLOCKS or LOBLKS (historical name), Lock Owner Blocks (caption), Lock_Owner_Blocks (attribute name), and LOBLKS (column name).

Lock Owner Blocks Allocated

The current number of allocated lock owner blocks. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOCK_OWNER_BLOCKS_ALLOCATED or LOWNBAL (historical name), Lock Owner Blocks Allocated (caption), Lock_Owner_Blocks_Allocated (attribute name), and LOWNBAL (column name).

Maximum Workspace Memory(MB)

The total amount of memory granted to running processes. This memory is used primarily for the hash, sort, and create index operations in megabytes. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_WORKSPACE_MEMORY or MAXWMEM (historical name), Maximum Workspace Memory(MB) (caption), Maximum_Workspace_Memory (attribute name), and MAXWMEM (column name).

Memory Grants Outstanding

The current number of processes that have successfully acquired a workspace memory. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_GRANTS_OUTSTANDING or MEMGOUT (historical name), Memory Grants Outstanding (caption), Memory_Grants_Outstanding (attribute name), and MEMGOUT (column name).

Memory Grants Pending

The current number of processes waiting for a workspace memory grant. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_GRANTS_PENDING or MEMGPEN (historical name), Memory Grants Pending (caption), Memory_Grants_Pending (attribute name), and MEMGPEN (column name).

Sample Timestamp

The time when data is collected. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Problem Detail data set

The MS SQL Problem Detail data set contains attributes that you can use to display detailed information about a selected error. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Error ID

The ID of the error message. The value format is an integer; for example, 2520. The unnumbered errors with text explanations are written to the SQL server message logs. The type is string with enumerated values. The following values are defined: Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_ID or ERRORID (historical name), *Error ID* (caption), Error_ID (attribute name), and ERRORID (column name).

Error SPID

The ID of the session when an event occurs. It is prefixed with spid. For example, spid34s, spid57. When the ID of the session is unavailable, the source of error is showed as it is. For example, logon. This attribute is not applicable for eventing thresholds. The type is string.

The following names are defined for this attribute: ERROR_SPID or ERRSPID (historical name), *Error SPID* (caption), Error_SPID (attribute name), and ERRSPID (column name).

Message Age (Min.)

The number of minutes that have elapsed since the error occurred. The value format is an integer; for example, 2. Monitor this value to track current messages. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE_AGE or MSGAGE (historical name), *Message Age (Min.)* (caption), Message_Age (attribute name), and MSGAGE (column name).

Message Text (Unicode)

The message text. This attribute is not available for use in eventing thresholds. The type is string with enumerated values. The following values are defined: Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE_TEXT_U or UMSGTEXT (historical name), *Message Text (Unicode)* (caption), Message_Text_U (attribute name), and UMSGTEXT (column name).

Message Timestamp

The timestamp that indicates the date and time the error occurred. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: Invalid (0000000000000000), N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE_TIMESTAMP or MSGTIME (historical name), *Message Timestamp* (caption), Message_Timestamp (attribute name), and MSGTIME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Severity Level

Indicates the severity level of the error. The type is string with enumerated values. The following values are defined: Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SEVERITY_LEVEL or SEVERITY (historical name), *Severity Level* (caption), Severity_Level (attribute name), and SEVERITY (column name).

SQL State Code

The SQL state value for the error message. The value format is an integer; for example, 37. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_STATE_CODE or SQLSTATE (historical name), *SQL State Code* (caption), SQL_State_Code (attribute name), and SQLSTATE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), Host_Name (attribute name), and HOSTNAME (column name).

Message Issuer

The source of the error message. The value format is an alphanumeric string with a maximum of 8 characters; for example, JSMITH. The type is string with enumerated values. The following values are defined: Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE_ISSUER or MSGISSUER (historical name), *Message Issuer* (caption), Message_Issuer (attribute name), and MSGISSUER (column name).

Message Text

The message text. This attribute is not available for use in eventing thresholds. The type is string.

The following names are defined for this attribute: MESSAGE_TEXT or MSGTEXT (historical name), *Message Text* (caption), Message_Text (attribute name), and MSGTEXT (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Problem Summary data set

The MS SQL Problem Summary data set contains attributes that you can use to monitor summary information about SQL server errors. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Age of Last Error (Min.)

The number of minutes that have elapsed since the last error message occurred. The value format is an integer; for example, 2. The type is integer (32-bit numeric property).

The following names are defined for this attribute: AGE_OF_LAST_ERROR or LASTERRAGE (historical name), *Age of Last Error (Min.)* (caption), Age_of_Last_Error (attribute name), and LASTERRAGE (column name).

Error Log Size (Bytes)

The number of bytes in the error log file. The value format is an integer; for example, 50000. The type is integer (32-bit gauge).

The following names are defined for this attribute: ERROR_LOG_SIZE or ERRLOGSZ (historical name), *Error Log Size (Bytes)* (caption), Error_Log_Size (attribute name), and ERRLOGSZ (column name).

Maximum Sev Current Interval

The error message of the highest severity level encountered during the current interval. The value format is an integer in the range 10 through 24; for example, 19. The type is string.

The following names are defined for this attribute: MAXIMUM_SEV_CURRENT_INTERVAL or MAXSEVCURR (historical name), *Maximum Sev Current Interval* (caption), Maximum_Sev_Current_Interval (attribute name), and MAXSEVCURR (column name).

Maximum Sev Timestamp

The timestamp that indicates the date and time the error message with the highest severity level occurred. The format is MM/DD/YY HH:mm:SS, where MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: Invalid (0000000000000000), N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_SEV_TIMESTAMP or MAXSEVTIME (historical name), *Maximum Sev Timestamp* (caption), Maximum_Sev_Timestamp (attribute name), and MAXSEVTIME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Error Messages

The number of error messages that have occurred since the SQL server was started. The value format is an integer; for example, 3. The type is integer (32-bit counter).

The following names are defined for this attribute: TOTAL_ERROR_MESSAGES or NUMERRORS (historical name), *Total Error Messages* (caption), Total_Error_Messages (attribute name), and NUMERRORS (column name).

Total Errors Cur Intvl

The number of error messages that occurred during the current interval. The value format is an integer; for example, 1. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_ERRORS_CURRENT_INTERVAL or NUMCURR (historical name), *Total Errors Cur Intvl* (caption), Total_Errors_Current_Interval (attribute name), and NUMCURR (column name).

Total Errors Other

The number of error messages with a severity level of less than 17 that have occurred since the SQL server was started. The value format is an integer; for example, 3. The type is integer (32-bit counter).

The following names are defined for this attribute: TOTAL_ERRORS_OTHER or NUMOTHER (historical name), *Total Errors Other* (caption), Total_Errors_Other (attribute name), and NUMOTHER (column name).

Current Interval (Sec.)

The number of seconds that have elapsed between the previous sample and the current sample. The value format is an integer; for example, 90. A sample contains the data that IBM Tivoli Monitoring for Microsoft SQL Server collects about each SQL server. New data becomes available if a new interval has occurred and data has been refreshed at the CMS hub. The type is integer (32-bit numeric property).

The following names are defined for this attribute: CURRENT_INTERVAL or CURRINTL (historical name), Current Interval (Sec.) (caption), Current_Interval (attribute name), and CURRINTL (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Maximum Sev Level

The level of highest severity encountered since the SQL server started. The value format is an integer in the range 10 through 24; for example, 22. The type is string.

The following names are defined for this attribute: MAXIMUM_SEV_LEVEL or MAXSEV (historical name), Maximum Sev Level (caption), Maximum_Sev_Level (attribute name), and MAXSEV (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, where MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Errors High Sev

The number of error messages with a severity level of 17 or higher that have occurred since the SQL server was started. The value format is an integer; for example, 2. The type is integer (32-bit counter).

The following names are defined for this attribute: TOTAL_ERRORS_HIGH_SEV or NUMSEVG17 (historical name), Total Errors High Sev (caption), Total_Errors_High_Sev (attribute name), and NUMSEVG17 (column name).

MS SQL Process Detail data set

The MS SQL Process Detail data set contains attributes that you can use to monitor detailed information about a selected SQL server process. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Blocking Process Duration(Mins.)

The total time (mins) for which the current blocking process has blocked a request since the agent start up. If a request is blocked by multiple processes during the data collection interval, the process

ID of the blocking process that first blocked a request, and the total blocking process duration of all the blocking processes are displayed. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_PROCESS_DURATION or BLKPROCDU (historical name), *Blocking Process Duration(Mins.)* (caption), Blocking_Process_Duration (attribute name), and BLKPROCDU (column name).

Blocking Process ID

The identifier for the process that is blocking a request for a lock. The value format is an alphanumeric string with a maximum of 8 characters; for example, 12. Some blocking processes can become stranded. Investigate eventing thresholds in which a process is being blocked for an extended period of time. For more information on a blocking process, query the sysprocesses and use the sp-lock procedure. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_PROCESS_ID or BLOCKID (historical name), *Blocking Process ID* (caption), Blocking_Process_ID (attribute name), and BLOCKID (column name).

Client Host Name

The name of the host for the client. The value format is an alphanumeric string with a maximum of 16 characters; for example, Rocket. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLIENT_HOST_NAME or CLNTHOST (historical name), *Client Host Name* (caption), Client_Host_Name (attribute name), and CLNTHOST (column name).

Client Process ID

The ID the client assigned to the process. The value format is an alphanumeric string with a maximum of 16 characters; for example, amc_2236. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLIENT_PROCESS_ID or CLNTPID (historical name), *Client Process ID* (caption), Client_Process_ID (attribute name), and CLNTPID (column name).

Client User ID

The ID of the user executing the process. The value format is an alphanumeric string with a maximum of 8 characters; for example, J_Kelly. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLIENT_USER_ID or CLNTUSER (historical name), *Client User ID* (caption), Client_User_ID (attribute name), and CLNTUSER (column name).

Command

The name of the command being executed by the process. The value format is an alphanumeric string with a maximum of 16 characters; for example, CREATE VIEW. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COMMAND or CMDNAME (historical name), *Command* (caption), Command (attribute name), and CMDNAME (column name).

Current CPU Pct Used

Relative percent of the CPU used by this SQL Server process compared to all other SQL Server processes. The value format is a percentage with two decimal places allowed; for example, 80.00. This dynamic information is from the sysprocesses table. Set alerts for processes using an abnormal amount of CPU, relative to all other SQL Server processes. The type is real number (32-bit gauge) with

two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_CPU_PERCENT_USED or PCTCPU (historical name), *Current CPU Pct Used* (caption), Current_CPU_Percent_Used (attribute name), and PCTCPU (column name).

Database Name (Unicode)

The database name. The value format is an alphanumeric string with a maximum of 384 bytes. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), *Database Name (Unicode)* (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Duration (Sec.)

The time (in seconds) for which the process has been running. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DURATION or PROCDURSEC (historical name), *Duration (Sec.)* (caption), Duration (attribute name), and PROCDURSEC (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Process ID

The ID of the process. The value format is an alphanumeric string with a maximum of 10 characters; for example, 42168. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCESS_ID or PROCESSID (historical name), *Process ID* (caption), Process_ID (attribute name), and PROCESSID (column name).

Process Status

Indicates the status of the process. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCESS_STATUS or STATUS (historical name), *Process Status* (caption), Process_Status (attribute name), and STATUS (column name).

Program Name (Unicode)

The Program Name. The value format is an alphanumeric string with a maximum of 384 bytes. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROGRAM_NAME_U or UPGMNAME (historical name), *Program Name (Unicode)* (caption), Program_Name_U (attribute name), and UPGMNAME (column name).

Server User ID

The SQL server-assigned ID for the user executing the process. The value format is an alphanumeric string with a maximum of 8 characters; for example, S. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_USER_ID or SRVUSER (historical name), *Server User ID* (caption), Server_User_ID (attribute name), and SRVUSER (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: *TIMESTAMP* (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and *TIMESTAMP* (column name).

Total CPU Time (Sec.)

The amount of CPU time, in seconds, the process has used on the host since the process started. The value format is an integer; for example, 60. This value is based on the statistics collected by the SQL server. Use this value to check for processes that use abnormal amounts of CPU time. The type is real number (32-bit gauge) with four decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *TOTAL_CPU_TIME* or *CPU* (historical name), *Total CPU Time (Sec.)* (caption), *Total_CPU_Time* (attribute name), and *CPU* (column name).

Total Disk IO

The number of accesses to hard disk since the process started. The value includes accesses to hard disk for physical reads and physical writes. The value format is an integer; for example, 10. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *TOTAL_DISK_IO* or *IO* (historical name), *Total Disk IO* (caption), *Total_Disk_IO* (attribute name), and *IO* (column name).

Total Memory Alloc (KB)

The number of KB allocated for this process, based on the number of pages in the procedure cache. A negative number indicates that the process is freeing memory allocated by another process. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *TOTAL_MEMORY_ALLOC* or *MEMORY* (historical name), *Total Memory Alloc (KB)* (caption), *Total_Memory_Alloc* (attribute name), and *MEMORY* (column name).

Blocking Resource Duration(Mins.)

The total time (mins) for which the blocking process is currently blocking a request for a resource type. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *BLOCKING_RESOURCE_DURATION* or *BLKRESDU* (historical name), *Blocking Resource Duration(Mins.)* (caption), *Blocking_Resource_Duration* (attribute name), and *BLKRESDU* (column name).

Blocking Resource Type

The resource type of the request that is currently blocked by the blocking process. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *BLOCKING_RESOURCE_TYPE* or *BLKRESTY* (historical name), *Blocking Resource Type* (caption), *Blocking_Resource_Type* (attribute name), and *BLKRESTY* (column name).

Client Group ID

The group ID of the user executing the process. This value does not apply to Microsoft SQL Server version 7. The value format is an alphanumeric string with a maximum of 8 characters; for example, ACCT_1. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *CLIENT_GROUP_ID* or *SRVGROUP* (historical name), *Client Group ID* (caption), *Client_Group_ID* (attribute name), and *SRVGROUP* (column name).

Command (Unicode)

The name of the executing command. The value format is an alphanumeric string with a maximum of 48 bytes. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COMMAND_U or UCMDNAME (historical name), Command (Unicode) (caption), Command_U (attribute name), and UCMDNAME (column name).

Database Name

The name of the database. The value format is an alphanumeric string with a maximum of 30 characters; for example, KOQ3. Each database name is unique. The SQL server also assigns each database its own identification number. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or DBNAME (historical name), Database Name (caption), Database_Name (attribute name), and DBNAME (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Login Name

ID used by the process to log in to the SQL Server. The value format is an alphanumeric string with a maximum of 30 characters. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOGIN_NAME or PROCESSNM (historical name), Login Name (caption), Login_Name (attribute name), and PROCESSNM (column name).

NT User

User name associated with the process. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NT_USER or PROCNTUSER (historical name), NT User (caption), NT_User (attribute name), and PROCNTUSER (column name).

OS Process ID

The Microsoft Windows thread ID. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OS_PROCESS_ID or OSPID (historical name), OS Process ID (caption), OS_Process_ID (attribute name), and OSPID (column name).

Program Name

The name of the program (front-end module) for the process. The value format is an alphanumeric string with a maximum of 16 characters; for example, ISQL. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROGRAM_NAME or PGMNAME (historical name), Program Name (caption), Program_Name (attribute name), and PGMNAME (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Process Summary data set

The MS SQL Process Summary data set contains attributes that you can use to monitor summary information about processes. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Percent Processes Blocked

The percentage of processes that are being blocked. The value format is a percentage with two decimal places allowed; for example, 4.00. This value is based on information from the sysprocesses table. The value includes all processes currently in a waiting state. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_BLOCKED or PCTBLOCK (historical name), *Percent Processes Blocked* (caption), Percent_Processes_Blocked (attribute name), and PCTBLOCK (column name).

Server CPU Pct System

The percentage of CPU time the SQL server processes are using on the host. The value format is a percentage with two decimal places allowed; for example, 25.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: SERVER_CPU_PERCENT_SYSTEM or SRVCPUPCT (historical name), *Server CPU Pct System* (caption), Server_CPU_Percent_System (attribute name), and SRVCPUPCT (column name).

Server CPU Percent Application

The percentage of CPU time the SQL server application processes are using on the host. The value format is a percentage with two decimal places allowed; for example, 20.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: SERVER_CPU_PERCENT_APPLICATION or APPCPUPCT (historical name), *Server CPU Percent Application* (caption), Server_CPU_Percent_Application (attribute name), and APPCPUPCT (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Log Suspend

The number of processes with a status of log suspend. A process with a status of log suspend is a process that cannot complete until there is free space in the transaction log. The value format is an integer; for example, 1. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_LOG_SUSPEND or NUMLOGSUS (historical name), *Total Log Suspend* (caption), Total_Log_Suspend (attribute name), and NUMLOGSUS (column name).

Total Other Sleep

The number of processes with a status of othersleep. The value format is an integer; for example, 2. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL_OTHER_SLEEP or NUMSLPOTH (historical name), *Total Other Sleep* (caption), Total_Other_Sleep (attribute name), and NUMSLPOTH (column name).

Total Processes Bad

The number of processes with a status of bad. The value format is an integer; for example, 1. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL PROCESSES_BAD or NUMBAD (historical name), *Total Processes Bad* (caption), Total_Processes_Bad (attribute name), and NUMBAD (column name).

Total Processes Blocked

The number of processes that are being blocked. The value format is an integer; for example, 2. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL PROCESSES_BLOCKED or NUMBLOCK (historical name), *Total Processes Blocked* (caption), Total_Processes_Blocked (attribute name), and NUMBLOCK (column name).

Total Processes Infected

The number of processes with a status of infected. A process with a status of infected is a process that cannot be completed. The value format is an integer; for example, 1. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL PROCESSES_INFECTED or NUMINFECT (historical name), *Total Processes Infected* (caption), Total_Processes_Infected (attribute name), and NUMINFECT (column name).

Total Processes Locksleep

The number of processes with a status of locksleep. A process with a status of locksleep is a process waiting for a lock on a resource to be released. The value format is an integer; for example, 5. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL PROCESSES_LOCKSLEEP or NUMSLPLCK (historical name), *Total Processes Locksleep* (caption), Total_Processes_Locksleep (attribute name), and NUMSLPLCK (column name).

Total Processes Stopped

The number of processes with a status of stopped. The value format is an integer; for example, 1. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL PROCESSES STOPPED or NUMSTOPPED (historical name), *Total Processes Stopped* (caption), Total_Processes_Stopped (attribute name), and NUMSTOPPED (column name).

Total Processes

The number of processes. The value includes background processes, processes for applications, and user processes. The value format is an integer; for example, 50. This value includes background processes, processes for applications, and user processes. The type is integer (32-bit gauge).

The following names are defined for this attribute: TOTAL PROCESSES or NUMPROCESS (historical name), *Total Processes* (caption), Total_Processes (attribute name), and NUMPROCESS (column name).

Current Interval (Sec.)

The number of seconds that have elapsed between the previous sample and the current sample. The value format is an integer; for example, 90. A sample contains the data that IBM Tivoli Monitoring for Microsoft SQL Server collects about each SQL server. New data becomes available if a new interval has occurred and data has been refreshed at the CMS hub. The type is integer (32-bit numeric property).

The following names are defined for this attribute: CURRENT_INTERVAL or CURRINTL (historical name), Current Interval (Sec.) (caption), Current_Interval (attribute name), and CURRINTL (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Percent Processes Bad

The percentage of processes with a status of bad. The value format is a percentage with two decimal places allowed; for example, 2.00. This value is based on information from the sysprocesses table. Bad processes are often associated with a process ID problem. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_BAD or PCTBAD (historical name), Percent Processes Bad (caption), Percent_Processes_Bad (attribute name), and PCTBAD (column name).

Percent Processes Infected

The percentage of processes with a status of infected. The value format is a percentage with two decimal places allowed; for example, 2.00. This value is based on information from the sysprocesses table. An infected process is associated with a serious error condition. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_INFECTED or PCTINFEKT (historical name), Percent Processes Infected (caption), Percent_Processes_Infected (attribute name), and PCTINFEKT (column name).

Percent Processes Locksleep

The percentage of processes with a status of locksleep. The value format is a percentage with two decimal places allowed; for example, 10.00. This value is based on information from the sysprocesses table. The processes are waiting to obtain locks on resources. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_LOCKSLEEP or PCTSLPLCK (historical name), Percent Processes Locksleep (caption), Percent_Processes_Locksleep (attribute name), and PCTSLPLCK (column name).

Percent Processes Othersleep

The percentage of processes with a status of alarm sleep, recv sleep, send sleep, or sleeping. The value format is a percentage with two decimal places allowed; for example, 10.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_OTHERSLEEP or PCTSLPOTH (historical name), Percent Processes Othersleep (caption), Percent_Processes_Othersleep (attribute name), and PCTSLPOTH (column name).

Percent Processes Sleeping

The percentage of processes with a status of sleep. The value format is a percentage with two decimal places allowed; for example, 20.00. This value is based on information from the sysprocesses table. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_SLEEPING or PCTSLEEP (historical name), Percent Processes Sleeping (caption), Percent_Processes_Sleeping (attribute name), and PCTSLEEP (column name).

Percent Processes Stopped

The percentage of processes with a status of stopped. The value format is a percentage with two decimal places allowed; for example, 2.00. This value is based on information from the sysprocesses table. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT PROCESSES_STOPPED or PCTSTOPPED (historical name), Percent Processes Stopped (caption), Percent_Processes_Stopped (attribute name), and PCTSTOPPED (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:ss, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Remote Servers data set

The MS SQL Remote Servers data set contains attributes that you can use to monitor summary information about remote SQL servers.

This data set contains the following attributes:

Connection Level

The connection level for the remote SQL server. The value format is an integer; for example, 3. If the SQL server is not a Microsoft SQL Server Secure SQL server, the value is 0. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONNECTION_LEVEL or RSRVCONNLV (historical name), Connection Level (caption), Connection_Level (attribute name), and RSRVCONNLV (column name).

Current Interval (Sec.)

The number of seconds that have elapsed between the previous sample and the current sample. The value format is an integer; for example, 90. A sample contains the data that IBM Tivoli Monitoring for Microsoft SQL Server collects about each SQL server. New data becomes available if a new interval has occurred and data has been refreshed at the CMS hub. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_INTERVAL or CURRINTL (historical name), Current Interval (Sec.) (caption), Current_Interval (attribute name), and CURRINTL (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Network Name

The name of the network for the remote SQL server. The value format is an alphanumeric string with a maximum of 16 characters; for example, REM_NET. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NETWORK_NAME or RSRVNETNM (historical name), Network Name (caption), Network_Name (attribute name), and RSRVNETNM (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Remote Server ID

The ID of the remote SQL server. The value format is an alphanumeric string with a maximum of 12 characters; for example, SERV_ID. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REMOTE_SERVER_ID or RSRVID (historical name), Remote Server ID (caption), Remote_Server_ID (attribute name), and RSRVID (column name).

Remote Server Name

The name of the remote SQL server. The value format is an alphanumeric string with a maximum of 12 characters; for example, REM_SVR5. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REMOTE_SERVER_NAME or RSRVNAME (historical name), Remote Server Name (caption), Remote_Server_Name (attribute name), and RSRVNAME (column name).

Remote Server Status

The status of the remote SQL server. The server status is displayed as Unknown when the collector process of the SQL Server agent is not running. The server status is displayed as Inactive when the collector process of the SQL Server agent is running, but the SQL Server is not responding to the request. The type is string with enumerated values. The following values are defined: Inactive (Inactive), Unknown (Unknown), Active (Active). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REMOTE_SERVER_STATUS or RSRVSTATUS (historical name), Remote Server Status (caption), Remote_Server_Status (attribute name), and RSRVSTATUS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

MS SQL Resource Pool Stats data set

The MS SQL Resource Pool Stats data set provides statistics associated with resource pools. This data set is not supported on the SQL Server 2005 and 2008 Express Edition.

This data set contains the following attributes:

Active Memory Grant Amount(MB)

The total amount of granted memory in megabytes in the resource pool. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTIVE_MEMORY_GRANT_AMOUNT or ACTMEMGA (historical name), Active Memory Grant Amount(MB) (caption), Active_Memory_Grant_Amount (attribute name), and ACTMEMGA (column name).

Active Memory Grants Count

The count of memory grants that are assigned to pool workers. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTIVE_MEMORY_GRANTS_COUNT or ACTMEMGC (historical name), Active Memory Grants Count (caption), Active_Memory_Grants_Count (attribute name), and ACTMEMGC (column name).

Cache Memory Target(MB)

The current memory target for cache memory in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CACHE_MEMORY_TARGET or CACHMEMTAR (historical name), Cache Memory Target(MB) (caption), Cache_Memory_Target (attribute name), and CACHMEMTAR (column name).

Compile Memory Target(MB)

The current memory target for compile memory in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COMPILE_MEMORY_TARGET or COMPMEMTAR (historical name), Compile Memory Target(MB) (caption), Compile_Memory_Target (attribute name), and COMPMEMTAR (column name).

CPU Control Effect Pct

The effect of the resource governor on the resource pool calculated as (CPU Usage Percent)/(CPU Usage Percent without RG). The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_CONTROL_EFFECT_PERCENT or CPUCONEFF (historical name), CPU Control Effect Pct (caption), CPU_Control_Effect_Percent (attribute name), and CPUCONEFF (column name).

CPU Usage Pct

The system CPU Usage by all requests in the specified instance of the performance object. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_USAGE_PERCENT or RPSCPUPER (historical name), CPU Usage Pct (caption), CPU_Usage_Percent (attribute name), and RPSCPUPER (column name).

CPU Usage Target Pct

The target value of CPU Usage Percent for the resource pool based on the configuration settings and the system load. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_USAGE_TARGET_PERCENT or CPUTARGET (historical name), CPU Usage Target Pct (caption), CPU_Usage_Target_Percent (attribute name), and CPUTARGET (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Max Memory(MB)

The maximum amount of memory in megabytes the resource pool can have based on the settings and the server state. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_MEMORY or MAXMEMKB (historical name), Max Memory(MB) (caption), Max_Memory (attribute name), and MAXMEMKB (column name).

Memory Grant Timeouts Per Sec

The number of query memory grant timeouts per second occurring in the resource pool. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_GRANT_TIMEOUTS_PER_SEC or MGTPERSEC (historical name), Memory Grant Timeouts Per Sec (caption), Memory_Grant_Timeouts_Per_Sec (attribute name), and MGTPERSEC (column name).

Memory Grants Per Sec

The number of query memory grants per second occurring in the resource pool. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_GRANTS_PER_SEC or MEMGPS (historical name), Memory Grants Per Sec (caption), Memory_Grants_Per_Sec (attribute name), and MEMGPS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Pending Memory Grants Count

The number of queries waiting for memory grants in the resource pool. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PENDING_MEMORY_GRANTS_COUNT or PENMEMGC (historical name), Pending Memory Grants Count (caption), Pending_Memory_Grants_Count (attribute name), and PENMEMGC (column name).

Query Exec Memory Target(MB)

The current memory target for query execution in megabytes. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_EXEC_MEMORY_TARGET or QUERYEMT (historical name), Query Exec Memory Target(MB) (caption), Query_Exec_Memory_Target (attribute name), and QUERYEMT (column name).

Resource Pool Name (Unicode)

The resource governor resource pool name. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RESOURCE_POOL_NAME or RPNAME (historical name), Resource Pool Name (Unicode) (caption), Resource_Pool_Name (attribute name), and RPNAME (column name).

Sample Timestamp

The time when data is collected. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Target Memory(MB)

The target amount of memory in megabytes the resource pool tries to attain based on the settings and the server state. The format is an integer. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TARGET_MEMORY or TARMEMKB (historical name), Target Memory(MB) (caption), Target_Memory (attribute name), and TARMEMKB (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Used Memory(MB)

The used amount of memory in megabytes in the resource pool. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: USED_MEMORY or USEMEMKB (historical name), Used Memory(MB) (caption), Used_Memory (attribute name), and USEMEMKB (column name).

MS SQL Running Queries Details data set

The MS SQL Running Queries Details data set contains attributes that provide information about the queries that are currently running on the SQL Server.

This data set contains the following attributes:

Blocking Session ID

The ID of the session that is currently blocking the execution of a query on the SQL Server. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKING_SESSION_ID or BLKSESSID (historical name), Blocking Session ID (caption), Blocking_Session_ID (attribute name), and BLKSESSID (column name).

Client Host Name

The name of the client workstation that is associated with the SQL query that is currently running on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLIENT_HOST_NAME or HSTNAME (historical name), Client Host Name (caption), Client_Host_Name (attribute name), and HSTNAME (column name).

Command Type

The type of command that is currently being processed by the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COMMAND_TYPE or CMDTYPE (historical name), Command Type (caption), Command_Type (attribute name), and CMDTYPE (column name).

CPU Time (Sec.)

The CPU time (in seconds) of the SQL query that is currently running on the SQL Server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_TIME_SEC or CPUTIMESEC (historical name), CPU Time (Sec.) (caption), CPU_Time_Sec (attribute name), and CPUTIMESEC (column name).

Database Name (Unicode)

The name of the database against which the SQL query is currently running on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or DBSNAME (historical name), Database Name (Unicode) (caption), Database_Name_U (attribute name), and DBSNAME (column name).

Granted Query Memory

The number of pages that are currently allocated to the execution of a query. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: GRANTED_QUERY_MEMORY or GRNTEDQMEM (historical name), Granted Query Memory (caption), Granted_Query_Memory (attribute name), and GRNTEDQMEM (column name).

Login Name

The name that is used to log on to the SQL Server on which the query is currently running. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LOGIN_NAME or LOGINNAME (historical name), Login Name (caption), Login_Name (attribute name), and LOGINNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Percent Work Complete

The percentage of completed work for the database command that is currently running. The name of the command that is currently running is displayed in the Command Type attribute. The following list displays the database commands:

- ALTER INDEX REORGANIZE
- AUTO_SHRINK option with ALTER DATABASE
- BACKUP DATABASE
- DBCC CHECKDB
- DBCC CHECKFILEGROUP
- DBCC CHECKTABLE
- DBCC INDEXDEFRAG
- DBCC SHRINKDATABASE
- DBCC SHRINKFILE
- RECOVERY
- RESTORE DATABASE
- ROLLBACK
- TDE ENCRYPTION

The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERCENT_WORK_COMPLETE or PERWRKCMPL (historical name), Percent Work Complete (caption), Percent_Work_Complete (attribute name), and PERWRKCMPL (column name).

Program Name

The name of the client program that initiated the request, which is currently running on the SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROGRAM_NAME or PROGNAME (historical name), Program Name (caption), Program_Name (attribute name), and PROGNAME (column name).

Query Reads

The number of query read operations completed by the SQL query that is currently running on the SQL Server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_READS or READS (historical name), Query Reads (caption), Query_Reads (attribute name), and READS (column name).

Query Start Time

The time when the query arrived on the SQL Server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL Server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_START_TIME or QSTRRTIME (historical name), Query Start Time (caption), Query_Start_Time (attribute name), and QSTRRTIME (column name).

Query Status

The status of the SQL query that is currently running on the SQL Server. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1),

Background (1), Running (2), Runnable (3), Sleeping (4), Suspended (5). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_STATUS or QSTATUS (historical name), Query Status (caption), Query_Status (attribute name), and QSTATUS (column name).

Query Writes

The number of query write operations completed by the SQL query that is currently running on the SQL Server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_WRITES or WRITES (historical name), Query Writes (caption), Query_Writes (attribute name), and WRITES (column name).

Row Count

The number of rows returned to the client by the SQL query that is currently running on the SQL Server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ROW_COUNT or ROWCOUNT (historical name), Row Count (caption), Row_Count (attribute name), and ROWCOUNT (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Running Query Text

The SQL text of the query that is currently running on the SQL Server. The value format is an alphanumeric string with a maximum of 2048 characters. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RUNNING_QUERY_TEXT or RUNSQLTEXT (historical name), Running Query Text (caption), Running_Query_Text (attribute name), and RUNSQLTEXT (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Session ID

The session ID of the query that is currently running on the SQL Server. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SESSION_ID or CURSESSID (historical name), Session ID (caption), Session_ID (attribute name), and CURSESSID (column name).

SQL Handle

The SQL handle of the query that is currently running on the SQL Server. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_HANDLE or SQLHANDLE (historical name), SQL Handle (caption), SQL_Handle (attribute name), and SQLHANDLE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Elapsed Time (Sec.)

The total time (in seconds) that has elapsed since the SQL query arrived on the SQL Server. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ELAPSED_TIME_SEC or TOTELTIMES (historical name), Total Elapsed Time (Sec.) (caption), Total_Elapsed_Time_Sec (attribute name), and TOTELTIMES (column name).

MS SQL Server Detail data set

The MS SQL Server Detail data set contains attributes that you can use to monitor detailed information about a selected SQL server. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Current Locks

The number of current locks for the SQL server. The value format is an integer; for example, 73. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_LOCKS or CURLOCKS (historical name), *Current Locks* (caption), Current_Locks (attribute name), and CURLOCKS (column name).

Data Cache Size (KB)

The number of kilobytes (KB) allocated for the data cache memory. The SQL server uses the data cache to store data and index pages. The value format is an integer; for example, 1000. The cache is sometimes referred to as the buffer cache. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_CACHE_SIZE or DATACSZ (historical name), *Data Cache Size (KB)* (caption), Data_Cache_Size (attribute name), and DATACSZ (column name).

Error Log Size (Bytes)

The number of bytes in the error log file. The value format is an integer; for example, 50000. The error log contains the fatal error and kernel error messages issued by the SQL server. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_LOG_SIZE or ERRLOGSZ (historical name), *Error Log Size (Bytes)* (caption), Error_Log_Size (attribute name), and ERRLOGSZ (column name).

Max Locks Allowed

The maximum number of allowable locks. The value format is an integer; for example, 2000. This value is a configuration parameter. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_LOCKS_ALLOWED or MAXLOCKS (historical name), *Max Locks Allowed* (caption), Max_Locks_Allowed (attribute name), and MAXLOCKS (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

OS Type

The operating system for the SQL server. The value format is an alphanumeric string; for example, AIX . The type is string with enumerated values. The following values are defined: Windows Srv 2008 (Win2008), Windows Srv 2003 (Win2003), Windows Srv 2000 (Win2000), Windows NT (WinNT), Windows Srv 2012 (Win2012), Windows Srv 2008 R2 (Win8R2), Windows Srv 2012 R2 (Win12R2), Windows Srv 2016 (win2016), Windows Srv 2019 (win2019), No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OS_TYPE or OSTYPE (historical name), *OS Type* (caption), OS_Type (attribute name), and OSTYPE (column name).

OS Version

The version of the operating system for the SQL server. The value format is the version in the format version.release; for example, 2. 5. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OS_VERSION or OSVERSN (historical name), *OS Version* (caption), OS_Version (attribute name), and OSVERSN (column name).

Procedure Buffers Pct Active

The percentage of slots with a procedure that is currently executing. Use this attribute to see how much of the cache these procedures use in relation to the cache size as a whole. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_BUFFERS_PERCENT_ACTIVE or PCTPROCBA (historical name), *Procedure Buffers Pct Active* (caption), Procedure_Buffers_Percent_Active (attribute name), and PCTPROCBA (column name).

Procedure Buffers Pct Used

The percentage of slots that have a procedure in them. A procedure buffer is considered used when it is associated with a procedure cache entry. A used procedure buffer can be active or not active. Use this attribute to see how much of the cache is being used by currently compiled procedures in relation to the cache size as a whole. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_BUFFERS_PERCENT_USED or PCTPROCBU (historical name), *Procedure Buffers Pct Used* (caption), Procedure_Buffers_Percent_Used (attribute name), and PCTPROCBU (column name).

Procedure Cache Pct Active

The total size of the procedure cache in pages. The size of the procedure cache can fluctuate depending on the activity of other database server processes that might require procedure cache slots, such as query plans. Use the attribute to see the current size of the procedure cache. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_CACHE_PERCENT_ACTIVE or PCTPROCCA (historical name), *Procedure Cache Pct Active* (caption), Procedure_Cache_Percent_Active (attribute name), and PCTPROCCA (column name).

Procedure Cache Pct Used

The percentage of the procedure cache that has procedures in it. Use this attribute to see how much of the cache currently executing procedures use in relation to the cache size as a whole. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_CACHE_PERCENT_USED or PCTPROCCU (historical name), *Procedure Cache Pct Used* (caption), Procedure_Cache_Percent_Used (attribute name), and PCTPROCCU (column name).

Procedure Cache Size (KB)

The number of kilobytes (KB) that are allocated for the procedure cache. The SQL server uses the procedure cache to compile queries and store procedures that are compiled. The value format is an integer; for example 1000. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_CACHE_SIZE or PROCCSZ (historical name), *Procedure Cache Size (KB)* (caption), Procedure_Cache_Size (attribute name), and PROCCSZ (column name).

Startup Timestamp

The timestamp that indicates the date and time the SQL server was started. The format is YYYY/MM/DD HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STARTUP_TIMESTAMP or SRVSTARTT (historical name), *Startup Timestamp* (caption), Startup_Timestamp (attribute name), and SRVSTARTT (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Error Log Name

The name of the file that contains the error log for the SQL server. The value format is an alphanumeric string with a maximum of 64 characters; for example, SQL_Error_Log. The error log tracks fatal and kernel errors. It also contains startup and system information. Check the error log for message content and details. This attribute is not available for use in eventing thresholds. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_LOG_NAME or ERRLOG (historical name), Error Log Name (caption), Error_Log_Name (attribute name), and ERRLOG (column name).

Error Log Name (Unicode)

The name of the file that contains the error log for the SQL server (Unicode). This attribute is not available for use in eventing thresholds. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_LOG_NAME_U or UERRLOG (historical name), Error Log Name (Unicode) (caption), Error_Log_Name_U (attribute name), and UERRLOG (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Percent Max Locks

The percentage of locks on resources of the maximum number of locks allowed by the SQL server. The value format is a percentage with two decimal places allowed; for example, 10.00. Setting a higher limit for the maximum number of locks does not impair performance. If your operations exceed the number of available locks, you can increase this limit. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERCENT_MAX_LOCKS or PCTLOCKS (historical name), Percent Max Locks (caption), Percent_Max_Locks (attribute name), and PCTLOCKS (column name).

Procedure Buffers Total

The number of procedure buffers that are in the procedure cache. This number is a fixed number based on the procedure cache that is allocated. Use this attribute to see how much of the cache currently executing procedures use in relation to the cache size as a whole. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_BUFFERS_TOTAL or PROCBTOT (historical name), Procedure Buffers Total (caption), Procedure_Buffers_Total (attribute name), and PROCBTOT (column name).

Procedure Cache Pages

The number of pages that are allocated for the procedure cache. The SQL server uses the procedure cache to compile queries and store procedures that are compiled. The value format is an integer; for example 1000. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_CACHE_PAGES or PROCCSIZ (historical name), Procedure Cache Pages (caption), Procedure_Cache_Pages (attribute name), and PROCCSIZ (column name).

Procedure Cache Percent

The percentage of cache memory the SQL server uses for the procedure cache. This value does not apply to Microsoft SQL Server version 7. The value format is a percentage with two decimal places allowed; for example, 20.00. This attribute is deprecated from the SQL Server agent V6. 2. 1. 1. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_CACHE_PERCENT or PROCCPCT (historical name), Procedure Cache Percent (caption), Procedure_Cache_Percent (attribute name), and PROCCPCT (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates

minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Server Status

Indicates the status of the SQL server. The server status is displayed as Unknown when the collector process of the SQL Server agent is not running. The server status is displayed as Inactive when the collector process of the SQL Server agent is running, but the SQL Server is not responding to the request. The type is string with enumerated values. The following values are defined: Active (Active), Active (1), Inactive (Inactive), Unknown (Unknown), Unknown (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_STATUS or SRVSTATUS (historical name), Server Status (caption), Server_Status (attribute name), and SRVSTATUS (column name).

Server Type

The type of SQL server. The value format is an alphanumeric string; for example, SQL Server. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_TYPE or SRVTYPER (historical name), Server Type (caption), Server_Type (attribute name), and SRVTYPER (column name).

Server Version

The version of the SQL Server. The value format is the version in the format w. x. y. z; for example, 10. 0. 2531. 0. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_SERVER_VER or SQLVER (historical name), Server Version (caption), SQL_Server_Ver (attribute name), and SQLVER (column name).

Server Version (Superceded)

The version of the SQL Server. The value format is the version in the format x. y. z; for example, 10. 0. 253. For the complete version of the SQL Server, see the Server Version attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_VERSION or SRVVERSN (historical name), Server Version (Superceded) (caption), Server_Version (attribute name), and SRVVERSN (column name).

Time Since Startup (Min.)

The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_SINCE_STARTUP or SRVAGE (historical name), Time Since Startup (Min.) (caption), Time_Since_Startup (attribute name), and SRVAGE (column name).

MS SQL Server Enterprise View data set

The MS SQL Server Enterprise View data set contains attributes that you can use to monitor the SQL servers in your enterprise. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Database Max blocks (Unicode)

The name of the database blocking the largest number of processes. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Database Max blocks (Unicode)* (caption), Database_Max_blocks_U (attribute name), and UMAXBLOCK (column name).

Database Max Locks (Unicode)

The name of the database locking the largest number of processes. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Database Max Locks (Unicode)* (caption), Database_Max_Locks_U (attribute name), and UMAXLOCK (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Table Max Locks (Unicode)

The name of the table with largest number of locks. Valid values include text strings with a maximum of 32 bytes. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Table Max Locks (Unicode)* (caption), Table_Max_Locks_U (attribute name), and UTBMXLOCK (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Lock Conflicts

The total number of processes involved in lock conflicts. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Total Lock Conflicts* (caption), Total_Lock_Conflicts (attribute name), and NUMCONFL (column name).

Total Locks

The total number of locks for the server. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Total Locks* (caption), *Total_Locks* (attribute name), and *NUMLOCKS* (column name).

Age of Last Error (Min.)

The age in minutes of the last error encountered. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Age of Last Error (Min.)* (caption), *Age_of_Last_Error* (attribute name), and *LASTERRAGE* (column name).

Collection Status

Indicates the status of the data collector. The data collector is the part of the product that collects information about the SQL server. The type is string with enumerated values. The following values are defined: Inactive (Inactive), Unknown (Unknown), Active (Active). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Collection Status* (caption), *Collection_Status* (attribute name), and *COLLSTATUS* (column name).

Current Logons

The number of logons and connections that are currently active. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Current Logons* (caption), *Current_Logons* (attribute name), and *LOGONCURR* (column name).

Database Max blocks

The name of the database blocking the largest number of processes. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Database Max blocks* (caption), *Database_Max_blocks* (attribute name), and *MAXBLOCK* (column name).

Database Max Locks

The name of the database with largest number of locks. The value format is an alphanumeric string with a maximum of 32 characters. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Database Max Locks* (caption), *Database_Max_Locks* (attribute name), and *MAXLOCK* (column name).

Error Log Size (Bytes)

The size in bytes of the error log. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *Error Log Size (Bytes)* (caption), *Error_Log_Size* (attribute name), and *ERRLOGSZ* (column name).

Host Name

Host on which MS SQLserver resides. The type is string.

The following names are defined for this attribute: *Host Name* (caption), *Host_Name* (attribute name), and *HOSTNAME* (column name).

IO Errors Current Interval

The number of disk errors in the current interval. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: *IO Errors Current Interval* (caption), *IO_Errors_Current_Interval* (attribute name), and *IOERRSCURR* (column name).

Maximum Sev Level

The highest severity level of an error message since startup. The type is string with enumerated values. The following values are defined: No Data (no). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Maximum Sev Level (caption), Maximum_Sev_Level (attribute name), and MAXSEV (column name).

Maximum Sev Timestamp

The timestamp of the of the error message with the highest severity level since startup. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute and SS indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp with enumerated values. The following values are defined: No Data (nodata), N/A (0000000000000001), N/C (0000000000000002), N/P (0000000000000003). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Maximum Sev Timestamp (caption), Maximum_Sev_Timestamp (attribute name), and MAXSEVTIME (column name).

Minimum Pct Data Freespace

The lowest percentage of free space in data only and data and log segments on a device allocated to a database. The value format is a percentage with two decimal places allowed; for example, 5.00. The amount of space needed by the database depends on its anticipated activity. Set alerts for abnormal conditions. Evaluate the possibility of adding more space to the database. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Minimum Pct Data Freespace (caption), Minimum_Pct_Data_Freespace (attribute name), and MINDBFREE (column name).

Minimum Pct Log Freespace

The lowest percentage of free space in the log. The value format is a decimal number 0.00 through 100.00. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Minimum Pct Log Freespace (caption), Minimum_Pct_Log_Freespace (attribute name), and MINLOGFREE (column name).

Pct Max Locks

The percentage of locks on resources of the maximum number of locks allowed by the SQL server. The value format is a percentage with two decimal places allowed; for example, 10.00. Setting a higher limit for the maximum number of locks does not impair performance. If your operations exceed the number of available locks, you can increase this limit. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Pct Max Locks (caption), Percent_Max_Locks (attribute name), and PCTLOCKS (column name).

Pct Max Logons Active

The percentage of the maximum user connections that are currently active. The value format is a decimal number in the range 0.00 to 100.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Pct Max Logons Active (caption), Percent_Max_Logons_Active (attribute name), and LOGONPCT (column name).

Percent Process Block

The percentage of total processes in conflict. The value format is a decimal number in the 0.00 through 100.00. The type is real number (32-bit gauge) with two decimal places of precision with

enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Percent Process Block (caption), Percent_Process_Block (attribute name), and PCTBLOCK (column name).

Physical Reads (per Sec.)

The number of physical reads per second during the current interval. The value format is a decimal number with 2 decimal places. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Physical Reads (per Sec.) (caption), Physical_Reads_per_Second (attribute name), and PHYRDPSEC (column name).

Physical Writes (per Sec.)

The Number of physical writes per second during the current interval. The value is a decimal number with 2 decimal places. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Physical Writes (per Sec.) (caption), Physical_Writes_per_Second (attribute name), and PHYWRPSEC (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the MSSQL Server. The type is string.

The following names are defined for this attribute: Server (caption), Server (attribute name), and SERVERID (column name).

Server CPU Pct System

The percentage of hCPU time the SQL server process is using on the host. The value format is a percentage with two decimal places allowed; for example, 20.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Server CPU Pct System (caption), Server_CPU_Percent_System (attribute name), and SRVCPUPCT (column name).

Server CPU Percent

The percentage of CPU time the SQL server process is using on the host. The value format is a percentage with two decimal places allowed; for example, 20.00. The type is real number (32-bit

gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Server CPU Percent (caption), Server_CPU_Percent (attribute name), and PCTCPU (column name).

Server CPU Percent Application

The percentage of CPU being used by application processes for the server. The value format is a decimal number in the range 0.00 through 100.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Server CPU Percent Application (caption), Server_CPU_Percent_Application (attribute name), and APPCPUPCT (column name).

Server Status

Indicates the status of the SQL server. The server status is displayed as Unknown when the collector process of the SQL Server agent is not running. The server status is displayed as Inactive when the collector process of the SQL Server agent is running, but the SQL Server is not responding to the request. The type is string with enumerated values. The following values are defined: Active (Active), Active (1), Inactive (Inactive), Unknown (Unknown), Unknown (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Server Status (caption), Server_Status (attribute name), and SRVSTATUS (column name).

Server Version

The version of the SQL Server. The value format is the version in the format w.x.y.z; for example, 10.0.2531.0. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Server Version (caption), SQL_Server_Ver (attribute name), and SQLVER (column name).

Server Version (Superceded)

The version of the SQL Server. The value format is the version in the format x.y.z; for example, 10.0.253. For the complete version of the SQL Server, see the Server Version attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Server Version (Superceded) (caption), Server_Version (attribute name), and SRVVERSN (column name).

Table Max Locks

The name of the table with largest number of locks. The type is string.

The following names are defined for this attribute: Table Max Locks (caption), Table_Max_Locks (attribute name), and TBMXLOCK (column name).

Time Since Startup (Min.)

The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Time Since Startup (Min.) (caption), Time_Since_Startup (attribute name), and SRVAGE (column name).

Total Databases

The number of databases for the SQL server. The value format is an integer; for example, 10. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Total Databases (caption), Total_Databases (attribute name), and NUMDBS (column name).

Total Errors High Sev

The total number of messages with a severity of 17 or higher since the server was started. The value format is an integer. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Total Errors High Sev (caption), Total_Errors_High_Sev (attribute name), and NUMSEVG17 (column name).

Total OS CPU Percent

The percentage of CPU time being used by all processes on the host. The value format is a percentage with two decimal places allowed; for example, 40.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Total OS CPU Percent (caption), Total_OS_CPU_Percent (attribute name), and PCTOSCPU (column name).

Total Processes

The total number of processes. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Total Processes (caption), Total_Processes (attribute name), and NUMPROCESS (column name).

Total Processes Blocked

The total number of processes blocked. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: Total Processes Blocked (caption), Total_Processes_Blocked (attribute name), and NUMBLOCK (column name).

MS SQL Server Properties data set

The MS SQL Server Properties data set contains attributes that you can use to monitor the properties of the SQL Server. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Collation ID

The identification number of this SQL Server collation. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COLLATION_ID or COLLATID (historical name), *Collation ID* (caption), Collation_ID (attribute name), and COLLATID (column name).

Collation

The name of the default collation for the server. The value format is a string. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COLLATION (historical name), *Collation* (caption), Collation (attribute name), and COLLATION (column name).

Edition

The installed product edition of this instance of the SQL Server. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EDITION (historical name), *Edition* (caption), Edition (attribute name), and EDITION (column name).

Engine Edition

The database engine edition of the instance of the SQL Server that is currently installed. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENGINE_EDITION or ENGDITION (historical name), *Engine Edition* (caption), Engine_Edition (attribute name), and ENGDITION (column name).

Filestream Level

The current level of FileStream support that is enabled for the SQL Server instance. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1), Disabled (1), Transactional Access (2), Local Full Access (3), Local And Remote Full Access (4). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTREAM_LEVEL or FILSTRLVL (historical name), *Filestream Level* (caption), Filestream_Level (attribute name), and FILSTRLVL (column name).

Filestream Share Name

The Windows share name where the FileStream data is stored. The value format is an alphanumeric string. The type is string with enumerated values. The following values are defined: Not Applicable (Not_applicable), Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILESTREAM_SHARE_NAME or FILSTRSHNM (historical name), *Filestream Share Name* (caption), Filestream_Share_Name (attribute name), and FILSTRSHNM (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), Host_Name (attribute name), and HOSTNAME (column name).

Is Clustered

Indicates whether the SQL Server instance is configured in the Windows Server Failover Clustering (WSFC) cluster. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), No (1), Yes (2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IS_CLUSTERED or ISCLUST (historical name), *Is Clustered* (caption), Is_Clustered (attribute name), and ISCLUST (column name).

Is Single User

The server is currently in the single-user mode, in which only a single user can connect to the server. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), No (1), Yes (2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IS_SINGLE_USER or ISSINUSER (historical name), *Is Single User* (caption), Is_Single_User (attribute name), and ISSINUSER (column name).

IsIntegrated Security Only

The server is in the integrated security mode. The value format is an integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), No (1), Yes (2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ISINTEGRATED_SECURITY_ONLY or ISINTSEC (historical name), *IsIntegrated Security Only* (caption), IsIntegrated_Security_Only (attribute name), and ISINTSEC (column name).

License Type

The mode of this instance of the SQL Server. The mode can be per-seat or per-processor. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LICENSE_TYPE or LICTYPE (historical name), *License Type* (caption), License_Type (attribute name), and LICTYPE (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Num Licenses

The number of client licenses that are currently registered for this instance of the SQL Server if the SQL Server is in the per-seat mode. The number of processors that are currently licensed for this instance of the SQL Server if the SQL Server is in the per-processor mode. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NUM_LICENSES or NUMLIC (historical name), *Num Licenses* (caption), Num_Licenses (attribute name), and NUMLIC (column name).

Physical NetBIOS Name

The NetBIOS name of the machine where this instance of the SQL Server is currently running. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PHYSICAL_NETBIOS_NAME or NETBIONAME (historical name), *Physical NetBIOS Name* (caption), Physical_NetBIOS_Name (attribute name), and NETBIONAME (column name).

Product Level

The level of the version of this instance of the SQL Server, for example, the original release version and the service pack version. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1), CTP (1), RTM (2), SP1 (3), SP2 (4), SP3 (5), SP4 (6), SP5 (7), SP6 (8), SP7 (9), SP8 (10), SP9 (11), SP10 (12). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PRODUCT_LEVEL or PRODLEVEL (historical name), *Product Level* (caption), Product_Level (attribute name), and PRODLEVEL (column name).

Qualified Server Name

The name of the qualified SQL Server instance in the format hostname:SQL server instance. For the default SQL Server instance, the format is hostname. The value format is an alphanumeric string with a maximum of 128 characters. The type is string with enumerated values. The following values are defined: Not Collected (Not_collected). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUALIFIED_SERVER_NAME or QUALISERV (historical name), *Qualified Server Name* (caption), Qualified_Server_Name (attribute name), and QUALISERV (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Edition ID

The identification number that represents the installed product edition of this instance of the SQL Server. The type is integer (64-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EDITION_ID or EDITIONID (historical name), Edition ID (caption), Edition_ID (attribute name), and EDITIONID (column name).

Fully Qualified Domain Name

The fully qualified DNS name that uniquely identifies the local computer. If the local computer is a node in a cluster, it provides the fully qualified DNS name of the cluster virtual server. The type is string.

The following names are defined for this attribute: FQDN_NAME or FQDN (historical name), Fully Qualified Domain Name (caption), FQDN_Name (attribute name), and FQDN (column name).

HADR Manager Status

Indicates whether the HADR manager for the AlwaysOn availability group has started. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1), Communication Pending (1), Running (2), Failed (3), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: HADR_MANAGER_STATUS or HADRMGST (historical name), HADR Manager Status (caption), HADR_Manager_Status (attribute name), and HADRMGST (column name).

Instance Path

Path of the sql server instance. Indicates the path of the SQL Server instance. The type is string.

The following names are defined for this attribute: INSTANCE_PATH or INSTPATH (historical name), Instance Path (caption), Instance_Path (attribute name), and INSTPATH (column name).

Instance Port

Indicates the port of the SQL Server instance. The type is integer (32-bit numeric property).

The following names are defined for this attribute: INSTANCE_PORT or INSTPORT (historical name), Instance Port (caption), Instance_Port (attribute name), and INSTPORT (column name).

Is HADR Enabled

Indicates whether the AlwaysOn Availability Group feature is enabled on the SQL Server instance. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1), No (1), Yes (2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IS_HADR_ENABLED or ISHADRENB (historical name), Is HADR Enabled (caption), Is_HADR_Enabled (attribute name), and ISHADRENB (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Indicate a *TEXT: Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server Version

The version of the SQL Server. The value format is the version in the format w. x. y. z; for example, 10.0.2531.0. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_VERSION or SQLVER (historical name), Server Version (caption), Server_Version (attribute name), and SQLVER (column name).

MS SQL Server Summary data set

The MS SQL Server Summary data set contains attributes that you can use to monitor summary information for SQL servers. This group of data represents the SQL system identification, revision level and status information. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Cache Average Free Scan

The average number of buffers scanned by the LazyWriter when the LazyWriter searches the data cache for an unused buffer to replenish the free buffer pool. If Microsoft SQL Server must read a large number of buffers to find a free one, server performance might degrade. A low value indicates optimal performance. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CACHE_AVERAGE_FREE_SCAN or CACHAVGFCN (historical name), *Cache Average Free Scan* (caption), Cache_Average_Free_Scan (attribute name), and CACHAVGFCN (column name).

Cache Hit Ratio

The current ratio of data cache hits to total requests. Use this attribute to check the effectiveness of the data cache. For optimal performance, the returned value must be approximately 90% or greater. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CACHE_HIT_RATIO or CACHHITR (historical name), *Cache Hit Ratio* (caption), Cache_Hit_Ratio (attribute name), and CACHHITR (column name).

Cache Maximum Free Page Scan

The maximum value for the number of buffers scanned by the LazyWriter when the LazyWriter searches the data cache for an unused buffer to replenish the free buffer pool. Use this attribute to check the effectiveness of the data cache. The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CACHE_MAXIMUM_FREE_PAGE_SCAN or CACHMAXFP (historical name), *Cache Maximum Free Page Scan* (caption), Cache_Maximum_Free_Page_Scan (attribute name), and CACHMAXFP (column name).

Collection Status

Indicates the status of the data collector on a remote node. The data collector is the part of the product that collects information about the SQL server. The type is string with enumerated values. The following values are defined: Inactive (Inactive), Unknown (Unknown), Active (Active). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: COLLECTION_STATUS or COLLSTATUS (historical name), *Collection Status* (caption), Collection_Status (attribute name), and COLLSTATUS (column name).

Data Cache Size (KB)

The number of kilobytes (KB) allocated for the data cache memory. The SQL server uses the data cache to store data and index pages. The value format is an integer; for example, 1000. The cache is sometimes referred to as the buffer cache. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATA_CACHE_SIZE or DATACHSZ (historical name), *Data Cache Size (KB)* (caption), Data_Cache_Size (attribute name), and DATACHSZ (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), Host_Name (attribute name), and HOSTNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Page Life Expectancy

The duration (in seconds) for which an SQL Server block or page is stored in the memory. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_LIFE_EXPECTANCY or PGLFEXPECT (historical name), *Page Life Expectancy* (caption), Page_Life_Expectancy (attribute name), and PGLFEXPECT (column name).

Pct Max Locks

The percentage of locks on resources of the maximum number of locks allowed by the SQL server. The value format is a percentage with two decimal places allowed; for example, 10.00. Setting a higher limit for the maximum number of locks does not impair performance. If your operations exceed the number of available locks, you can increase this limit. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERCENT_MAX_LOCKS or PCTLOCKS (historical name), *Pct Max Locks* (caption), Percent_Max_Locks (attribute name), and PCTLOCKS (column name).

Procedure Cache Size (KB)

The number of kilobytes (KB) allocated for the procedure cache. The SQL server uses the procedure cache to compile queries and store procedures that are compiled. The value format is an integer; for example, 1000. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PROCEDURE_CACHE_SIZE or PROCCHSZ (historical name), *Procedure Cache Size (KB)* (caption), Procedure_Cache_Size (attribute name), and PROCCHSZ (column name).

Server CPU Pct IO

The percentage of time used for I/O operations during the current monitoring interval. Use this attribute to gauge how much of the CPU resource the database server uses for I/O so you can allocate resources more efficiently. You also can use this attribute to determine how I/O resource-intensive certain operations are. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_CPU_PERCENT_IO or PCTCPUIO (historical name), *Server CPU Pct IO* (caption), Server_CPU_Percent_IO (attribute name), and PCTCPUIO (column name).

Server CPU Percent

The percentage of CPU time the SQL server process is using on the host. The value format is a percentage with two decimal places allowed; for example, 20.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_CPU_PERCENT or PCTCPU (historical name), *Server CPU Percent* (caption), Server_CPU_Percent (attribute name), and PCTCPU (column name).

Server Status

Indicates the status of the SQL server. The server status is displayed as Unknown when the collector process of the SQL Server agent is not running. The server status is displayed as Inactive when the collector process of the SQL Server agent is running, but the SQL Server is not responding to the request. The type is string with enumerated values. The following values are defined: Active (Active), Active (1), Inactive (Inactive), Unknown (Unknown), Unknown (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_STATUS or SRVSTATUS (historical name), *Server Status* (caption), Server_Status (attribute name), and SRVSTATUS (column name).

Server Version

The version of the SQL Server. The value format is the version in the format w. x. y. z; for example, 10. 0. 2531. 0. The type is string with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_SERVER_VER or SQLVER (historical name), *Server Version* (caption), SQL_Server_Ver (attribute name), and SQLVER (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), *Server* (caption), Server (attribute name), and SERVERID (column name).

SQLServerAgent Failed Jobs

Reports any jobs run by the SQLServerAgent service that have failed in the last monitoring interval. The first time you retrieve this attribute, it returns 0. The next time you retrieve it, it returns the total number of SQLServerAgent failed jobs found in the system history tables since the first run. These jobs include replication and user-defined jobs, such as maintenance or backup tasks. Use this attribute to alert you when scheduled tasks have failed. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQLSERVERAGENT_FAILED_JOBS or SQLFAILJ (historical name), *SQLServerAgent Failed Jobs* (caption), SQLServerAgent_Failed_Jobs (attribute name), and SQLFAILJ (column name).

Time Since Startup (Min.)

The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TIME_SINCE_STARTUP or SRVAGE (historical name), *Time Since Startup (Min.)* (caption), Time_Since_Startup (attribute name), and SRVAGE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Blocking Lock Requests

The total number of current locks blocking other processes. A blocking lock must be released before other processes requesting competing locks can progress. (For more information on locks, see the Microsoft SQL Server Books online.) Use this attribute to show the number of blocking locks active during server activity. This attribute can indicate that processes are being held up through lock contention rather than hardware performance issues. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_BLOCKING_LOCK_REQUESTS or BLKLCKREQ (historical name), *Total Blocking Lock Requests* (caption), Total_Blocking_Lock_Requests (attribute name), and BLKLCKREQ (column name).

Total Locks Remaining

The total number of locks that can still be taken out. The maximum number of locks is configurable. In Microsoft SQL Server Version 8.0, you can have the number of locks configured automatically. Use this attribute to show the number of locks active during server activity. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_LOCKS_REMAINING or LOCKSREM (historical name), *Total Locks Remaining* (caption), Total_Locks_Remaining (attribute name), and LOCKSREM (column name).

Total OS CPU Percent

The percentage of CPU time being used by all processes on the host. The value format is a percentage with two decimal places allowed; for example, 40.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_OS_CPU_PERCENT or PCTOSCPU (historical name), *Total OS CPU Percent* (caption), Total_OS_CPU_Percent (attribute name), and PCTOSCPU (column name).

Total Server Memory (KB)

The total amount of dynamic memory (in KB) that the server is using currently. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SERVER_MEMORY or TOTSVRMEM (historical name), *Total Server Memory (KB)* (caption), Total_Server_Memory (attribute name), and TOTSVRMEM (column name).

Cache Free Buffers

The current number of cache buffers in the free buffer pool. Use this attribute as an alert when the number of buffers is getting low. A low number of free buffers might degrade server performance. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following

values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CACHE_FREE_BUFFERS or CACHFREEB (historical name), Cache Free Buffers (caption), Cache_Free_Buffers (attribute name), and CACHFREEB (column name).

Checkpoint Pages per Sec

The number of pages flushed to disk per second by a checkpoint or other operation that require all dirty pages to be flushed. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CHECKPOINT_PAGES_PER_SEC or CKPTPGSEC (historical name), Checkpoint Pages per Sec (caption), Checkpoint_Pages_per_Sec (attribute name), and CKPTPGSEC (column name).

CPU Idle Delta

The difference in the CPU Idle between the current and the previous samples. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_IDLE_DELTA or CPUIDLEDLT (historical name), CPU Idle Delta (caption), CPU_Idle_Delta (attribute name), and CPUIDLEDLT (column name).

CPU Idle Sec

The time (in seconds) that the SQL Server has been idle since it was last started. The value is in seconds even though the column name would imply milliseconds. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_IDLE_SEC or CPUIDLEMS (historical name), CPU Idle Sec (caption), CPU_Idle_Sec (attribute name), and CPUIDLEMS (column name).

CPU Pct Idle

The percentage of time that a database server has been idle during the current monitoring interval. Use this attribute to gauge how much of the CPU resource the database server uses so that you can allocate resources more efficiently. Also, you can use this attribute to determine how resource-intensive certain operations are. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_PERCENT_IDLE or PCTCPUIDLE (historical name), CPU Pct Idle (caption), CPU_Percent_Idle (attribute name), and PCTCPUIDLE (column name).

CPU Usage Delta

The difference in the CPU Usage between the current and the previous samples. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_USAGE_DELTA or CPUUSAGDLT (historical name), CPU Usage Delta (caption), CPU_Usage_Delta (attribute name), and CPUUSAGDLT (column name).

CPU Usage Sec

The time (in seconds) that the CPU has spent working since the SQL Server was last started. The value is in seconds even though the column name would imply milliseconds. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_USAGE_SEC or CPUUSAGEMS (historical name), CPU Usage Sec (caption), CPU_Usage_Sec (attribute name), and CPUUSAGEMS (column name).

Current Interval (Sec.)

The number of seconds that have elapsed between the previous sample and the current sample. The value format is an integer; for example, 90. A sample contains the data that IBM Tivoli Monitoring for Microsoft SQL Server collects about each SQL server. New data becomes available if a new interval has occurred and data has been refreshed at the CMS hub. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_INTERVAL or CURRINTL (historical name), Current Interval (Sec.) (caption), Current_Interval (attribute name), and CURRINTL (column name).

Lazy Writes per Sec

The number of buffers written per second by the buffer manager lazy writer. The lazy writer is a system process that flushes out batches of dirty, aged buffers and make them available to user processes. The lazy writer eliminates the need to perform frequent checkpoints to create available buffers. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAZY_WRITES_PER_SEC or LAZYWRTSEC (historical name), Lazy Writes per Sec (caption), Lazy_Writes_per_Sec (attribute name), and LAZYWRTSEC (column name).

Mixed Page Allocations Per Sec

The number of pages that are allocated from mixed extents per second. These pages can be used for storing the Index Allocation Map (IAM) pages. The first eight pages from mixed extents are allocated to an allocation unit. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MIXED_PAGE_ALLOCATIONS_PER_SEC or MIXPGALCPS (historical name), Mixed Page Allocations Per Sec (caption), Mixed_Page_Allocations_Per_Sec (attribute name), and MIXPGALCPS (column name).

Page Reads per Sec

The number of physical database page reads that are issued per second. This statistic displays the total number of physical page reads across all databases. Because physical IO is expensive, you might be able to minimize the cost, either by using a larger data cache, intelligent indexes, and more efficient queries, or by changing the database design. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_READS_PER_SEC or PAGERDSSEC (historical name), Page Reads per Sec (caption), Page_Reads_per_Sec (attribute name), and PAGERDSSEC (column name).

Page Splits per Sec

The number of page splits that occur per second because the space in an index page is insufficient. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_SPLITS_PER_SEC or PGSPLTSSC (historical name), Page Splits per Sec (caption), Page_Splits_per_Sec (attribute name), and PGSPLTSSC (column name).

Page Writes per Sec

The number of physical database page writes issued. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PAGE_WRITES_PER_SEC or PAGEWRTSEC (historical name), Page Writes per Sec (caption), Page_Writes_per_Sec (attribute name), and PAGEWRTSEC (column name).

Plan Cache Hit Ratio

The ratio between the cache hits and the cache lookups. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PLAN_CACHE_HIT_RATIO or PCACHHITR (historical name), Plan Cache Hit Ratio (caption), Plan_Cache_Hit_Ratio (attribute name), and PCACHHITR (column name).

Readahead Pages per Sec

The number of pages read in anticipation of use. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: READAHEAD_PAGES_PER_SEC or RDAHDPGSEC (historical name), Readahead Pages per Sec (caption), Readahead_Pages_per_Sec (attribute name), and RDAHDPGSEC (column name).

Repl Dist Delivered Cmds per Sec

The number of distribution commands delivered per second to the Subscriber. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

REPL_DISTRIBUTION_DELIVERED_CMDS_PER_SEC or RDDELCMDS (historical name), Repl Dist Delivered Cmds per Sec (caption), Repl_Distribution_Delivered_Cmds_per_Sec (attribute name), and RDDELCMDS (column name).

Repl Dist Delivered Trans per Sec

The number of distribution transactions delivered per second to the Subscriber. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

REPL_DISTRIBUTION_DELIVERED_TRANS_PER_SEC or RDDELTRANS (historical name), Repl Dist Delivered Trans per Sec (caption), Repl_Distribution_Delivered_Trans_per_Sec (attribute name), and RDDELTRANS (column name).

Repl Dist Delivery Latency

The distribution latency (in MS), the time it takes for transactions to be delivered to the Distributor and applied at the Subscriber. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPL_DISTRIBUTION_DELIVERY_LATENCY or RDDELLATEN (historical name), Repl Dist Delivery Latency (caption), Repl_Distribution_Delivery_Latency (attribute name), and RDDELLATEN (column name).

Repl Logreader Delivered Cmds per Sec

The number of Log Reader Agent commands delivered per second to the Distributor. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPL_LOGREADER_DELIVERED_CMDS_PER_SEC or RLDELCMDS (historical name), Repl Logreader Delivered Cmds per Sec (caption), Repl_Logreader_Delivered_Cmds_per_Sec (attribute name), and RLDELCMDS (column name).

Repl Logreader Delivered Trans per Sec

The number of Log Reader Agent transactions delivered per second to the Distributor. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPL_LOGREADER_DELIVERED_TRANS_PER_SEC or RLDELTRANS (historical name), Repl Logreader Delivered Trans per Sec (caption), Repl_Logreader_Delivered_Trans_per_Sec (attribute name), and RLDELTRANS (column name).

Repl Logreader Delivery Latency

The current amount of time (in MS) elapsed from when transactions are applied at the Publisher to when they are delivered to the Distributor. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPL_LOGREADER_DELIVERY_LATENCY or RLDELLATEN (historical name), Repl Logreader Delivery Latency (caption), Repl_Logreader_Delivery_Latency (attribute name), and RLDELLATEN (column name).

Repl Merge Conflicts per Sec

The number of conflicts per second that occurred in the Publisher or Subscriber upload and download. If the value is not zero, the value might require notifying the losing side, overriding the conflict, and so on. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPL_MERGE_CONFLICTS_PER_SEC or RMCONFLCTS (historical name), Repl Merge Conflicts per Sec (caption), Repl_Merge_Conflicts_per_Sec (attribute name), and RMCONFLCTS (column name).

Repl Merge Downloaded Changes per Sec

The number of rows merged (inserted, updated, and deleted) per second from the Publisher to the Subscriber. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

REPL_MERGE_DOWNLOADED_CHANGES_PER_SEC or RMDNLDCHGS (historical name), Repl Merge Downloaded Changes per Sec (caption), Repl_Merge_Downloaded_Changes_per_Sec (attribute name), and RMDNLDCHGS (column name).

Repl Merge Uploaded Changes per Sec

The number of rows merged (inserted, updated, and deleted) per second from the Subscriber to the Publisher. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REPL_MERGE_UPLOADED_CHANGES_PER_SEC or RMUPLDCHGS (historical name), Repl Merge Uploaded Changes per Sec (caption), Repl_Merge_Uploaded_Changes_per_Sec (attribute name), and RMUPLDCHGS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product

collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server Type

The type of SQL server. The value format is an alphanumeric string; for example, SQL Server. The type is string.

The following names are defined for this attribute: SERVER_TYPE or SRVTYPE (historical name), Server Type (caption), Server_Type (attribute name), and SRVTYPE (column name).

Server Version (Superceded)

The version of the SQL Server. The value format is the version in the format x. y. z; for example, 10. 0. 253. For the complete version of the SQL Server, see the Server Version attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVER_VERSION or SRVVERSN (historical name), Server Version (Superceded) (caption), Server_Version (attribute name), and SRVVERSN (column name).

Stolen Pages

The number of pages used for miscellaneous server purposes. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STOLEN_PAGES or STOLENPGS (historical name), Stolen Pages (caption), Stolen_Pages (attribute name), and STOLENPGS (column name).

Stolen Pages Growth

The growth of the number of stolen pages between the current sample and previous sample. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STOLEN_PAGES_GROWTH or STOLENPGGR (historical name), Stolen Pages Growth (caption), Stolen_Pages_Growth (attribute name), and STOLENPGGR (column name).

Total Pages

The total number of stolen, free, and database pages in the buffer pool. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_PAGES or TOTPAGE (historical name), Total Pages (caption), Total_Pages (attribute name), and TOTPAGE (column name).

MS SQL Server Transactions Summary data set

The MS SQL Server Transactions Summary data set contains attributes that you can use to monitor the transaction summary information for SQL Servers. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Active Transactions

The number of transaction enlistments (local, dtc, and bound) that are currently active. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ACTIVE_TRANSACTIONS or TACTTRN (historical name), *Total Active Transactions* (caption), Total_Active_Transactions (attribute name), and TACTTRN (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Longest Transaction Running Time (Sec.)

The transaction that is currently active for the longest duration of time (in seconds). The value format is an integer. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LONGEST_TRANSACTION_RUNNING_TIME or LNGTRNRNTM (historical name), Longest Transaction Running Time (Sec.) (caption), Longest_Transaction_Running_Time (attribute name), and LNGTRNRNTM (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Tempdb Free Space (KB)

The amount of free space (in KB) that is currently available in the tempdb database. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TEMPDB_FREE_SPACE or TEMDBFRESP (historical name), Tempdb Free Space (KB) (caption), Tempdb_Free_Space (attribute name), and TEMDBFRESP (column name).

Total Non-Snapshot Version Transactions

The number of currently active non-snapshot transactions that generate version records. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_NON-SNAPSHOT_VERSION_TRANSACTIONS or TNSNAPVTRN (historical name), Total Non-Snapshot Version Transactions (caption), Total_Non-Snapshot_Version_Transactions (attribute name), and TNSNAPVTRN (column name).

Total Snapshot Transactions

The number of snapshot transactions that are currently active. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SNAPSHOT_TRANSACTIONS or TSNAPTRN (historical name), Total Snapshot Transactions (caption), Total_Snapshot_Transactions (attribute name), and TSNAPTRN (column name).

Total Update Snapshot Transactions

The number of currently active snapshot transactions that perform update operations. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_UPDATE_SNAPSHOT_TRANSACTIONS or TUPSNPTRN (historical name), Total Update Snapshot Transactions (caption), Total_Update_Snapshot_Transactions (attribute name), and TUPSNPTRN (column name).

Update Conflict Ratio

The percentage of snapshot transactions that currently have update conflicts. The value format is an integer. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: UPDATE_CONFLICT_RATIO or UPCONRTO (historical name), Update Conflict Ratio (caption), Update_Conflict_Ratio (attribute name), and UPCONRTO (column name).

Version Cleanup Rate (KB/sec)

The current rate (in KB per second) at which versions are removed from the version store. The value format is an integer. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_CLEANUP_RATE or VERCLNRT (historical name), Version Cleanup Rate (KB/sec) (caption), Version_Cleanup_Rate (attribute name), and VERCLNRT (column name).

Version Generation Rate (KB/sec)

The current rate (in KB per second) at which versions are generated in the version store. The value format is an integer. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_GENERATION_RATE or VERGENRT (historical name), Version Generation Rate (KB/sec) (caption), Version_Generation_Rate (attribute name), and VERGENRT (column name).

Version Store Size

The amount of space (in KB) that is currently available in the version store. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_STORE_SIZE or VERSTRSZ (historical name), Version Store Size (caption), Version_Store_Size (attribute name), and VERSTRSZ (column name).

Version Store Unit Count

The number of allocation units that are currently active in the version store. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_STORE_UNIT_COUNT or VERSTRUNCT (historical name), Version Store Unit Count (caption), Version_Store_Unit_Count (attribute name), and VERSTRUNCT (column name).

Version Store Unit Creation

The number of units that are currently created in the version store. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_STORE_UNIT_CREATION or VERSTRUNCR (historical name), Version Store Unit Creation (caption), Version_Store_Unit_Creation (attribute name), and VERSTRUNCR (column name).

Version Store Unit Truncation

The number of units that are currently truncated in the version store. The value format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VERSION_STORE_UNIT_TRUNCATION or VERSTRUNTR (historical name), Version Store Unit Truncation (caption), Version_Store_Unit_Truncation (attribute name), and VERSTRUNTR (column name).

MS SQL Service Broker Activation data set

The MS SQL Service Broker Activation data set monitors the details of the Service Broker activation for each database in an SQL Server instance.

This data set contains the following attributes:

Database Name (Unicode)

The name of the database in the Unicode format. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME_U or UDBNAME (historical name), Database Name (Unicode) (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Stored Procedure Invoked Per sec

The number of stored procedures that are currently called per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STORED_PROCEDURE_INVOKED_PER_SEC or STPROCINVS (historical name), Stored Procedure Invoked Per sec (caption), Stored_Procedure_Invoked_Per_sec (attribute name), and STPROCINVS (column name).

Task Limit Reached Per sec

The number of times that the task limit that is currently activated for a queue has been reached per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TASK_LIMIT_REACHED_PER_SEC or TSKLMTRCHS (historical name), Task Limit Reached Per sec (caption), Task_Limit_Reached_Per_sec (attribute name), and TSKLMTRCHS (column name).

Tasks Aborted Per sec

The number of activated tasks that are currently ended per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TASKS_ABORTED_PER_SEC or TSKABORS (historical name), Tasks Aborted Per sec (caption), Tasks_Aborted_Per_sec (attribute name), and TSKABORS (column name).

Tasks Started Per sec

The number of activated tasks that are currently being started per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TASKS_STARTED_PER_SEC or TSKSTRTS (historical name), Tasks Started Per sec (caption), Tasks_Starts_Per_sec (attribute name), and TSKSTRTS (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Task Limit Reached

The number of times that the task limit that is currently activated for a queue has been reached. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_TASK_LIMIT_REACHED or TTSKLMTRCH (historical name), Total Task Limit Reached (caption), Total_Task_Limit_Reached (attribute name), and TTSKLMTRCH (column name).

Total Tasks Started

The total number of activated tasks that are currently started. The type is integer (32-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_TASKS_STARTED or TTSKSTRT (historical name), Total Tasks Started (caption), Total_Tasks_Start (attribute name), and TTSKSTRT (column name).

MS SQL Service Broker Statistics data set

The MS SQL Service Broker Statistics data set displays the summary of the Service Broker statistics.

This data set contains the following attributes:

Broker Transaction Rollbacks

The current number of transactions related to the Service Broker that have been rolled back. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BROKER_TRANSACTION_ROLLBACKS or BRTRANRLBK (historical name), Broker Transaction Rollbacks (caption), Broker_Transaction_Rollbacks (attribute name), and BRTRANRLBK (column name).

Enqueued Local Messages Per Sec

The number of messages per second from local endpoints that are currently being delivered into the queues. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENQUEUED_LOCAL_MESSAGES_PER_SEC or ENQLMSPSEC (historical name), Enqueued Local Messages Per Sec (caption), Enqueued_Local_Messages_Per_Sec (attribute name), and ENQLMSPSEC (column name).

Enqueued Messages Per Sec

The number of messages from local endpoints and the transport that are currently delivered into the queues per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENQUEUED_MESSAGES_PER_SEC or ENQMSGPSEC (historical name), Enqueued Messages Per Sec (caption), Enqueued_Messages_Per_Sec (attribute name), and ENQMSGPSEC (column name).

Enqueued Transport Messages Per Sec

The number of messages from the transport that are being delivered into the queues per second. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENQUEUED_TRANSPORT_MESSAGES_PER_SEC or ENQTRMPSEC (historical name), Enqueued Transport Messages Per Sec (caption), Enqueued_Transport_Messages_Per_Sec (attribute name), and ENQTRMPSEC (column name).

Forwarded Messages Discarded Per Sec

The number of forwarded messages that are currently discarded per second because of memory limits for forwarded messages, age limits, and so on. The type is real number (64-bit gauge) with two

decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FORWARDED_MESSAGES_DISCARDED_PER_SEC or FMSGDCPSEC (historical name), Forwarded Messages Discarded Per Sec (caption), Forwarded_Messages_Discarded_Per_Sec (attribute name), and FMSGDCPSEC (column name).

Forwarded Pending Messages

The number of forwarded messages that are currently in the queue, and that are not yet sent. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FORWARDED_PENDING_MESSAGES or FRPENDMSG (historical name), Forwarded Pending Messages (caption), Forwarded_Pending_Messages (attribute name), and FRPENDMSG (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

SQL Receives Per Sec

The number of SQL RECEIVE commands that are processed per second by the Service Broker. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_RECEIVES_PER_SEC or SQLRCVPSEC (historical name), SQL Receives Per Sec (caption), SQL_Receives_Per_Sec (attribute name), and SQLRCVPSEC (column name).

SQL Sends Per Sec

The number of SQL SEND commands that are processed per second by the Service Broker. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SQL_SENDS_PER_SEC or SQLSENPSEC (historical name), SQL Sends Per Sec (caption), SQL_Sends_Per_Sec (attribute name), and SQLSENPSEC (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Enqueued Local Messages

The total number of message fragments from the local endpoints that are currently being delivered into the queues. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ENQUEUED_LOCAL_MESSAGES or TOTENQLMSG (historical name), Total Enqueued Local Messages (caption), Total_Enqueued_Local_Messages (attribute name), and TOTENQLMSG (column name).

Total Enqueued Messages

The total number of message fragments from the local endpoints and the transport that are currently being delivered into the queues. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ENQUEUED_MESSAGES or TOTENQMSG (historical name), Total Enqueued Messages (caption), Total_Enqueued_Messages (attribute name), and TOTENQMSG (column name).

Total Enqueued Transport Messages

The total number of messages from the transport that are currently delivered into the queues. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_ENQUEUED_TRANSPORT_MESSAGES or TOTENQTMMSG (historical name), Total Enqueued Transport Messages (caption), Total_Enqueued_Transport_Messages (attribute name), and TOTENQTMMSG (column name).

Total Forwarded Messages

The total number of forwarded messages from the local endpoints and the transport that are currently being sent. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_FORWARDED_MESSAGES or TOTFRMSG (historical name), Total Forwarded Messages (caption), Total_Forwarded_Messages (attribute name), and TOTFRMSG (column name).

Total Forwarded Messages Discarded

The total number of forwarded messages that are currently discarded because of memory limits for forwarded messages, age limits, and so on. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_FORWARDED_MESSAGES_DISCARDED or TOTFMSGDSC (historical name), Total Forwarded Messages Discarded (caption), Total_Forwarded_Messages_Dispatched (attribute name), and TOTFMSGDSC (column name).

Total SQL Receives

The total number of SQL RECEIVE commands that are currently processed by the Service Broker. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SQL_RECEIVES or TOTSQLRECV (historical name), Total SQL Receives (caption), Total_SQL_Receives (attribute name), and TOTSQLRECV (column name).

Total SQL Sends

The total number of SQL SEND commands that are currently processed by the Service Broker. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_SQL_SENDS or TOTSQLSEND (historical name), Total SQL Sends (caption), Total_SQL_Sends (attribute name), and TOTSQLSEND (column name).

MS SQL Service Broker Transport data set

The MS SQL Service Broker Transport data set monitors the details of the Service Broker transport for the SQL Server.

This data set contains the following attributes:

Avg Size Message Fragments Received

The average byte size of the message fragments that are received by the transport receive I/O operations. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVG_SIZE_MESSAGE_FRAGMENTS_RECEIVED or SZMSGFRREC (historical name), Avg Size Message Fragments Received (caption), Avg_Size_Message_Fragments_Received (attribute name), and SZMSGFRREC (column name).

Avg Size Message Fragments Sent

The average byte size of the message fragments that are sent by the transport send I/O operations. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVG_SIZE_MESSAGE_FRAGMENTS_SEND or SZMSGFRSND (historical name), Avg Size Message Fragments Sent (caption), Avg_Size_Message_Fragments_Send (attribute name), and SZMSGFRSND (column name).

Current Bytes For Receive I/O

The number of bytes that are to be read by the transport receive I/O operations that are currently running. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_BYTES_RECEIVED_IO or CURBYTRECV (historical name), Current Bytes For Receive I/O (caption), Current_Bytes_Received_IO (attribute name), and CURBYTRECV (column name).

Current Bytes For Send I/O

The number of buffer bytes that are to be read by the transport send I/O operations that are currently running. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_BYTES_SEND_IO or CURBYTSEND (historical name), Current Bytes For Send I/O (caption), Current_Bytes_Send_IO (attribute name), and CURBYTSEND (column name).

Current Message Fragments For Send I/O

The number of message fragments that are to be read by the transport send I/O operations that are currently running. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CURRENT_MESSAGE_FRAGMENTS_SEND_IO or MSGFRGSEND (historical name), Current Message Fragments For Send I/O (caption), Current_Message_Fragments_Send_IO (attribute name), and MSGFRGSEND (column name).

Message Fragment Received Per Sec

The number of message fragments that are currently received per second in the transport receive I/O operations. The type is real number (64-bit gauge) with two decimal places of precision with

enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE_FRAGMENT_RECEIVED_PER_SEC or MGFRRECSEC (historical name), Message Fragment Received Per Sec (caption), Message_Fragment_Received_Per_Sec (attribute name), and MGFRRECSEC (column name).

Message Fragments Sent Per Sec

The number of message fragments that are currently sent per second in the transport send I/O operations. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE_FRAGMENTS_SEND_PER_SEC or MGFRSNDSSEC (historical name), Message Fragments Sent Per Sec (caption), Message_Fragments_Send_Per_Sec (attribute name), and MGFRSNDSSEC (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Pending Bytes For Receive I/O

The number of bytes that are currently present in the completed transport receive I/O operations whose message fragments are not yet sent to the queue, or are not yet discarded. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PENDING_BYTES_RECEIVED_IO or PENDBYREC (historical name), Pending Bytes For Receive I/O (caption), Pending_Bytes_Received_IO (attribute name), and PENDBYREC (column name).

Pending Bytes For Send I/O

The current number of buffer bytes of the message fragments that are being sent to the queue, and that are ready to be sent by the transport send I/O operations. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PENDING_BYTES_SEND_IO or PENDBYSND (historical name), Pending Bytes For Send I/O (caption), Pending_Bytes_Send_IO (attribute name), and PENDBYSND (column name).

Pending Message Fragments For Receive I/O

The current number of message fragments that are received by transport receive I/O operations, and that are not yet sent to the queue or are not yet discarded. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PENDING_MESSAGE_FRAGMENTS_RECEIVED_IO or PNDMSGREC (historical name), Pending Message Fragments For Receive I/O (caption), Pending_Message_Fragments_Received_IO (attribute name), and PNDMSGREC (column name).

Pending Message Fragments For Send I/O

The current number of message fragments that are being sent to the queue, and that are ready to be sent through the transport layer. The type is integer (64-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PENDING_MESSAGE_FRAGMENTS_SEND_IO or PNDMSGSEND (historical name), Pending Message Fragments For Send I/O (caption), Pending_Message_Fragments_Send_IO (attribute name), and PNDMSGSEND (column name).

Receive I/Os Per Sec

The current number of transport receive I/O operations per second. A transport receive I/O might contain more than one message fragment. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: RECEIVE_IOS_PER_SEC or RECIOSSEC (historical name), Receive I/Os Per Sec (caption), Receive_IOS_Per_Sec (attribute name), and RECIOSSEC (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Send I/Os Per Sec

The current number of transport send I/O operations per second. A transport send I/O operation might contain more than one message fragment. The type is real number (64-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SEND_IOS_PER_SEC or SNDIOSEC (historical name), Send I/Os Per Sec (caption), Send_IOS_Per_Sec (attribute name), and SNDIOSEC (column name).

Server

The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Open Connection Count

The total number of transport connections that are currently open. The type is integer (64-bit counter) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_OPEN_CONNECTION_COUNT or TOTOPCONN (historical name), Total Open Connection Count (caption), Total_Open_Connection_Count (attribute name), and TOTOPCONN (column name).

MS SQL Services Detail data set

The MS SQL Services Detail data set provides details about the services that make up the SQL Server. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Service Name

The service name for the SQL Server. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: SERVICE_NAME or SRVCNAME (historical name), *Service Name* (caption), Service_Name (attribute name), and SRVCNAME (column name).

Service Start Mode

Defined start mode for the service. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Automatic (1), Manual (2), Disabled (3), Unknown (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVICE_START_MODE or STARTMODE (historical name), *Service Start Mode* (caption), Service_Start_Mode (attribute name), and STARTMODE (column name).

Service State

Current service state. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Stopped (0), Start Pending (1), Stop Pending (2), Running (3), Continue Pending (4), Pause Pending (5), Paused (6), Unknown (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVICE_STATE or SRVCSTATE (historical name), *Service State* (caption), Service_State (attribute name), and SRVCSTATE (column name).

Service Status

Current service status. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: OK (0), Error (1), Degraded (2), Unknown (3), Pred Fail (4), Starting (5), Stopping (6), Service (7), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVICE_STATUS or SRVCSTATUS (historical name), *Service Status* (caption), Service_Status (attribute name), and SRVCSTATUS (column name).

Service Type

The service type for this service to the Microsoft SQL Server. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: SQLServer (1), SQLAgent (2), DTC (3), Analysis OLAP (4), Search (5), ADHelper (6), Browser (7), DTS (8), SQLWriter (9), Report (10), Distributed Replay Controller (11), Distributed Replay Client (12), Launchpad (13), PolyBase Data Movement (14), PolyBase Engine (15), Analysis Services CEIP (16), SQL Server CEIP (17), Integration Services CEIP (18), Integration Services Scale Out Worker (19), Integration Services Scale Out Master (20), Unknown (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SERVICE_TYPE or SRVCTYPE (historical name), *Service Type* (caption), Service_Type (attribute name), and SRVCTYPE (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), *Host Name* (caption), Host_Name (attribute name), and HOSTNAME (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Statistics Detail data set

The MS SQL Statistics Detail data set contains attributes that you can use to monitor detailed information about SQL server statistics. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 3 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Average Value (per Sec.)

The average value per second for the statistic since the SQL server was started. The value format is a decimal with two decimal places allowed; for example, 5.00. This value indicates the norm for the statistic during the current interval. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: AVERAGE_VALUE_PER_SECOND or STATAVGS (historical name), *Average Value (per Sec.)* (caption), Average_Value_per_Second (attribute name), and STATAVGS (column name).

Maximum Seen

The greatest value per second for the statistic since the SQL server was started. The value format is an integer; for example, 2000. This is a benchmark value. The type is integer (32-bit numeric property).

The following names are defined for this attribute: MAXIMUM_SEEN or STATMAX (historical name), *Maximum Seen* (caption), Maximum_Seen (attribute name), and STATMAX (column name).

Minimum Seen

The smallest value per second for the statistic since the SQL server was stated. The value format is an integer; for example, 10. This is a benchmark value. The type is integer (32-bit numeric property).

The following names are defined for this attribute: MINIMUM_SEEN or STATMIN (historical name), *Minimum Seen* (caption), Minimum_Seen (attribute name), and STATMIN (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Statistic Name

The name of the statistic. The value format is an alphanumeric string with a maximum of 32 characters; for example, cpu_busy. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: CPU busy (cpu_busy), IO busy (io_busy), connections

(connections), idle (idle), packets error (packets_error), packets received (packets_received), packets sent (packets_sent). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STATISTIC_NAME or STATNAME (historical name), *Statistic Name* (caption), Statistic_Name (attribute name), and STATNAME (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Current Value

The value for the statistic during the current interval. The value format is an integer; for example, 40. This is a cumulative value for the statistic. The type is integer (32-bit gauge).

The following names are defined for this attribute: CURRENT_VALUE or STATCURR (historical name), Current Value (caption), Current_Value (attribute name), and STATCURR (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Since Startup

The total of all the values for the statistic since the SQL server was started. The value format is an integer; for example, 9000. The type is integer (32-bit counter).

The following names are defined for this attribute: TOTAL_SINCE_STARTUP or STATCUM (historical name), Total Since Startup (caption), Total_Since_Startup (attribute name), and STATCUM (column name).

MS SQL Statistics Summary data set

The MS SQL Statistics Summary data set contains attributes that you can use to monitor summary information about SQL server statistics. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every minute and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Client Count Pct Used

The number of client workstations currently connected to the database server and returns the ratio percentage of the number of possible connections. For example, if you have 10 client connections and eight are currently in use, this attribute returns a value of 80%. Running out of client connections can result in server access problems. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: CLIENT_COUNT_PERCENT_USED or CLNTPCT (historical name), *Client Count Pct Used* (caption), Client_Count_Percent_Used (attribute name), and CLNTPCT (column name).

Current Logons

The number of active connections (logons). The value format is an integer; for example, 50. The type is integer (32-bit gauge).

The following names are defined for this attribute: CURRENT_LOGONS or LOGONCURR (historical name), *Current Logons* (caption), Current_Logons (attribute name), and LOGONCURR (column name).

Max User Connections Allowed

The maximum number of active connections (logons) allowed for the SQL server. The value format is an integer in the range 5 through 2147483647; for example, 100. This is a configuration value for the SQL server. To determine the number of connections that can be configured for the SQL server, use the select @@max_connections command. The type is integer (32-bit numeric property).

The following names are defined for this attribute: MAX_USER_CONNECTIONS_ALLOWED or LOGONMAX (historical name), *Max User Connections Allowed* (caption), Max_User_Connections_Allowed (attribute name), and LOGONMAX (column name).

Network Read Rate (per Sec.)

The rate at which tabular data stream (TDS) packets are read from the network, in packets per second. This statistic is an indicator of network throughput. When this statistic is high, it indicates heavy network traffic. The value format is an integer. The type is integer (32-bit gauge).

The following names are defined for this attribute: NETWORK_READ_RATE or NTWKRD RATE (historical name), *Network Read Rate (per Sec.)* (caption), Network_Read_Rate (attribute name), and NTWKRD RATE (column name).

Network Write Rate (per Sec.)

The rate at which tabular data stream (TDS) packets are written to the network, in packets per second. This statistic is an indicator of network throughput. When this statistic is high, it indicates heavy network traffic. The value format is an integer. The type is integer (32-bit gauge).

The following names are defined for this attribute: NETWORK_WRITE_RATE or NTWKWR RATE (historical name), *Network Write Rate (per Sec.)* (caption), Network_Write_Rate (attribute name), and NTWKWR RATE (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Pct Max Logons Active

The percentage of active connections (logons) of the maximum number of active connections allowed for the SQL server. The value format is a percentage with two decimal places allowed; for example, 50.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT_MAX_LOGONS_ACTIVE or LOGONPCT (historical name), *Pct Max Logons Active* (caption), Percent_Max_Logons_Active (attribute name), and LOGONPCT (column name).

Physical Reads (per Sec.)

The average number of physical reads per second during the current interval. The value format is a decimal with two decimal places allowed; for example, 5.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PHYSICAL_READS_PER_SECOND or PHYRDPSEC (historical name), *Physical Reads (per Sec.)* (caption), Physical_Reads_per_Second (attribute name), and PHYRDPSEC (column name).

Physical Writes (per Sec.)

The average number of physical writes per second during the current interval. The value format is a decimal with two decimal places allowed; for example, 5.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PHYSICAL_WRITES_PER_SECOND or PHYWRPSEC (historical name), *Physical Writes (per Sec.)* (caption), Physical_Writes_per_Second (attribute name), and PHYWRPSEC (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Total Logins (Per Sec.)

The number of login operations that started per second in the current interval. The value format is a decimal with two decimal places allowed; for example, 15.00. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_LOGINS_PER_SEC or LOGINPSEC (historical name), *Total Logins (Per Sec.)* (caption), Total_Logins_Per_Sec (attribute name), and LOGINPSEC (column name).

Total Logouts (per Sec.)

Total number of logout operations started per second. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: TOTAL_LOGOUTS_PER_SECOND or LOGOUTPSEC (historical name), *Total Logouts (per Sec.)* (caption), Total_Logouts_per_Second (attribute name), and LOGOUTPSEC (column name).

Client Count

The number of client workstations currently connected to the database server. Use this attribute as an alert when you are running out of client connections, which can result in server access problems. This attribute is also useful for tracking client access patterns. The value format is an integer. The type is integer (32-bit gauge).

The following names are defined for this attribute: CLIENT_COUNT or CLNTCNT (historical name), Client Count (caption), Client_Count (attribute name), and CLNTCNT (column name).

Current Interval (Sec.)

The number of seconds that have elapsed between the previous sample and the current sample. The value format is an integer; for example, 90. A sample contains the data the product collects about each SQL server. New data becomes available if a new interval has occurred and data has been refreshed at the CMS hub. The type is integer (32-bit numeric property).

The following names are defined for this attribute: CURRENT_INTERVAL or CURRINTL (historical name), Current Interval (Sec.) (caption), Current_Interval (attribute name), and CURRINTL (column name).

Disk IO Current Interval

The number of times the SQL server accessed hard disk during the current interval. The value includes access to hard disk for physical reads and physical writes. The value format is an integer; for example, 50. The type is integer (32-bit gauge).

The following names are defined for this attribute: DISK_IO_CURRENT_INTERVAL or DISKIO (historical name), Disk IO Current Interval (caption), Disk_IO_Current_Interval (attribute name), and DISKIO (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

IO Errors Current Interval

The number of I/O errors that occurred when the SQL server accessed hard disk during the current interval. The value format is an integer; for example, 5. The type is integer (32-bit gauge).

The following names are defined for this attribute: IO_ERRORS_CURRENT_INTERVAL or IOERRSCURR (historical name), IO Errors Current Interval (caption), IO_Errors_Current_Interval (attribute name), and IOERRSCURR (column name).

IO Errors Since Startup

The number of I/O errors that have occurred when the SQL server accessed hard disk since startup. The value format is an integer; for example, 2. The type is integer (32-bit counter).

The following names are defined for this attribute: IO_ERRORS_SINCE_STARTUP or IOERRS (historical name), IO Errors Since Startup (caption), IO_Errors_Since_Startup (attribute name), and IOERRS (column name).

Pct IO Errors Cur Intvl

The percentage of the accesses to hard disk that had errors occur during the current interval. The value format is a percentage with two decimal places allowed; for example, 1.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT_IO_ERRORS_CURRENT_INTERVAL or IOERRSPCT (historical name), Pct IO Errors Cur Intvl (caption), Percent_IO_Errors_Current_Interval (attribute name), and IOERRSPCT (column name).

Pct User Connections Remaining

The current number of remaining user connections as a percentage of the maximum number of available user connections for the database server. Use this attribute to gain an overview of high and low access periods and to warn you of impending availability problems. The value format is an integer. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: PERCENT_USER_CONNECTIONS_REMAINING or PCTUSRCREM (historical name), Pct User Connections Remaining (caption), Percent_User_Connections_Remaining (attribute name), and PCTUSRCREM (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS,

where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Total Logons (per Sec.)

The average number of active connections (logons) per second during the current interval. The value format is a decimal with two decimal places allowed; for example, 15.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: TOTAL_LOGONS_PER_SECOND or LOGONPSEC (historical name), Total Logons (per Sec.) (caption), Total_Logons_per_Second (attribute name), and LOGONPSEC (column name).

Total OS CPU Pct Busy

The percentage of CPU seconds the SQL server has used during the current interval of all the CPU seconds used since the SQL server was started. The value format is a percentage with two decimal places allowed; for example, 20.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: TOTAL_OS_CPU_PERCENT_BUSY or PCTCPU (historical name), Total OS CPU Pct Busy (caption), Total_OS_CPU_Percent_Busy (attribute name), and PCTCPU (column name).

Total OS IO Pct Busy

The percentage of I/O the SQL server used during the current interval of all the I/O used since the SQL server was started. The value format is a percentage with two decimal places allowed; for example, 25.00. The type is real number (32-bit gauge) with two decimal places of precision.

The following names are defined for this attribute: TOTAL_OS_IO_PERCENT_BUSY or PCTIO (historical name), Total OS IO Pct Busy (caption), Total_OS_IO_Percent_Busy (attribute name), and PCTIO (column name).

MS SQL Table Detail data set

The MS SQL Table Detail data set contains attributes that you can use to monitor tables within one or more databases. The data collection for this data set is limited to 10000 rows. This data set is configured for historical collection. Thresholds for this data set are associated with the Microsoft SQL Server component. A data sample is sent to the server every 3 minutes and is maintained for 8 days by default. The attributes shown in italic are visible in the UI. All attributes are available for thresholds.

This data set contains the following attributes:

Database Id

ID number of this database. The format is an integer. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_ID or DBID (historical name), *Database Id* (caption), Database_Id (attribute name), and DBID (column name).

Database Name

The Database name. The value format is an alphanumeric string with a maximum of 30 characters.

The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATABASE_NAME or UDBNAME (historical name), *Database Name* (caption), Database_Name (attribute name), and UDBNAME (column name).

FileTable

Indicates whether the current table is a FileTable. The format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Applicable (-2), Not Collected (-1), No (1), Yes (2), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FILETABLE or ISFILTBL (historical name), *FileTable* (caption), FileTable (attribute name), and ISFILTBL (column name).

Fragmentation

The degree of fragmentation for the table. Use this attribute to determine if fragmentation is reaching a level that causes performance degradation. Performance degradation is due to non-contiguous table reads that require additional extent switches. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FRAGMENTATION or PCTFRAG (historical name), *Fragmentation* (caption), Fragmentation (attribute name), and PCTFRAG (column name).

Index Id

The ID number of this table index. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INDEX_ID or IDXID (historical name), *Index Id* (caption), Index_Id (attribute name), and IDXID (column name).

Index Name

The table index name. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: INDEX_NAME or UIDXNAME (historical name), *Index Name* (caption), Index_Name (attribute name), and UIDXNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), *Node* (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number Of Rows

The number of rows that are present in a database table. The format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), No Data (-3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NUMBER_OF_ROWS or NUMROWS (historical name), *Number Of Rows* (caption), Number_Of_Rows (attribute name), and NUMROWS (column name).

Optimizer Statistics Age (in Min.)

The time (in minutes) since statistics were last updated for the table. Use this attribute to ensure that queries base their query plans on up-to-date information. If query plans are based on old information, they might be inefficient. The value for this attribute is displayed as 0 (zero) when the time since statistics was last updated for the table is less than 1 minute and is displayed as Not Collected when the statistics for the table was never updated. The format is an integer. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OPTIMIZER_STATISTICS_AGE or OPTSTATAGE (historical name), *Optimizer Statistics Age (in Min.)* (caption), Optimizer_Statistics_Age (attribute name), and OPTSTATAGE (column name).

Space Used (MB)

The amount of space (in megabytes) used by the specified table. Use this attribute to determine how much of the database space is used by a table. You can monitor the growth of individual tables and compare the actual growth to the expected growth. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SPACE_USED or SPCUSED (historical name), *Space Used (MB)* (caption), Space_Used (attribute name), and SPCUSED (column name).

Table Id

ID number of this table. The format is an integer. This attribute is a key attribute. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_ID or TBLID (historical name), *Table Id* (caption), Table_Id (attribute name), and TBLID (column name).

Table Name

The table name. The value format is an alphanumeric string with a maximum of 30 characters. The type is string with enumerated values. The following values are defined: No Data (n/d). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TABLE_NAME or UTBLNAME (historical name), *Table Name* (caption), Table_Name (attribute name), and UTBLNAME (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), *Timestamp* (caption), *Timestamp* (attribute name), and TIMESTAMP (column name).

Durability

A numeric value that indicates whether the table is durable. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1), Not Applicable (-2), No Data (-3), SCHEMA AND DATA (1), SCHEMA ONLY (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DURABILITY (historical name), Durability (caption), Durability (attribute name), and DURABILITY (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Memory Optimized

A numeric value that indicates whether the table is optimized for memory. The type is integer (32-bit gauge) with enumerated values. The following values are defined: No Data (-3), Not Applicable (-2), Not Collected (-1), Not Memory Optimized (1), Is memory optimized (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_OPTIMIZED or ISMEMOPT (historical name), Memory Optimized (caption), Memory_Optimized (attribute name), and ISMEMOPT (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and SS indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value format is an alphanumeric string with a maximum of 32 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

MS SQL Table Summary data set

The MS SQL Table Summary data set provides the summary of all the tables that are present in an SQL Server instance. Note that empty tables are not monitored. If the table detail collection settings are configured for the SQL Server agent, the attributes in this data set display the last collected value of data.

This data set contains the following attributes:

Last Collection Duration (Min.)

The time (in minutes) that is required to collect data of the MS SQL Table Detail data set. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_COLLECTION_DURATION or COLLTIMDUR (historical name), Last Collection Duration (Min.) (caption), Last_Collection_Duration (attribute name), and COLLTIMDUR (column name).

Last Collection Start Time

The date and time when data collection of the MS SQL Table Detail data set last started. The type is timestamp with enumerated values. The following values are defined: Not Collected (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_COLLECTION_START_TIME or STCOLLTIM (historical name), Last Collection Start Time (caption), Last_Collection_Start_Time (attribute name), and STCOLLTIM (column name).

Maximum Fragmentation Percent

The highest fragmentation value (in a percent) among all the tables in the current interval. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_FRAGMENTATION_PERCENT or MAXFRAGPCT (historical name), Maximum Fragmentation Percent (caption), Maximum_Fragmentation_Percent (attribute name), and MAXFRAGPCT (column name).

Maximum Fragmentation Percent Since Startup

The highest value of fragmentation (in a percent) among all tables since the agent was started. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute:

MAXIMUM_FRAGMENTATION_PERCENT_SINCE_STARTUP or MAXFRGSTRT (historical name), Maximum Fragmentation Percent Since Startup (caption), Maximum_Fragmentation_Percent_Since_Startup (attribute name), and MAXFRGSTRT (column name).

Maximum Optimizer Statistics Age (Min.)

The maximum age of the optimizer statistics in minutes among all the tables in the current instance of the SQL Server. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_OPTIMIZER_STATISTICS_AGE or MAXOPTSTAG (historical name), Maximum Optimizer Statistics Age (Min.) (caption), Maximum_Optimizer_Statistics_Age (attribute name), and MAXOPTSTAG (column name).

Maximum Table Size (MB)

The maximum space in MB that is used among all the tables in the current instance of the SQL Server. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_TABLE_SIZE_MB or MAXTBSZMB (historical name), Maximum Table Size (MB) (caption), Maximum_Table_Size_MB (attribute name), and MAXTBSZMB (column name).

Minimum Fragmentation Percent

The lowest value of fragmentation (in a percent) among all the tables in the current interval. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MINIMUM_FRAGMENTATION_PERCENT or MINFRAGPCT (historical name), Minimum Fragmentation Percent (caption), Minimum_Fragmentation_Percent (attribute name), and MINFRAGPCT (column name).

Minimum Optimizer Statistics Age (Min.)

The minimum age of the optimizer statistic in minutes among all the tables in the current instance of the SQL Server. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MINIMUM_OPTIMIZER_STATISTICS_AGE or MINOPTSTAG (historical name), Minimum Optimizer Statistics Age (Min.) (caption), Minimum_Optimizer_Statistics_Age (attribute name), and MINOPTSTAG (column name).

Minimum Table Size (MB)

The minimum space in MB that is used among all the tables of the SQL Server. The type is real number (32-bit numeric property) with two decimal places of precision with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MINIMUM_TABLE_SIZE_MB or MINTBSZMB (historical name), Minimum Table Size (MB) (caption), Minimum_Table_Size_MB (attribute name), and MINTBSZMB (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (historical name), ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (historical name), SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL Server. The value format is an alphanumeric string with a maximum of 32 characters. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total FileTables

The total number of FileTables that are present in the current instance of the SQL Server. This attribute is not supported in MSSQL Linux agent. The type is integer (32-bit gauge) with enumerated values.

The following values are defined: Not Applicable (-2), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_FILETABLES or NUMFILTBL (historical name), Total FileTables (caption), Total_FileTables (attribute name), and NUMFILTBL (column name).

Total Tables

The total number of tables that are present in the current instance of the SQL Server. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TOTAL_TABLES or NUMTBL (historical name), Total Tables (caption), Total_Tables (attribute name), and NUMTBL (column name).

MS SQL Text data set

The MS SQL Text data set contains attributes that you can use to monitor information about SQL text strings associated with a selected process. For custom queries, this data set displays data only for the specified available process ID. This data set is not supported in MSSQL APM agent.

This data set contains the following attributes:

Client Host Name

The host name of the client where the command was issued. The value format is an alphanumeric string with a maximum of 16 characters. The type is string.

The following names are defined for this attribute: Client Host Name (caption), Client_Host_Name (attribute name), and CLNTHOST (column name).

Client Process ID

The process ID assigned by the host client. The value format is an alphanumeric string with a maximum of 16 characters. The type is string.

The following names are defined for this attribute: Client Process ID (caption), Client_Process_ID (attribute name), and CLNTPID (column name).

Database Name

The name of the database. The value format is an alphanumeric string with a maximum of 30 characters; for example, KOQ3. Each database name is unique. The SQL server also assigns each database its own identification number. The type is string.

The following names are defined for this attribute: Database Name (caption), Database_Name (attribute name), and DBNAME (column name).

Database Name (Unicode)

The Database name. The type is string.

The following names are defined for this attribute: Database Name (Unicode) (caption), Database_Name_U (attribute name), and UDBNAME (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Process ID

The ID of the process that is requesting or holding the lock on the resource. The value format is an alphanumeric string with a maximum of 10 characters; for example, 42168. This attribute is a key attribute. The type is integer (32-bit numeric property).

The following names are defined for this attribute: Process ID (caption), Process_ID (attribute name), and PID (column name).

Process Status

Indicates the status of the process. The type is string.

The following names are defined for this attribute: Process Status (caption), Process_Status (attribute name), and STATUS (column name).

ROWNO

The row number in sample. The type is integer (32-bit numeric property).

The following names are defined for this attribute: ROWNO (caption), ROWNO (attribute name), and ROWNO (column name).

Sample Timestamp

The timestamp that indicates the date and time the product collected the sample for the SQL server. A sample is the data the product collects about the SQL server. The format is MM/DD/YY HH:mm:SS, where, MM indicates month, DD indicates day, YY indicates year, HH indicates hour, mm indicates minute, and ss indicates second. For example, 01/25/02 08:00:00 indicates that the product collected the data from the SQL server on Friday, January 25, 2002 at 8:00 a.m. The type is timestamp.

The following names are defined for this attribute: Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

SAMPLENO

Sample number. The type is integer (32-bit numeric property).

The following names are defined for this attribute: SAMPLENO (caption), SAMPLENO (attribute name), and SAMPLENO (column name).

Server

The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. The type is string.

The following names are defined for this attribute: Server (caption), Server (attribute name), and SERVERID (column name).

Text

The SQL text for the process. The value format is an alphanumeric string with a maximum of 255 characters. The type is string.

The following names are defined for this attribute: Text (caption), Text (attribute name), and SQLTEXT (column name).

Text (Unicode)

The SQL text for the process. The type is string.

The following names are defined for this attribute: Text (Unicode) (caption), Text_U (attribute name), and USQLTEXT (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total CPU Time (Sec.)

The amount of CPU time, in seconds, the process has used on the host since the process started. The value format is an integer; for example, 60. This value is based on the statistics collected by the SQL server. Use this value to check for processes that use abnormal amounts of CPU time. The type is real number (32-bit gauge) with four decimal places of precision.

The following names are defined for this attribute: Total CPU Time (Sec.) (caption), Total_CPU_Time (attribute name), and CPU (column name).

MS SQL Workload Group Stats data set

The MS SQL Workload Group Stats data set provides statistics associated with workload groups. This data set is not supported on the SQL Server 2005 and 2008 Express Edition.

This data set contains the following attributes:

Active Parallel Threads

The number of threads used by parallel queries in the workload group. Serial queries and the main thread of parallel queries are not included in this number. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTIVE_PARALLEL_THREADS or APTHREADS (historical name), Active Parallel Threads (caption), Active_Parallel_Threads (attribute name), and APTHREADS (column name).

Active Requests

The number of requests currently running in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACTIVE_REQUESTS or ACTREQS (historical name), Active Requests (caption), Active_Requests (attribute name), and ACTREQS (column name).

Blocked Tasks

The number of blocked tasks in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BLOCKED_TASKS or BLKTASK (historical name), Blocked Tasks (caption), Blocked_Tasks (attribute name), and BLKTASK (column name).

CPU Usage Pct

The system CPU usage by all requests in the specified instance of the performance object. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_USAGE_PERCENT or WGSCPUPER (historical name), CPU Usage Pct (caption), CPU_Usage_Percent (attribute name), and WGSCPUPER (column name).

Host Name

The name of the computer on which the SQL Server is running. The value format is an alphanumeric string with a maximum of 64 characters, for example, Voyager. This attribute is not available for use in eventing thresholds or for historical data collection. The type is string.

The following names are defined for this attribute: HOST_NAME or HOSTNAME (historical name), Host Name (caption), Host_Name (attribute name), and HOSTNAME (column name).

Max Requests Cpu Time(ms)

The maximum requests for CPU time in milliseconds. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_REQUESTS_CPU_TIME or MAXREQCT (historical name), Max Requests Cpu Time(ms) (caption), Max_Requests_Cpu_Time (attribute name), and MAXREQCT (column name).

Max Requests Memory Grant(MB)

The maximum value of memory grant in megabytes used by a query in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_REQUESTS_MEMORY_GRANT or MAXREQMG (historical name), Max Requests Memory Grant(MB) (caption), Max_Requests_Memory_Grant (attribute name), and MAXREQMG (column name).

Node

The managed system name of the agent. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODE (historical name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Query Optimization Per Sec

The number of query optimizations per second occurring in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUERY_OPTIMIZATION_PER_SEC or QOPERSEC (historical name), Query Optimization Per Sec (caption), Query_Optimization_Per_Sec (attribute name), and QOPERSEC (column name).

Queued Requests

The number of requests waiting in the queue due to resource governor limits in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: QUEUED_REQUESTS or QUEUEREC (historical name), Queued Requests (caption), Queued_Requests (attribute name), and QUEUEREC (column name).

Reduced Memory Grants Per Sec

The number of queries reduced per second from the memory in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REDUCED_MEMORY_GRANTS_PER_SEC or RMGPERSEC (historical name), Reduced Memory Grants Per Sec (caption), Reduced_Memory_Grants_Per_Sec (attribute name), and RMGPERSEC (column name).

Requests Completed Per Sec

The number of completed requests per second in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REQUESTS_COMPLETED_PER_SEC or RCPERSEC (historical name), Requests Completed Per Sec (caption), Requests_Completed_Per_Sec (attribute name), and RCPERSEC (column name).

Sample Timestamp

The time when data is collected. The format is an integer. The type is timestamp.

The following names are defined for this attribute: SAMPLE_TIMESTAMP or SAMPTIME (historical name), Sample Timestamp (caption), Sample_Timestamp (attribute name), and SAMPTIME (column name).

Server

The name of the SQL Server. The type is string.

The following names are defined for this attribute: SERVER or SERVERID (historical name), Server (caption), Server (attribute name), and SERVERID (column name).

Suboptimal Plans Per Sec

The number of suboptimal query plans generated per second in the workload group. The format is an integer. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Value Exceeds Maximum (2147483647), Not Collected (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SUBOPTIMAL_PLANS_PER_SEC or SOPERSEC (historical name), Suboptimal Plans Per Sec (caption), Suboptimal_Plans_Per_Sec (attribute name), and SOPERSEC (column name).

Timestamp

The local time at the agent when the data was collected. The type is string.

The following names are defined for this attribute: TIMESTAMP (historical name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Workload Group Name (Unicode)

The name of the workload group. This attribute is a key attribute. The type is string with enumerated values. The following values are defined: No Data (nodata). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: WORKLOAD_GROUP_NAME or WLGNAME (historical name), Workload Group Name (Unicode) (caption), Workload_Group_Name (attribute name), and WLGNAME (column name).

Accessibility features

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Accessibility features

The web-based interface of IBM® Cloud Application Performance Management is the Cloud APM console. The console includes the following major accessibility features:

- Enables users to use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen.¹ Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Enables users to operate specific or equivalent features using only the keyboard.
- Communicates all information independently of color.²

The Cloud APM console uses the latest W3C Standard, [WAI-ARIA 1.0](#), [US Section 508](#), and [Web Content Accessibility Guidelines \(WCAG\) 2.0](#). To take advantage of accessibility features, use the latest release of your screen reader in combination with the latest web browser that is supported by this product.

The Cloud APM console online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described at [IBM Knowledge Center release notes](#).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The Cloud APM console web user interface does not rely on cascading style sheets to render content properly and to provide a usable experience. However, the product documentation does rely on cascading style sheets. IBM Knowledge Center provides an equivalent way for low-vision users to use their custom display settings, including high-contrast mode. You can control font size by using the device or browser settings.

The Cloud APM console web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

The Cloud APM console user interface does not have content that flashes 2 - 55 times per second.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility](#).

¹ Exceptions include some of the **Agent Configuration** pages and historical line charts in the Cloud APM console.

² Exceptions include some **Agent Configuration** pages of the Cloud APM console.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2014, 2015.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux® is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth in the following paragraphs.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies can be disabled, but disabling them will also likely eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek

your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Index

A

accessibility features [251](#)
Accessible attribute [124](#)
Action ID attribute [79](#)
Active Jobs attribute [162](#)
Active Memory Grant Amount(MB) attribute [195](#)
Active Memory Grants Count attribute [195](#)
Active Parallel Threads attribute [248](#)
Active Requests attribute [248](#)
Active Transactions attribute [124](#)
additional information
 attributes [43](#)
Age of Last Error (Min.) attribute [182, 207](#)
Agent Type attribute [175, 176](#)
Aggregate Cache (KB) attribute [59](#)
Aggregate Database Freespace (MB) attribute [124](#)
Aggregate Database Freespace Pct attribute [131](#)
Aggregation Hits Per Sec attribute [71](#)
Aggregation Lookups Per Sec attribute [71](#)
Aggregation Map Files attribute [59](#)
attributes
 Accessible [124](#)
 Action ID [79](#)
 Active Jobs [162](#)
 Active Memory Grant Amount(MB) [195](#)
 Active Memory Grants Count [195](#)
 Active Parallel Threads [248](#)
 Active Requests [248](#)
 Active Transactions [124](#)
 Age of Last Error (Min.) [182, 207](#)
 Agent Type [175, 176](#)
 Aggregate Cache (KB) [59](#)
 Aggregate Database Freespace (MB) [124](#)
 Aggregate Database Freespace Pct [131](#)
 Aggregation Hits Per Sec [71](#)
 Aggregation Lookups Per Sec [71](#)
 Aggregation Map Files [59](#)
 Audit Action Name [79](#)
 Audit Name [79](#)
 Audit Specification Name [79](#)
 Audit Specification Type [79](#)
 Auto-Param Attempts Per Sec [111](#)
Autogrowth [141](#)
Autogrowth Unit [141](#)
Autogrowth Value [141](#)
Availability Mode [100](#)
Average CPU Time(ms) [146](#)
Average Duration(ms) [146](#)
Average Logical Reads [146](#)
Average Logical Writes [147](#)
Average Physical Reads [147](#)
Average Time Per Query [71](#)
Average Value (per Sec.) [235](#)
Average Wait Time locks (ms) [169](#)
Avg CPU Time (Sec.) [152](#)
Avg Logical Reads [152](#)

attributes (continued)

Avg Logical Writes [152](#)
Avg Physical Reads [153](#)
Avg Size Message Fragments Received [231](#)
Avg Size Message Fragments Sent [231](#)
Backup Preference [94, 96](#)
Backup Priority [100](#)
Batch Requests Per Sec [112](#)
Blocked Database Name (Unicode) [114](#)
Blocked Parent Query Text [114](#)
Blocked Query Text [114](#)
Blocked Session ID [114](#)
Blocked Tasks [248](#)
Blocked User Name (Unicode) [114](#)
Blocking Process Duration(Mins.) [184](#)
Blocking Process ID [164, 185](#)
Blocking Query Text [114](#)
Blocking Reason [114](#)
Blocking Resource [115](#)
Blocking Resource Duration(Mins.) [187](#)
Blocking Resource Type [187](#)
Blocking Session ID [115, 198](#)
Blocking User Name [115](#)
Broker Transaction Rollbacks [228](#)
Buffer In Use [156](#)
Buffer Memory (MB) [156](#)
Buffer Spooled [156](#)
Bytes Received Per Sec [104, 133](#)
Bytes Sent Per Sec [71, 104, 134](#)
Bytes Sent To Transport Per Sec [104](#)
Cache Average Free Scan [215](#)
Cache Free Buffers [218](#)
Cache Hit Ratio [215](#)
Cache Maximum Free Page Scan [215](#)
Cache Memory Target(MB) [195](#)
Calculation Cache Hits Per Sec [72](#)
Calculation Cache Lookups Per Sec [72](#)
Cases Per sec [52](#)
Checkpoint Pages per Sec [219](#)
Class Type [79](#)
Cleaner Balance Per Sec [60](#)
Cleaner Current Price [60](#)
Cleaner Memory (KB) [60](#)
Cleaner Memory nonshrinkable (KB) [60](#)
Cleaner Memory shrinkable (KB) [60](#)
Cleaner Memory Shrunk KB Per Sec [60](#)
Client Count [238](#)
Client Count Pct Used [237](#)
Client Group ID [165, 187](#)
Client Host Name [185, 198, 246](#)
Client Process ID [185, 246](#)
Client User ID [165, 185](#)
Cluster Node [103](#)
Cluster Resource ID [96](#)
Collation [211](#)
Collation ID [211](#)
Collection Status [207, 216](#)

attributes (continued)
 Command 185
 Command (Unicode) 188
 Command Type 198
 Compile Memory Target(MB) 195
 Concurrent Data Mining Queries 53
 Config Parameter 116
 Config Value 116
 Connection Level 193
 Connection Memory(MB) 177
 Connection State 106
 Containment 51
 CPU Control Effect Pct 195
 CPU Idle Delta 219
 CPU Idle Sec 219
 CPU Pct Idle 219
 CPU Rate (percentage) 148
 CPU Time (Sec.) 198
 CPU Usage Delta 219
 CPU Usage Pct 195, 248
 CPU Usage Sec 219
 CPU Usage Target Pct 196
 Create Date 100
 Current Bytes For Receive I/O 231
 Current Bytes For Send I/O 231
 Current Cached Evaluation Nodes 54
 Current CPU Pct Used 185
 Current Dimension Queries 72
 Current Evaluation Nodes 54
 Current Execution Status 160
 Current Execution Step 158
 Current Interval (Sec.) 139, 183, 191, 193, 220, 238
 Current Locks 201
 Current Logons 207, 237
 Current Measure Group Queries 72
 Current Message Fragments For Send I/O 231
 Current Models Processing 53
 Current Pyramid Operations 72
 Current Retry Attempt 158
 Current Status 158
 Current Value 236
 Customized SQL Definition File 117
 Data Bytes Per Sec 72
 Data Cache Size (KB) 201, 216
 Data File Names 125
 Data Files Freespace (MB) 131
 Data Freespace (MB) 131
 Data Freespace Percent 131
 Data Reads Per Sec 73
 Data Size (MB) 125
 Database Count 94
 Database Growth Percent 125
 Database Id 151, 165, 166, 241
 Database ID 51, 84, 173, 175, 176
 Database Max blocks 207
 Database Max blocks (Unicode) 206
 Database Max Growth Size(MB) 125
 Database Max Locks 207
 Database Max Locks (Unicode) 206
 Database Name 49, 51, 79, 82, 84, 86, 117, 125, 143, 165, 168, 173, 175, 176, 188, 241, 246
 Database Name (Unicode) 125, 141, 147, 149, 153, 165, 167, 186, 198, 226, 246
 Database Principal Name 80

attributes (continued)
 Database Space Pct Used 126
 Database State 82, 126
 Database Status 126
 DB ID 126
 DB Owner 131
 DB Owner (Unicode) 126
 DBO Only Access 126
 Definition File Last Modified Time 118
 Device Free Space (MB) 141
 Device Free Space Percent 142
 Device Name 143
 Device Name (Unicode) 142
 Device Size (MB) 142
 Device Status 142
 Device Type 142
 Dimension Cache Hits Per Sec 73
 Dimension Cache Lookups Per Sec 73
 Dimension Index(hash) Files 60
 Dimension Property Files 61
 Dimension Queries Per Sec 73
 Dimension String Files 61
 Disk IO Current Interval 239
 Dump Tran Date 131
 Durability 242
 Duration (Sec.) 186
 Edition 212
 Edition ID 214
 Enabled 158
 Endpoint URL 100
 Engine Edition 212
 Enqueued Local Messages Per Sec 228
 Enqueued Messages Per Sec 228
 Enqueued Transport Messages Per Sec 228
 Error ID 144, 180
 Error Log Name 203
 Error Log Name (Unicode) 203
 Error Log Size (Bytes) 182, 201, 207
 Error Source 175
 Error SPID 144, 180
 Error Status 126
 Event Time 80
 Exclusive Extent Locks 171
 Exclusive Intent Locks 171
 Exclusive Locks 170
 Exclusive Page Locks 171
 Exclusive Table Locks 171
 Execution rate 147
 Executions Per Interval 148
 Extent Locks 171
 Fact Aggregation Files 61
 Fact Data Files 61
 Fact String Files 61
 Failed Auto-Params Per Sec 112
 Failed Jobs 162
 Failover Mode 100
 Failover Ready 84
 Failure Condition 94, 96
 File Bytes Received Per Sec 86
 File Name 80
 Filegroup ID 149
 Filegroup Max Growth Size 149
 Filegroup Name 149
 Filegroup Size 149

attributes (continued)

Filegroup Space Pct Used [150](#)
Filegroup Status [150](#)
Filestore (KB) [61](#)
Filestore Clock Pages Examined Per Sec [61](#)
Filestore Clock Pages HaveRef Per Sec [61](#)
Filestore Clock Pages Valid Per Sec [62](#)
Filestore KB Reads Per Sec [62](#)
Filestore KB Writes Per Sec [62](#)
Filestore Memory Pinned (KB) [62](#)
Filestore Page Faults Per Sec [62](#)
Filestore Reads Per Sec [62](#)
Filestore Writes Per Sec [63](#)
Filestream Directory Name [51](#)
Filestream Level [212](#)
Filestream Send Rate (KB Per Sec) [49](#)
Filestream Share Name [212](#)
FileTable [241](#)
FileTable Directory [151](#)
FileTable Enabled [151](#)
Filter Rows Excluded Per Sec [73](#)
Filter Rows Included Per Sec [73](#)
Filtered Rows Per Sec [73](#)
Flat Buffer Memory (MB) [156](#)
Flat Buffers In Use [156](#)
Flat Cache Hits Per Sec [74](#)
Flat Cache Lookups Per Sec [74](#)
Flow Control Per Sec [104](#)
Flow Control Time (ms) [105](#)
Forced Parameterization Per Sec [112](#)
Forwarded Messages Discarded Per Sec [228](#)
Forwarded Pending Messages [229](#)
Fragmentation [241](#)
Free Space Accounting Suppressed [127](#)
Fully Qualified Domain Name [214](#)
Granted Query Memory [198](#)
Granted Workspace Memory(MB) [178](#)
Group Database ID [49, 82, 84](#)
Group ID [49, 82, 85, 91, 93, 96, 100, 103, 106](#)
Group Name [82, 85, 91, 93, 97, 101, 103, 108](#)
Group Role [94](#)
HADR Manager Status [214](#)
Host Name [51, 95, 98, 112, 116, 127, 134, 140, 143, 144, 150, 151, 160, 163, 166, 168, 170, 171, 178, 181, 183, 188, 191, 193, 196, 204, 207, 212, 216, 224, 234, 236, 239, 242, 246, 248](#)
In Memory Aggregation Map File (KB) [63](#)
In Memory Aggregation Map File KB Per Sec [63](#)
In Memory Dimension Index(hash) File (KB) [63](#)
In Memory Dimension Index(hash) File KB Per Sec [63](#)
In Memory Fact Aggregation File (KB) [63](#)
In Memory Fact Aggregation File KB Per Sec [63](#)
In Memory Fact Data File (KB) [64](#)
In Memory Fact Data File KB Per Sec [64](#)
In Memory Fact String File (KB) [64](#)
In Memory Fact String File (KB) Per Sec [64](#)
In Memory Map File (KB) [64](#)
In Memory Map File (KB) Per Sec [64](#)
In Memory Other File (KB) [65](#)
In Memory Other File (KB) Per Sec [65](#)
In Memory Property File (KB) [65](#)
In Memory Property File KB Per Sec [65](#)
In Memory String File (KB) [65](#)

attributes (continued)

In Memory String File KB Per Sec [65](#)
Index Bytes Per Sec [74](#)
Index Id [241](#)
Index Name [241](#)
Index Reads Per Sec [74](#)
Instance Name [105](#)
Instance Path [214](#)
Instance Port [214](#)
Integer Column Name 1 [119](#)
Integer Column Name 2 [119](#)
Integer Column Name 3 [119](#)
Integer Column Name 4 [120](#)
Integer Column Name 5 [120](#)
Integer Column Value 1 [120](#)
Integer Column Value 2 [120](#)
Integer Column Value 3 [120](#)
Integer Column Value 4 [120](#)
Integer Column Value 5 [120](#)
Intent Locks [170](#)
IO Errors Current Interval [207, 239](#)
IO Errors Since Startup [239](#)
Is Clustered [212](#)
Is Conformant [91](#)
Is DHCP [92](#)
Is HADR Enabled [214](#)
Is Single User [212](#)
Is Integrated Security Only [213](#)
Job Category ID [160](#)
Job Category Name [160](#)
Job Description [160](#)
Job Duration [158](#)
Job Error Code [158](#)
Job Execution Duration [158](#)
Job Failure Count [159](#)
Job Id [159](#)
Job Name [159, 176](#)
Job Owner [159](#)
Job Status [159, 176](#)
Job Success Count [159](#)
Job Success Rate [162](#)
Job Type [160](#)
Jobs Executed Current Interval [162](#)
Jobs Failed Current Interval [162](#)
Jobs Failed Since Startup [163](#)
Join State [85, 103](#)
Last Collection Duration (Min.) [243](#)
Last Collection Start Time [243](#)
Last Connect Error Number [108](#)
Last Database Size(MB) [132](#)
Last Execution Error Message [118](#)
Last Execution Status [118](#)
Last Execution Time [118, 153](#)
Last File Processed [173](#)
Last File Processed Time [173](#)
Last Group Role [95](#)
Last Run Outcome [159](#)
Last Run Timestamp [159](#)
Lazy Writes per Sec [220](#)
License Type [213](#)
Listener ID [92](#)
Listener IP Address [92](#)
Listener Name [92](#)
Listener Subnet IP [92](#)

attributes (continued)

Local 82, 107
Local Replica Recovery Health 95
Lock Blocks 179
Lock Blocks Allocated 179
Lock Memory(MB) 177
Lock Owner Blocks 179
Lock Owner Blocks Allocated 179
Lock Request Status 167
Lock Requests per Sec. 169
Lock Resource Type 167, 169
Lock Timeouts per Sec. 169
Lock Wait Time (ms) 169
Log Bytes Flushed Per Sec 132
Log Bytes Received Per Sec 86, 134
Log Bytes Redone From Cache Per Sec 134
Log Bytes Sent From Cache Per Sec 134
Log Bytes Sent Per Sec 134
Log Bytes Sent Rate (KB Per Sec) 49
Log Compressed Bytes Rcvd Per Sec 135
Log Compressed Bytes Sent Per Sec 135
Log File Names 127
Log Freespace (MB) 127
Log Freespace Percent 132
Log Growths 127
Log Harden Time (ms) 135
Log Max Growth Size(MB) 127
Log Remaining For Undo(MB) 135
Log Scanned For Undo(MB) 135
Log Send Flow Control Time (ms) 135
Log Send Queue (KB) 86
Log Send Queue Size (KB) 50
Log Send Queue(MB) 136
Log Size (MB) 127
Log Space Pct Used 128
Log Time 175, 176
Login Name 188, 198
Long Running Transaction Name 128
Long Running Transaction Process ID 128
Long Running Transaction Time (Sec.) 128
Longest Transaction Running Time (Sec.) 224
Map Bytes Per Sec 74
Map Files 65
Map Reads Per Sec 74
Max CLR Time(Sec.) 153
Max CPU Time(Sec.) 153
Max Disconnected Replicas 109
Max Jobs Failed Interval 163
Max Local Failed Replicas 109
Max Locks Allowed 201
Max logical Reads 153
Max logical Writes 153
Max Memory(MB) 196
Max Non Healthy Databases 88
Max Non Healthy Groups 98
Max Non Healthy Replicas 109
Max Non Online Local Replica Groups 98
Max Non-Online Databases 89
Max Physical Reads 154
Max Requests Cpu Time(ms) 248
Max Requests Memory Grant(MB) 248
Max Running Jobs Interval 164
Max Secondary Replicas 109
Max Unsynchronized Commit DB Replicas 89

attributes (continued)

Max Unsynchronized Databases 89
Max User Connections Allowed 237
Max Waiting Time(Sec.) 154
Maximum Fragmentation Percent 244
Maximum Fragmentation Percent Since Startup 244
Maximum Optimizer Statistics Age (Min.) 244
Maximum Seen 235
Maximum Sev Current Interval 182
Maximum Sev Level 184, 208
Maximum Sev Timestamp 183, 208
Maximum Table Size (MB) 244
Maximum Value 116
Maximum Workspace Memory(MB) 179
Measure Group Cache Hits Per Sec 75
Measure Group Lookups Per Sec 75
Measure Group Queries Per Sec 75
Memory Grant Timeouts Per Sec 196
Memory Grants Outstanding 179
Memory Grants Pending 180
Memory Grants Per Sec 196
Memory Limit High (KB) 66
Memory Limit Low (KB) 66
Memory Optimized 243
Memory Usage (KB) 66
Message 174–176
Message Age (Min.) 145, 180
Message Fragment Received Per Sec 231
Message Fragments Sent Per Sec 232
Message Issuer 145, 181
Message Text 181
Message Text (Unicode) 145, 180
Message Timestamp 145, 181
Messages Sent To Transport Per Sec 105
Min Disconnected Replicas 109
Min Local Failed Replicas 110
Min Non Healthy Databases 89
Min Non Healthy Groups 98
Min Non Healthy Replicas 110
Min Non Online Local Replica Groups 99
Min Non-Online Databases 89
Min Secondary Replicas 110
Min Unsynchronized Commit DB Replicas 89
Min Unsynchronized Databases 90
Minimum Fragmentation Percent 244
Minimum Optimizer Statistics Age (Min.) 244
Minimum Pct Data Freespace 140, 208
Minimum Pct Log Freespace 140, 208
Minimum Seen 235
Minimum Table Size (MB) 245
Minimum Value 116
Mirror Device Name 142
Mirror Name (Unicode) 136
Mirrored Write Transaction Per Sec 86
Mirrored Write Transactions Per Sec 136
Mirroring Role 132, 136
Mirroring State 136
Mixed Page Allocations Per Sec 220
Modify Date 101
More Columns 121
MS SQL Additional Availability Database Details 49
MS SQL Additional Database Detail 51
MS SQL Analysis Services Data Mining 52
MS SQL Analysis Services MDX Processing 54

attributes (*continued*)

MS SQL Analysis Services Memory Statistics 59
MS SQL Analysis Services Rows Processing 69
MS SQL Analysis Services Storage Engine Query 71
MS SQL Audit Details 79
MS SQL Availability Database Details 81
MS SQL Availability Database Details In Cluster 84
MS SQL Availability Database Statistics 86
MS SQL Availability Database Summary 88
MS SQL Availability Group Listener Details 91
MS SQL Availability Groups Details 93
MS SQL Availability Groups Details In Cluster 96
MS SQL Availability Groups Summary 98
MS SQL Availability Replicas Details 100
MS SQL Availability Replicas Details In Cluster 102
MS SQL Availability Replicas Statistics 104
MS SQL Availability Replicas Status 106
MS SQL Availability Replicas Status Summary 108
MS SQL Batch Stats 111
MS SQL Blocked Queries Details 113
MS SQL Configuration 116
MS SQL Customized SQL Query Details 117
MS SQL Customized SQL Query Result 119
MS SQL Database Detail 124
MS SQL Database Mirroring 133
MS SQL Database Summary 138
MS SQL Device Detail 141
MS SQL Error Event Detail 144
MS SQL Expensive Query Plans 146
MS SQL Filegroup Detail 149
MS SQL FileTable Detail 150
MS SQL Individual Queries Details 152
MS SQL Integration Service Details 156
MS SQL Job Detail 157
MS SQL Job Summary 162
MS SQL Lock Conflict Detail 164
MS SQL Lock Detail 166
MS SQL Lock Resource Type Summary 169
MS SQL Lock Summary 170
MS SQL Log Shipping DB Details 173
MS SQL Log Shipping Errors 174
MS SQL Log Shipping Jobs Detail 176
MS SQL Memory Manager 177
MS SQL Problem Detail 180
MS SQL Problem Summary 182
MS SQL Process Detail 184
MS SQL Process Summary 189
MS SQL Remote Servers 193
MS SQL Resource Pool Stats 194
MS SQL Running Queries Details 197
MS SQL Server Detail 201
MS SQL Server Enterprise View 206
MS SQL Server Properties 211
MS SQL Server Summary 215
MS SQL Server Transactions Summary 223
MS SQL Service Broker Activation 226
MS SQL Service Broker Statistics 228
MS SQL Service Broker Transport 231
MS SQL Services Detail 233
MS SQL Statistics Detail 235
MS SQL Statistics Summary 237
MS SQL Table Detail 240
MS SQL Table Summary 243
MS SQL Text 246

attributes (*continued*)

MS SQL Workload Group Stats 247
Network Name 193
Network Read Rate (per Sec.) 237
Network Round Trips Per Sec 75
Network Write Rate (per Sec.) 237
Next Run Timestamp 161
No CKPT After Recovery 132
Node 50, 51, 53, 55, 66, 69, 75, 80, 82, 85, 87, 88, 92, 93, 97, 98, 101, 103, 105, 107, 109, 112, 115, 116, 118, 121, 128, 136, 138, 143, 145, 147, 150, 151, 154, 156, 160, 163, 165, 167, 169, 170, 174, 175, 177, 181, 183, 186, 189, 193, 196, 198, 202, 206, 213, 216, 224, 226, 229, 232, 234, 235, 237, 241, 245, 246, 249
Non Transactional Access Level 52
Notify Level Eventlog 161
NT User 188
Num Licenses 213
Number Deadlocks per Second 169
Number Of Rows 241
Number of Steps 161
Object Name 80
Oldest Open Transaction(min) 128
Operational State 107
Optimizer Memory(MB) 178
Optimizer Statistics Age (in Min.) 242
OS Process ID 188
OS Type 202
OS Version 202
Other Files 66
overview 43
Owner SID 101
Page Life Expectancy 216
Page Locks 170
Page Number 167
Page Pool 1 Alloc (KB) 66
Page Pool 1 Lookaside (KB) 66
Page Pool 64 Alloc (KB) 66
Page Pool 64 Lookaside (KB) 67
Page Pool 8 Alloc (KB) 67
Page Pool 8 Lookaside (KB) 67
Page Reads per Sec 220
Page Splits per Sec 220
Page Verify 129
Page Writes per Sec 220
Pages Sent Per Sec 136
Parm Type 116
Pct IO Errors Cur Intvl 239
Pct Max Locks 171, 208, 216
Pct Max Logons Active 208, 238
Pct User Connections Remaining 239
Pending Bytes For Receive I/O 232
Pending Bytes For Send I/O 232
Pending Memory Grants Count 196
Pending Message Fragments For Receive I/O 232
Pending Message Fragments For Send I/O 232
Percent Max Locks 204
Percent Process Block 208
Percent Processes Bad 191
Percent Processes Blocked 189
Percent Processes Infected 191
Percent Processes Locksleep 191
Percent Processes Othersleep 192

attributes (continued)

Percent Processes Sleeping 192
Percent Processes Stopped 192
Percent Work Complete 199
Persisted Cache Hits Per Sec 75
Persisted Cache Lookups Per Sec 75
Physical Device Name 143
Physical Device Name (Unicode) 143
Physical NetBIOS Name 213
Physical Reads (per Sec.) 209, 238
Physical Writes (per Sec.) 209, 238
Plan Cache Hit Ratio 221
Port 92
Potential In-memory Aggregation Map File (KB) 67
Potential In-memory Dimension Index (Hash) File (KB) 67
Potential In-memory Dimension Property File (KB) 67
Potential In-memory Dimension String File (KB) 68
Potential In-memory Fact Aggregation File (KB) 68
Potential In-memory Fact Data File (KB) 68
Potential In-memory Fact String File (KB) 68
Potential In-memory Map File (KB) 68
Potential In-memory Other File (KB) 68
Predictions Per Sec 53
Primary Connection Type 101
Primary Replica 95
Primary Replica ID 94
Private Buffer Memory (MB) 156
Private Buffers In Use 157
Procedure Buffers Pct Active 202
Procedure Buffers Pct Used 202
Procedure Buffers Total 204
Procedure Cache Pages 204
Procedure Cache Pct Active 202
Procedure Cache Pct Used 203
Procedure Cache Percent 204
Procedure Cache Size (KB) 203, 216
Process Holding Lock 167
Process ID 186, 246
Process Status 186, 246
Product Level 213
Program Name 188, 199
Program Name (Unicode) 186
Pyramid Operations Per Sec 76
Qualified Server Name 213
Queries Answered Per Sec 76
Queries From Cache Direct Per Sec 76
Queries From Cache Filtered Per Sec 76
Queries From Files Per Sec 76
Queries Per Sec 53
Query Age in Minutes 147
Query Exec Memory Target(MB) 197
Query Execution Count 154
Query Hash 154
Query Optimization Per Sec 249
Query Reads 199
Query Start Time 199
Query State 147
Query Status 199
Query Text 154
Query Text (Unicode) 148
Query Writes 200
Queued Jobs 163

attributes (continued)

Queued Requests 249
Quota (KB) 68
Quota Blocked 69
Read Only Access 129
Readahead Pages per Sec 221
Receive I/Os Per Sec 233
Received Messages Per Sec 105
Receives Per Sec 137
Recovery Health 107
Recovery Model 129
Recovery Queue 87
Redo Blocked Per Sec 87
Redo Bytes Per Sec 137
Redo Bytes Remaining (KB) 87
Redo Queue Size (KB) 50
Redo Queue(MB) 137
Redo Rate (KB Per Sec) 50
Redone Bytes Per Sec 87
Reduced Memory Grants Per Sec 249
Remote Server ID 193
Remote Server Name 194
Remote Server Status 194
Repl Dist Delivered Cmds per Sec 221
Repl Dist Delivered Trans per Sec 221
Repl Dist Delivery Latency 221
Repl Logreader Delivered Cmds per Sec 221
Repl Logreader Delivered Trans per Sec 222
Repl Logreader Delivery Latency 222
Repl Merge Conflicts per Sec 222
Repl Merge Downloaded Changes per Sec 222
Repl Merge Uploaded Changes per Sec 222
Replica ID 83, 85, 101, 103, 107
Replica Server Name 101, 103
Replicated Transaction Rate 132
Replicated Transactions 129
Replication Latency (ms) 129
Requestor Process ID 165
Requests Completed Per Sec 249
Resent Messages Per Sec 105
Resource Group ID 97
Resource Pool Name (Unicode) 197
Role 50, 83, 102, 107
Row Count 200
ROWNO 50, 52, 53, 55, 69, 70, 76, 80, 84, 85, 87, 90, 92, 95, 97, 99, 102, 103, 105, 108, 110, 112, 115, 117, 118, 121, 133, 140, 143, 145, 148, 151, 154, 157, 161, 164, 166, 168, 172, 182, 184, 188, 192, 194, 200, 204, 209, 214, 222, 224, 226, 229, 233, 236, 239, 243, 245, 247
Rows Converted Per Sec 70
Rows Per Sec 53
Rows Read Per Sec 70, 76
Rows Sent Per Sec 77
Rows Written Per Sec 70
Run Value 117
Running Jobs Current Interval 164
Running Query Text 200
Safe Auto-Params Per Sec 112
Sample Timestamp 117, 133, 137, 140, 143, 145, 150, 161, 164, 166, 168, 170, 172, 174, 175, 177, 180, 182, 184, 189, 192, 194, 197, 204, 209, 222, 235, 236, 239, 243, 247, 249

attributes (continued)

SAMPLENO 50, 52, 53, 55, 69, 70, 77, 80, 84, 85, 87, 90, 93, 95, 97, 99, 102, 103, 106, 108, 110, 112, 115, 117, 118, 121, 133, 140, 144, 145, 148, 151, 154, 157, 161, 164, 166, 168, 172, 182, 184, 189, 192, 194, 200, 205, 209, 215, 223, 224, 227, 229, 233, 236, 240, 243, 245, 247
Schema Name 80
Secondary Connection Type 102
Secondary Replicas Count 95
Secondary Suspension Pending 85
Select Into Bulkcopy Enabled 133
Send I/Os Per Sec 233
Send Receive Ack Time 137
Sends Per Sec 137
Sent Messages Per Sec 106
Sequence Number 175
Server 50, 52, 54, 55, 69, 70, 77, 80, 84, 85, 87, 90, 93, 95, 97, 99, 102, 104, 106, 107, 110, 113, 115, 117, 119, 121, 129, 137, 141, 144, 145, 148, 150, 151, 155, 157, 161, 164, 166, 168, 170, 172, 174, 175, 177, 180, 182, 184, 189, 192, 194, 197, 200, 205, 209, 214, 217, 224, 227, 229, 233, 235, 236, 240, 243, 245, 247, 249
Server Attention Rate 113
Server Compilations Per Sec 113
Server CPU Pct IO 217
Server CPU Pct System 189, 209
Server CPU Percent 209, 217
Server CPU Percent Application 190, 210
Server Principal Name 81
Server Re-Compilations Per Sec 113
Server Status 205, 210, 217
Server Type 174, 205, 223
Server User ID 166, 186
Server Version 144, 205, 210, 215, 217
Server Version (Superceded) 144, 205, 210, 223
Service Name 234
Service Start Mode 234
Service State 234
Service Status 234
Service Type 234
Session ID 155, 177, 200
Session Timeout 102
Severity Level 146, 181
Shared Intent Locks 172
Shared Locks 171
Shared Page Locks 172
Shared Table Locks 172
Single User Access 129
Sorting Criteria 148
Source 174
Space Used (MB) 242
SQL Cache Memory(MB) 178
SQL Handle 148, 155, 200
SQL ID 119, 121
SQL Query 119
SQL Receives Per Sec 229
SQL Sends Per Sec 229
SQL State Code 146, 181
SQLServerAgent Failed Jobs 217
Start Step ID 161
Startup Timestamp 203
State 93

attributes (continued)

Statement 81
Statistic Name 235
Stolen Pages 223
Stolen Pages Growth 223
Stored Procedure Invoked Per sec 227
String Column Name 1 121
String Column Name 2 121
String Column Name 3 122
String Column Name 4 122
String Column Name 5 122
String Column Value 1 122
String Column Value 2 122
String Column Value 3 122
String Column Value 4 122
String Column Value 5 123
Suboptimal Plans Per Sec 250
Succeeded 81
Successful Jobs 163
Suspect Database 133
Suspend Reason 83
Synchronization Health 83, 94, 107
Synchronization State 83, 138
Synchronized Commit 83
Table Count 130
Table Id 151, 242
Table Lock Escalations per Sec 172
Table Locks 171
Table Max Locks 210
Table Max Locks (Unicode) 206
Table Name 152, 167, 242
Table Name (Unicode) 168
Target Memory(MB) 197
Target Server Memory(MB) 178
Task Limit Reached Per sec 227
Tasks Aborted Per sec 227
Tasks Started Per sec 227
Tempdb Free Space (KB) 224
Text 247
Text (Unicode) 247
Threshold 174
Threshold Exceeded 174
Time Column Name 1 123
Time Column Name 2 123
Time Column Name 3 123
Time Column Name 4 123
Time Column Value 1 123
Time Column Value 2 123
Time Column Value 3 124
Time Column Value 4 124
Time Since Startup (Min.) 205, 210, 218
Timestamp 51, 52, 54, 55, 69, 70, 77, 81, 83, 86–88, 93, 94, 97, 98, 102, 104, 106, 108, 109, 113, 115, 117, 119, 124, 130, 138, 143, 146, 148, 150, 152, 155, 157, 160, 163, 165, 167, 170, 171, 174, 176–178, 181, 183, 187, 190, 194, 197, 201, 203, 206, 214, 218, 224, 227, 229, 233, 234, 236, 238, 242, 245, 247, 250
Total Active Transactions 224
Total Autoexist 55
Total Blocking Lock Requests 218
Total Bulk Mode Evaluation Nodes 55
Total Bytes Sent 77
Total Cached Bulk Mode Evaluation Nodes 55

attributes (continued)

Total Cached Other Evaluation Nodes 56
Total Cached Storage Engine Evaluation Nodes 56
Total Calculation Cache Registered 56
Total Calculations Covered 56
Total Cell By Cell Evaluation nodes 56
Total Cell By Cell Hits In Cache of Evaluation nodes 56
Total Cell By Cell Misses In Cache of Evaluation nodes 57
Total Cells Calculated 57
Total Connected Replicas 110
Total CPU Time (Sec.) 155, 187, 247
Total Databases 90, 108, 139, 210
Total Databases Critical 138
Total Databases Warning 139
Total DBs DBO Only 139
Total DBs in Error 139
Total DBs No Free Space Accounting 139
Total DBs Read Only 139
Total DBs Single User 139
Total Devices 130
Total Dimension Queries 77
Total Disconnected Replicas 111
Total Disk IO 187
Total Elapsed Time (Sec.) 201
Total Elapsed Time(ms) 149
Total Enqueued Local Messages 230
Total Enqueued Messages 230
Total Enqueued Transport Messages 230
Total Error Messages 183
Total Errors Cur Intvl 183
Total Errors High Sev 184, 211
Total Errors Other 183
Total Evaluation Nodes Calculations at Same Granularity 57
Total Evaluation Nodes Covered Single Cell 57
Total Evictions Of Evaluation Nodes 57
Total EXISTING Operators 57
Total Filestore IO Errors 69
Total Filestore IO Errors Per Sec 69
Total FileTables 245
Total Flat Cache Inserts 58
Total Forwarded Messages 230
Total Forwarded Messages Discarded 230
Total Groups 98
Total Hash Index Hits In Cache Of Evaluation Nodes 58
Total Local Failed Replicas 111
Total Local Replicas 111
Total Lock Conflicts 173, 206
Total Locks 206
Total Locks Remaining 218
Total Log Requiring Undo (KB) 88
Total Log Suspend 190
Total Logins (Per Sec.) 238
Total Logons (per Sec.) 240
Total Logouts (per Sec.) 238
Total Measure Group Queries 77
Total Memory Alloc (KB) 187
Total Network Round Trips 77
Total Non Healthy Databases 88
Total Non Healthy Groups 98
Total Non Healthy Replicas 109

attributes (continued)

Total Non Online Local Replica Groups 99
Total Non Transactional Handles Opened 152
Total Non-Online Databases 90
Total Non-Snapshot Version Transactions 225
Total Nonempty Algorithms 58
Total Nonempty Algorithms For Calculated Members 58
Total Nonempty Unoptimized Algorithms 58
Total Open Connection Count 233
Total OS CPU Pct Busy 240
Total OS CPU Percent 211, 218
Total OS IO Pct Busy 240
Total Other Sleep 190
Total Pages 223
Total Percent CPU Used 155
Total Percent Waiting 155
Total Predictions 54
Total Primary Databases 90
Total Primary Replicas 111
Total Primary Role Group 99
Total Processes 191, 211
Total Processes Bad 190
Total Processes Blocked 190, 211
Total Processes Infected 190
Total Processes Locksleep 190
Total Processes Stopped 191
Total Pyramid Operations 78
Total Queries 54
Total Queries Answered 78
Total Queries From Cache Direct 78
Total Queries From Cache Filtered 78
Total Queries From File 78
Total Recomputes 58
Total Rows 54
Total Rows Converted 70
Total Rows Read 71, 157
Total Rows Sent 78
Total Rows Written 71, 157
Total Secondary Databases 90
Total Secondary Replicas 111
Total Secondary Role Group 99
Total Server Memory (KB) 218
Total Server Memory(MB) 178
Total Since Startup 236
Total Snapshot Transactions 225
Total Sonar Subcubes 58
Total SQL Receives 230
Total SQL Sends 230
Total Storage Engine Evaluation Nodes 59
Total Subcube Hits In Cache Of Evaluation Nodes 59
Total Subcube Misses In Cache Of Evaluation Nodes 59
Total Tables 245
Total Task Limit Reached 227
Total Tasks Started 228
Total Unsynchronized Commit DB Replicas 91
Total Unsynchronized Databases 91
Total Update Snapshot Transactions 225
Total Wait Time (Sec.) 115
Transaction Delay (ms) 138
Transaction Isolation Level of Memory Optimized Tables 52
Transactions per Sec. 130

attributes (*continued*)

Type [168](#)
Undo Log Remaining (KB) [88](#)
Unsafe Auto-Params Per Sec [113](#)
Update Conflict Ratio [225](#)
Update Extent Locks [173](#)
Update Page Locks [173](#)
Used Memory(MB) [197](#)
User Defined Event ID [81](#)
User Defined Information [81](#)
Version Cleanup Rate (KB/sec) [225](#)
Version Generation Rate (KB/sec) [225](#)
Version number [162](#)
Version Store Size [226](#)
Version Store Unit Count [226](#)
Version Store Unit Creation [226](#)
Version Store Unit Truncation [226](#)
Wait Timeout (ms) [96, 97](#)
Warning Jobs Current Interval [163](#)
Workload Group Name (Unicode) [250](#)
Write Transactions per sec [130](#)
Audit Action Name attribute [79](#)
Audit Name attribute [79](#)
Audit Specification Name attribute [79](#)
Audit Specification Type attribute [79](#)
Auto-Param Attempts Per Sec attribute [111](#)
Autogrowth attribute [141](#)
Autogrowth Unit attribute [141](#)
Autogrowth Value attribute [141](#)
Availability Mode attribute [100](#)
Average CPU Time(ms) attribute [146](#)
Average Duration(ms) attribute [146](#)
Average Logical Reads attribute [146](#)
Average Logical Writes attribute [147](#)
Average Physical Reads attribute [147](#)
Average Time Per Query attribute [71](#)
Average Value (per Sec.) attribute [235](#)
Average Wait Time locks (ms) attribute [169](#)
Avg CPU Time (Sec.) attribute [152](#)
Avg Logical Reads attribute [152](#)
Avg Logical Writes attribute [152](#)
Avg Physical Reads attribute [153](#)
Avg Size Message Fragments Received attribute [231](#)
Avg Size Message Fragments Sent attribute [231](#)

B

Backup Preference attribute [94, 96](#)
Backup Priority attribute [100](#)
Batch Requests Per Sec attribute [112](#)
Blocked Database Name (Unicode) attribute [114](#)
Blocked Parent Query Text attribute [114](#)
Blocked Query Text attribute [114](#)
Blocked Session ID attribute [114](#)
Blocked Tasks attribute [248](#)
Blocked User Name (Unicode) attribute [114](#)
Blocking Process Duration(Mins.) attribute [184](#)
Blocking Process ID attribute [164, 185](#)
Blocking Query Text attribute [114](#)
Blocking Reason attribute [114](#)
Blocking Resource attribute [115](#)
Blocking Resource Duration(Mins.) attribute [187](#)
Blocking Resource Type attribute [187](#)
Blocking Session ID attribute [115, 198](#)

Blocking User Name attribute [115](#)
Broker Transaction Rollbacks attribute [228](#)
Buffer In Use attribute [156](#)
Buffer Memory (MB) attribute [156](#)
Buffer Spooled attribute [156](#)
Bytes Received Per Sec attribute [104, 133](#)
Bytes Sent Per Sec attribute [71, 104, 134](#)
Bytes Sent To Transport Per Sec attribute [104](#)

C

Cache Average Free Scan attribute [215](#)
Cache Free Buffers attribute [218](#)
Cache Hit Ratio attribute [215](#)
Cache Maximum Free Page Scan attribute [215](#)
Cache Memory Target(MB) attribute [195](#)
Calculation Cache Hits Per Sec attribute [72](#)
Calculation Cache Lookups Per Sec attribute [72](#)
Cases Per sec attribute [52](#)
Checkpoint Pages per Sec attribute [219](#)
Class Type attribute [79](#)
Cleaner Balance Per Sec attribute [60](#)
Cleaner Current Price attribute [60](#)
Cleaner Memory (KB) attribute [60](#)
Cleaner Memory nonshrinkable (KB) attribute [60](#)
Cleaner Memory shrinkable (KB) attribute [60](#)
Cleaner Memory Shrunk KB Per Sec attribute [60](#)
Client Count attribute [238](#)
Client Count Pct Used attribute [237](#)
Client Group ID attribute [165, 187](#)
Client Host Name attribute [185, 198, 246](#)
Client Process ID attribute [185, 246](#)
Client User ID attribute [165, 185](#)
Cluster Node attribute [103](#)
Cluster Resource ID attribute [96](#)
Collation attribute [211](#)
Collation ID attribute [211](#)
Collection Status attribute [207, 216](#)
Command (Unicode) attribute [188](#)
Command attribute [185](#)
Command Type attribute [198](#)
Compile Memory Target(MB) attribute [195](#)
Concurrent Data Mining Queries attribute [53](#)
Config Parameter attribute [116](#)
Config Value attribute [116](#)
Connection Level attribute [193](#)
Connection Memory(MB) attribute [177](#)
Connection State attribute [106](#)
Containment attribute [51](#)
CPU Control Effect Pct attribute [195](#)
CPU Idle Delta attribute [219](#)
CPU Idle Sec attribute [219](#)
CPU Pct Idle attribute [219](#)
CPU Rate (percentage) attribute [148](#)
CPU Time (Sec.) attribute [198](#)
CPU Usage Delta attribute [219](#)
CPU Usage Pct attribute [195, 248](#)
CPU Usage Sec attribute [219](#)
CPU Usage Target Pct attribute [196](#)
Create Date attribute [100](#)
Current Bytes For Receive I/O attribute [231](#)
Current Bytes For Send I/O attribute [231](#)
Current Cached Evaluation Nodes attribute [54](#)
Current CPU Pct Used attribute [185](#)

Current Dimension Queries attribute 72
Current Evaluation Nodes attribute 54
Current Execution Status attribute 160
Current Execution Step attribute 158
Current Interval (Sec.) attribute 139, 183, 191, 193, 220, 238
Current Locks attribute 201
Current Logons attribute 207, 237
Current Measure Group Queries attribute 72
Current Message Fragments For Send I/O attribute 231
Current Models Processing attribute 53
Current Pyramid Operations attribute 72
Current Retry Attempt attribute 158
Current Status attribute 158
Current Value attribute 236
Customized SQL Definition File attribute 117

D

dashboard 3
Data Bytes Per Sec attribute 72
Data Cache Size (KB) attribute 201, 216
Data File Names attribute 125
Data Files Freespace (MB) attribute 131
Data Freespace (MB) attribute 131
Data Freespace Percent attribute 131
Data Reads Per Sec attribute 73
data set
 attributes 49
data sets
 list of all 44
 MS SQL Additional Availability Database Details 49
 MS SQL Additional Database Detail 51
 MS SQL Analysis Services Data Mining 52
 MS SQL Analysis Services MDX Processing 54
 MS SQL Analysis Services Memory Statistics 59
 MS SQL Analysis Services Rows Processing 69
 MS SQL Analysis Services Storage Engine Query 71
 MS SQL Audit Details 79
 MS SQL Availability Database Details 81
 MS SQL Availability Database Details In Cluster 84
 MS SQL Availability Database Statistics 86
 MS SQL Availability Database Summary 88
 MS SQL Availability Group Listener Details 91
 MS SQL Availability Groups Details 93
 MS SQL Availability Groups Details In Cluster 96
 MS SQL Availability Groups Summary 97
 MS SQL Availability Replicas Details 100
 MS SQL Availability Replicas Details In Cluster 102
 MS SQL Availability Replicas Statistics 104
 MS SQL Availability Replicas Status 106
 MS SQL Availability Replicas Status Summary 108
 MS SQL Batch Stats 111
 MS SQL Blocked Queries Details 113
 MS SQL Configuration 116
 MS SQL Customized SQL Query Details 117
 MS SQL Customized SQL Query Result 119
 MS SQL Database Detail 124
 MS SQL Database Mirroring 133
 MS SQL Database Summary 138
 MS SQL Device Detail 141
 MS SQL Error Event Detail 144
 MS SQL Expensive Query Plans 146
 MS SQL Filegroup Detail 149

data sets (*continued*)
 MS SQL FileTable Detail 150
 MS SQL Individual Queries Details 152
 MS SQL Integration Service Details 156
 MS SQL Job Detail 157
 MS SQL Job Summary 162
 MS SQL Lock Conflict Detail 164
 MS SQL Lock Detail 166
 MS SQL Lock Resource Type Summary 168
 MS SQL Lock Summary 170
 MS SQL Log Shipping DB Details 173
 MS SQL Log Shipping Errors 174
 MS SQL Log Shipping Jobs Detail 176
 MS SQL Memory Manager 177
 MS SQL Problem Detail 180
 MS SQL Problem Summary 182
 MS SQL Process Detail 184
 MS SQL Process Summary 189
 MS SQL Remote Servers 193
 MS SQL Resource Pool Stats 194
 MS SQL Running Queries Details 197
 MS SQL Server Detail 201
 MS SQL Server Enterprise View 206
 MS SQL Server Properties 211
 MS SQL Server Summary 215
 MS SQL Server Transactions Summary 223
 MS SQL Service Broker Activation 226
 MS SQL Service Broker Statistics 228
 MS SQL Service Broker Transport 231
 MS SQL Services Detail 233
 MS SQL Statistics Detail 235
 MS SQL Statistics Summary 237
 MS SQL Table Detail 240
 MS SQL Table Summary 243
 MS SQL Text 246
 MS SQL Workload Group Stats 247
 overview 43
Data Size (MB) attribute 125
Database Count attribute 94
Database Growth Percent attribute 125
Database Id attribute 151, 165, 166, 241
Database ID attribute 51, 84, 173, 175, 176
Database Max blocks (Unicode) attribute 206
Database Max blocks attribute 207
Database Max Growth Size(MB) attribute 125
Database Max Locks (Unicode) attribute 206
Database Max Locks attribute 207
Database Name (Unicode) attribute 125, 141, 147, 149, 153, 165, 167, 186, 198, 226, 246
Database Name attribute 49, 51, 79, 82, 84, 86, 117, 125, 143, 165, 168, 173, 175, 176, 188, 241, 246
Database Principal Name attribute 80
Database Space Pct Used attribute 126
Database State attribute 82, 126
Database Status attribute 126
DB ID attribute 126
DB Owner (Unicode) attribute 126
DB Owner attribute 131
DBO Only Access attribute 126
Definition File Last Modified Time attribute 118
Device Free Space (MB) attribute 141
Device Free Space Percent attribute 142
Device Name (Unicode) attribute 142
Device Name attribute 143

Device Size (MB) attribute 142
Device Status attribute 142
Device Type attribute 142
Dimension Cache Hits Per Sec attribute 73
Dimension Cache Lookups Per Sec attribute 73
Dimension Index(hash) Files attribute 60
Dimension Property Files attribute 61
Dimension Queries Per Sec attribute 73
Dimension String Files attribute 61
Disk IO Current Interval attribute 239
Dump Tran Date attribute 131
Durability attribute 242
Duration (Sec.) attribute 186

E

Edition attribute 212
Edition ID attribute 214
Enabled attribute 158
Endpoint URL attribute 100
Engine Edition attribute 212
Enqueued Local Messages Per Sec attribute 228
Enqueued Messages Per Sec attribute 228
Enqueued Transport Messages Per Sec attribute 228
Error ID attribute 144, 180
Error Log Name (Unicode) attribute 203
Error Log Name attribute 203
Error Log Size (Bytes) attribute 182, 201, 207
Error Source attribute 175
Error SPID attribute 144, 180
Error Status attribute 126
Event Time attribute 80
Exclusive Extent Locks attribute 171
Exclusive Intent Locks attribute 171
Exclusive Locks attribute 170
Exclusive Page Locks attribute 171
Exclusive Table Locks attribute 171
Execution rate attribute 147
Executions Per Interval attribute 148
Extent Locks attribute 171

F

Fact Aggregation Files attribute 61
Fact Data Files attribute 61
Fact String Files attribute 61
Failed Auto-Params Per Sec attribute 112
Failed Jobs attribute 162
Failover Mode attribute 100
Failover Ready attribute 84
Failure Condition attribute 94, 96
File Bytes Received Per Sec attribute 86
File Name attribute 80
Filegroup ID attribute 149
Filegroup Max Growth Size attribute 149
Filegroup Name attribute 149
Filegroup Size attribute 149
Filegroup Space Pct Used attribute 150
Filegroup Status attribute 150
Filestore (KB) attribute 61
Filestore Clock Pages Examined Per Sec attribute 61
Filestore Clock Pages HaveRef Per Sec attribute 61
Filestore Clock Pages Valid Per Sec attribute 62

Filestore KB Reads Per Sec attribute 62
Filestore KB Writes Per Sec attribute 62
Filestore Memory Pinned (KB) attribute 62
Filestore Page Faults Per Sec attribute 62
Filestore Reads Per Sec attribute 62
Filestore Writes Per Sec attribute 63
Filestream Directory Name attribute 51
Filestream Level attribute 212
Filestream Send Rate (KB Per Sec) attribute 49
Filestream Share Name attribute 212
FileTable attribute 241
FileTable Directory attribute 151
FileTable Enabled attribute 151
Filter Rows Excluded Per Sec attribute 73
Filter Rows Included Per Sec attribute 73
Filtered Rows Per Sec attribute 73
Flat Buffer Memory (MB) attribute 156
Flat Buffers In Use attribute 156
Flat Cache Hits Per Sec attribute 74
Flat Cache Lookups Per Sec attribute 74
Flow Control Per Sec attribute 104
Flow Control Time (ms) attribute 105
Forced Parameterization Per Sec attribute 112
Forwarded Messages Discarded Per Sec attribute 228
Forwarded Pending Messages attribute 229
Fragmentation attribute 241
Free Space Accounting Suppressed attribute 127
Fully Qualified Domain Name attribute 214

G

Granted Query Memory attribute 198
Granted Workspace Memory(MB) attribute 178
Group Database ID attribute 49, 82, 84
Group ID attribute 49, 82, 85, 91, 93, 96, 100, 103, 106
Group Name attribute 82, 85, 91, 93, 97, 101, 103, 108
Group Role attribute 94

H

HADR Manager Status attribute 214
Host Name attribute 51, 95, 98, 112, 116, 127, 134, 140, 143, 144, 150, 151, 160, 163, 166, 168, 170, 171, 178, 181, 183, 188, 191, 193, 196, 204, 207, 212, 216, 224, 234, 236, 239, 242, 246, 248

I

In Memory Aggregation Map File (KB) attribute 63
In Memory Aggregation Map File KB Per Sec attribute 63
In Memory Dimension Index(hash) File (KB) attribute 63
In Memory Dimension Index(hash) File KB Per Sec attribute 63
In Memory Fact Aggregation File (KB) attribute 63
In Memory Fact Aggregation File KB Per Sec attribute 63
In Memory Fact Data File (KB) attribute 64
In Memory Fact Data File KB Per Sec attribute 64
In Memory Fact String File (KB) attribute 64
In Memory Fact String File (KB) Per Sec attribute 64
In Memory Map File (KB) attribute 64
In Memory Map File (KB) Per Sec attribute 64
In Memory Other File (KB) attribute 65
In Memory Other File (KB) Per Sec attribute 65

In Memory Property File (KB) attribute 65
In Memory Property File KB Per Sec attribute 65
In Memory String File (KB) attribute 65
In Memory String File KB Per Sec attribute 65
Index Bytes Per Sec attribute 74
Index Id attribute 241
Index Name attribute 241
Index Reads Per Sec attribute 74
Instance Name attribute 105
Instance Path attribute 214
Instance Port attribute 214
Integer Column Name 1 attribute 119
Integer Column Name 2 attribute 119
Integer Column Name 3 attribute 119
Integer Column Name 4 attribute 120
Integer Column Name 5 attribute 120
Integer Column Value 1 attribute 120
Integer Column Value 2 attribute 120
Integer Column Value 3 attribute 120
Integer Column Value 4 attribute 120
Integer Column Value 5 attribute 120
Intent Locks attribute 170
Introduction 1
IO Errors Current Interval attribute 207, 239
IO Errors Since Startup attribute 239
Is Clustered attribute 212
Is Conformant attribute 91
Is DHCP attribute 92
Is HADR Enabled attribute 214
Is Single User attribute 212
IsIntegrated Security Only attribute 213

J

Job Category ID attribute 160
Job Category Name attribute 160
Job Description attribute 160
Job Duration attribute 158
Job Error Code attribute 158
Job Execution Duration attribute 158
Job Failure Count attribute 159
Job Id attribute 159
Job Name attribute 159, 176
Job Owner attribute 159
Job Status attribute 159, 176
Job Success Count attribute 159
Job Success Rate attribute 162
Job Type attribute 160
Jobs Executed Current Interval attribute 162
Jobs Failed Current Interval attribute 162
Jobs Failed Since Startup attribute 163
Join State attribute 85, 103

K

KPIs 3

L

Last Collection Duration (Min.) attribute 243
Last Collection Start Time attribute 243
Last Connect Error Number attribute 108
Last Database Size(MB) attribute 132

Last Execution Error Message attribute 118
Last Execution Status attribute 118
Last Execution Time attribute 118, 153
Last File Processed attribute 173
Last File Processed Time attribute 173
Last Group Role attribute 95
Last Run Outcome attribute 159
Last Run Timestamp attribute 159
Lazy Writes per Sec attribute 220
License Type attribute 213
Listener ID attribute 92
Listener IP Address attribute 92
Listener Name attribute 92
Listener Subnet IP attribute 92
Local attribute 82, 107
Local Replica Recovery Health attribute 95
Lock Blocks Allocated attribute 179
Lock Blocks attribute 179
Lock Memory(MB) attribute 177
Lock Owner Blocks Allocated attribute 179
Lock Owner Blocks attribute 179
Lock Request Status attribute 167
Lock Requests per Sec. attribute 169
Lock Resource Type attribute 167, 169
Lock Timeouts per Sec. attribute 169
Lock Wait Time (ms) attribute 169
Log Bytes Flushed Per Sec attribute 132
Log Bytes Received Per Sec attribute 86, 134
Log Bytes Redone From Cache Per Sec attribute 134
Log Bytes Sent From Cache Per Sec attribute 134
Log Bytes Sent Per Sec attribute 134
Log Bytes Sent Rate (KB Per Sec) attribute 49
Log Compressed Bytes Rcvd Per Sec attribute 135
Log Compressed Bytes Sent Per Sec attribute 135
Log File Names attribute 127
Log Freespace (MB) attribute 127
Log Freespace Percent attribute 132
Log Growths attribute 127
Log Harden Time (ms) attribute 135
Log Max Growth Size(MB) attribute 127
Log Remaining For Undo(MB) attribute 135
Log Scanned For Undo(MB) attribute 135
Log Send Flow Control Time (ms) attribute 135
Log Send Queue (KB) attribute 86
Log Send Queue Size (KB) attribute 50
Log Send Queue(MB) attribute 136
Log Size (MB) attribute 127
Log Space Pct Used attribute 128
Log Time attribute 175, 176
Login Name attribute 188, 198
Long Running Transaction Name attribute 128
Long Running Transaction Process ID attribute 128
Long Running Transaction Time (Sec.) attribute 128
Longest Transaction Running Time (Sec.) attribute 224

M

Map Bytes Per Sec attribute 74
Map Files attribute 65
Map Reads Per Sec attribute 74
Max CLR Time(Sec.) attribute 153
Max CPU Time(Sec.) attribute 153
Max Disconnected Replicas attribute 109
Max Jobs Failed Interval attribute 163

Max Local Failed Replicas attribute [109](#)
Max Locks Allowed attribute [201](#)
 Max logical Reads attribute [153](#)
 Max logical Writes attribute [153](#)
 Max Memory(MB) attribute [196](#)
 Max Non Healthy Databases attribute [88](#)
 Max Non Healthy Groups attribute [98](#)
 Max Non Healthy Replicas attribute [109](#)
 Max Non Online Local Replica Groups attribute [98](#)
 Max Non-Online Databases attribute [89](#)
 Max Physical Reads attribute [154](#)
 Max Requests Cpu Time(ms) attribute [248](#)
 Max Requests Memory Grant(MB) attribute [248](#)
 Max Running Jobs Interval attribute [164](#)
 Max Secondary Replicas attribute [109](#)
 Max Unsynchonized Commit DB Replicas attribute [89](#)
 Max Unsynchonized Databases attribute [89](#)
Max User Connections Allowed attribute [237](#)
 Max Waiting Time(Sec.) attribute [154](#)
 Maximum Fragmentation Percent attribute [244](#)
 Maximum Fragmentation Percent Since Startup attribute [244](#)
 Maximum Optimizer Statistics Age (Min.) attribute [244](#)
Maximum Seen attribute [235](#)
Maximum Sev Current Interval attribute [182](#)
 Maximum Sev Level attribute [184, 208](#)
Maximum Sev Timestamp attribute [183, 208](#)
 Maximum Table Size (MB) attribute [244](#)
 Maximum Value attribute [116](#)
 Maximum Workspace Memory(MB) attribute [179](#)
 Measure Group Cache Hits Per Sec attribute [75](#)
 Measure Group Lookups Per Sec attribute [75](#)
 Measure Group Queries Per Sec attribute [75](#)
 Memory Grant Timeouts Per Sec attribute [196](#)
 Memory Grants Outstanding attribute [179](#)
 Memory Grants Pending attribute [180](#)
 Memory Grants Per Sec attribute [196](#)
 Memory Limit High (KB) attribute [66](#)
 Memory Limit Low (KB) attribute [66](#)
 Memory Optimized attribute [243](#)
 Memory Usage (KB) attribute [66](#)
 Message Age (Min.) attribute [145, 180](#)
 Message attribute [174–176](#)
 Message Fragment Received Per Sec attribute [231](#)
 Message Fragments Sent Per Sec attribute [232](#)
 Message Issuer attribute [145, 181](#)
 Message Text (Unicode) attribute [145, 180](#)
 Message Text attribute [181](#)
 Message Timestamp attribute [145, 181](#)
 Messages Sent To Transport Per Sec attribute [105](#)
 Min Disconnected Replicas attribute [109](#)
 Min Local Failed Replicas attribute [110](#)
 Min Non Healthy Databases attribute [89](#)
 Min Non Healthy Groups attribute [98](#)
 Min Non Healthy Replicas attribute [110](#)
 Min Non Online Local Replica Groups attribute [99](#)
 Min Non-Online Databases attribute [89](#)
 Min Secondary Replicas attribute [110](#)
 Min Unsynchonized Commit DB Replicas attribute [89](#)
 Min Unsynchonized Databases attribute [90](#)
 Minimum Fragmentation Percent attribute [244](#)
 Minimum Optimizer Statistics Age (Min.) attribute [244](#)
 Minimum Pct Data Freespace attribute [140, 208](#)
 Minimum Pct Log Freespace attribute [140, 208](#)
Minimum Seen attribute [235](#)
 Minimum Table Size (MB) attribute [245](#)
 Minimum Value attribute [116](#)
Mirror Device Name attribute [142](#)
 Mirror Name (Unicode) attribute [136](#)
 Mirrored Write Transaction Per Sec attribute [86](#)
 Mirrored Write Transactions Per Sec attribute [136](#)
 Mirroring Role attribute [132, 136](#)
 Mirroring State attribute [136](#)
 Mixed Page Allocations Per Sec attribute [220](#)
 Modify Date attribute [101](#)
 More Columns attribute [121](#)
 MS SQL Additional Availability Database Details data set [49](#)
 MS SQL Additional Database Detail data set [51](#)
 MS SQL Analysis Services Data Mining data set [52](#)
 MS SQL Analysis Services MDX Processing data set [54](#)
 MS SQL Analysis Services Memory Statistics data set [59](#)
 MS SQL Analysis Services Rows Processing data set [69](#)
 MS SQL Analysis Services Storage Engine Query data set [71](#)
 MS SQL Audit Details data set [79](#)
 MS SQL Availability Database Details data set [81](#)
 MS SQL Availability Database Details In Cluster data set [84](#)
 MS SQL Availability Database Statistics data set [86](#)
 MS SQL Availability Database Summary data set [88](#)
 MS SQL Availability Group Listener Details data set [91](#)
 MS SQL Availability Groups Details data set [93](#)
 MS SQL Availability Groups Details In Cluster data set [96](#)
 MS SQL Availability Groups Summary data set [97](#)
 MS SQL Availability Replicas Details data set [100](#)
 MS SQL Availability Replicas Details In Cluster data set [102](#)
 MS SQL Availability Replicas Statistics data set [104](#)
 MS SQL Availability Replicas Status data set [106](#)
 MS SQL Availability Replicas Status Summary data set [108](#)
 MS SQL Batch Stats data set [111](#)
 MS SQL Blocked Queries Details data set [113](#)
 MS SQL Configuration data set [116](#)
 MS SQL Customized SQL Query Details data set [117](#)
 MS SQL Customized SQL Query Result data set [119](#)
 MS SQL Database Detail data set [124](#)
 MS SQL Database Mirroring data set [133](#)
 MS SQL Database Summary data set [138](#)
 MS SQL Device Detail data set [141](#)
 MS SQL Error Event Detail data set [144](#)
 MS SQL Expensive Query Plans data set [146](#)
 MS SQL Filegroup Detail data set [149](#)
 MS SQL FileTable Detail data set [150](#)
 MS SQL Individual Queries Details data set [152](#)
 MS SQL Integration Service Details data set [156](#)
 MS SQL Job Detail data set [157](#)
 MS SQL Job Summary data set [162](#)
 MS SQL Lock Conflict Detail data set [164](#)
 MS SQL Lock Detail data set [166](#)
 MS SQL Lock Resource Type Summary data set [168](#)
 MS SQL Lock Summary data set [170](#)
 MS SQL Log Shipping DB Details data set [173](#)
 MS SQL Log Shipping Errors data set [174](#)
 MS SQL Log Shipping Jobs Detail data set [176](#)
 MS SQL Memory Manager data set [177](#)
 MS SQL Problem Detail data set [180](#)
 MS SQL Problem Summary data set [182](#)
 MS SQL Process Detail data set [184](#)
 MS SQL Process Summary data set [189](#)
 MS SQL Remote Servers data set [193](#)
 MS SQL Resource Pool Stats data set [194](#)
 MS SQL Running Queries Details data set [197](#)

MS SQL Server Detail data set 201
MS SQL Server Enterprise View data set 206
MS SQL Server Properties data set 211
MS SQL Server Summary data set 215
MS SQL Server Transactions Summary data set 223
MS SQL Service Broker Activation data set 226
MS SQL Service Broker Statistics data set 228
MS SQL Service Broker Transport data set 231
MS SQL Services Detail data set 233
MS SQL Statistics Detail data set 235
MS SQL Statistics Summary data set 237
MS SQL Table Detail data set 240
MS SQL Table Summary data set 243
MS SQL Text data set 246
MS SQL Workload Group Stats data set 247

N

Network Name attribute 193
Network Read Rate (per Sec.) attribute 237
Network Round Trips Per Sec attribute 75
Network Write Rate (per Sec.) attribute 237
Next Run Timestamp attribute 161
No CKPT After Recovery attribute 132
Node attribute 50, 51, 53, 55, 66, 69, 75, 80, 82, 85, 87, 88, 92, 93, 97, 98, 101, 103, 105, 107, 109, 112, 115, 116, 118, 121, 128, 136, 138, 143, 145, 147, 150, 151, 154, 156, 160, 163, 165, 167, 169, 170, 174, 175, 177, 181, 183, 186, 189, 193, 196, 198, 202, 206, 213, 216, 224, 226, 229, 232, 234, 235, 237, 241, 245, 246, 249
Non Transactional Access Level attribute 52
Notify Level Eventlog attribute 161
NT User attribute 188
Num Licenses attribute 213
Number Deadlocks per Second attribute 169
Number Of Rows attribute 241
Number of Steps attribute 161

O

Object Name attribute 80
Oldest Open Transaction(min) attribute 128
Operational State attribute 107
Optimizer Memory(MB) attribute 178
Optimizer Statistics Age (in Min.) attribute 242
OS Process ID attribute 188
OS Type attribute 202
OS Version attribute 202
Other Files attribute 66
Owner SID attribute 101

P

Page Life Expectancy attribute 216
Page Locks attribute 170
Page Number attribute 167
Page Pool 1 Alloc (KB) attribute 66
Page Pool 1 Lookaside (KB) attribute 66
Page Pool 64 Alloc (KB) attribute 66
Page Pool 64 Lookaside (KB) attribute 67
Page Pool 8 Alloc (KB) attribute 67
Page Pool 8 Lookaside (KB) attribute 67
Page Reads per Sec attribute 220
Page Splits per Sec attribute 220
Page Verify attribute 129
Page Writes per Sec attribute 220
Pages Sent Per Sec attribute 136
Parm Type attribute 116
Pct IO Errors Cur Intvl attribute 239
Pct Max Locks attribute 171, 208, 216
Pct Max Logons Active attribute 208, 238
Pct User Connections Remaining attribute 239
Pending Bytes For Receive I/O attribute 232
Pending Bytes For Send I/O attribute 232
Pending Memory Grants Count attribute 196
Pending Message Fragments For Receive I/O attribute 232
Pending Message Fragments For Send I/O attribute 232
Percent Max Locks attribute 204
Percent Process Block attribute 208
Percent Processes Bad attribute 191
Percent Processes Blocked attribute 189
Percent Processes Infected attribute 191
Percent Processes Locksleep attribute 191
Percent Processes Othersleep attribute 192
Percent Processes Sleeping attribute 192
Percent Processes Stopped attribute 192
Percent Work Complete attribute 199
Persisted Cache Hits Per Sec attribute 75
Persisted Cache Lookups Per Sec attribute 75
Physical Device Name (Unicode) attribute 143
Physical Device Name attribute 143
Physical NetBIOS Name attribute 213
Physical Reads (per Sec.) attribute 209, 238
Physical Writes (per Sec.) attribute 209, 238
Plan Cache Hit Ratio attribute 221
Port attribute 92
Potential In-memory Aggregation Map File (KB) attribute 67
Potential In-memory Dimension Index (Hash) File (KB) attribute 67
Potential In-memory Dimension Property File (KB) attribute 67
Potential In-memory Dimension String File (KB) attribute 68
Potential In-memory Fact Aggregation File (KB) attribute 68
Potential In-memory Fact Data File (KB) attribute 68
Potential In-memory Fact String File (KB) attribute 68
Potential In-memory Map File (KB) attribute 68
Potential In-memory Other File (KB) attribute 68
Predictions Per Sec attribute 53
Primary Connection Type attribute 101
Primary Replica attribute 95
Primary Replica ID attribute 94
Private Buffer Memory (MB) attribute 156
Private Buffers In Use attribute 157
Procedure Buffers Pct Active attribute 202
Procedure Buffers Pct Used attribute 202
Procedure Buffers Total attribute 204
Procedure Cache Pages attribute 204
Procedure Cache Pct Active attribute 202
Procedure Cache Pct Used attribute 203
Procedure Cache Percent attribute 204
Procedure Cache Size (KB) attribute 203, 216
Process Holding Lock attribute 167
Process ID attribute 186, 246
Process Status attribute 186, 246
Product Level attribute 213
Program Name (Unicode) attribute 186
Program Name attribute 188, 199

Pyramid Operations Per Sec attribute [76](#)

Q

Qualified Server Name attribute [213](#)

Queries Answered Per Sec attribute [76](#)

Queries From Cache Direct Per Sec attribute [76](#)

Queries From Cache Filtered Per Sec attribute [76](#)

Queries From Files Per Sec attribute [76](#)

Queries Per Sec attribute [53](#)

Query Age in Minutes attribute [147](#)

Query Exec Memory Target(MB) attribute [197](#)

Query Execution Count attribute [154](#)

Query Hash attribute [154](#)

Query Optimization Per Sec attribute [249](#)

Query Reads attribute [199](#)

Query Start Time attribute [199](#)

Query State attribute [147](#)

Query Status attribute [199](#)

Query Text (Unicode) attribute [148](#)

Query Text attribute [154](#)

Query Writes attribute [200](#)

Queued Jobs attribute [163](#)

Queued Requests attribute [249](#)

Quota (KB) attribute [68](#)

Quota Blocked attribute [69](#)

R

Read Only Access attribute [129](#)

Readahead Pages per Sec attribute [221](#)

Receive I/Os Per Sec attribute [233](#)

Received Messages Per Sec attribute [105](#)

Receives Per Sec attribute [137](#)

Recovery Health attribute [107](#)

Recovery Model attribute [129](#)

Recovery Queue attribute [87](#)

Redo Blocked Per Sec attribute [87](#)

Redo Bytes Per Sec attribute [137](#)

Redo Bytes Remaining (KB) attribute [87](#)

Redo Queue Size (KB) attribute [50](#)

Redo Queue(MB) attribute [137](#)

Redo Rate (KB Per Sec) attribute [50](#)

Redone Bytes Per Sec attribute [87](#)

Reduced Memory Grants Per Sec attribute [249](#)

Remote Server ID attribute [193](#)

Remote Server Name attribute [194](#)

Remote Server Status attribute [194](#)

Repl Dist Delivered Cmds per Sec attribute [221](#)

Repl Dist Delivered Trans per Sec attribute [221](#)

Repl Dist Delivery Latency attribute [221](#)

Repl Logreader Delivered Cmds per Sec attribute [221](#)

Repl Logreader Delivered Trans per Sec attribute [222](#)

Repl Logreader Delivery Latency attribute [222](#)

Repl Merge Conflicts per Sec attribute [222](#)

Repl Merge Downloaded Changes per Sec attribute [222](#)

Repl Merge Uploaded Changes per Sec attribute [222](#)

Replica ID attribute [83, 85, 101, 103, 107](#)

Replica Server Name attribute [101, 103](#)

Replicated Transaction Rate attribute [132](#)

Replicated Transactions attribute [129](#)

Replication Latency (ms) attribute [129](#)

Requestor Process ID attribute [165](#)

Requests Completed Per Sec attribute [249](#)

Resent Messages Per Sec attribute [105](#)

Resource Group ID attribute [97](#)

Resource Pool Name (Unicode) attribute [197](#)

Role attribute [50, 83, 102, 107](#)

Row Count attribute [200](#)

ROWNO attribute [50, 52, 53, 55, 69, 70, 76, 80, 84, 85, 87, 90, 92, 95, 97, 99, 102, 103, 105, 108, 110, 112, 115, 117, 118, 121, 133, 140, 143, 145, 148, 151, 154, 157, 161, 164, 166, 168, 172, 182, 184, 188, 192, 194, 200, 204, 209, 214, 222, 224, 226, 229, 233, 236, 239, 243, 245, 247](#)

Rows Converted Per Sec attribute [70](#)

Rows Per Sec attribute [53](#)

Rows Read Per Sec attribute [70, 76](#)

Rows Sent Per Sec attribute [77](#)

Rows Written Per Sec attribute [70](#)

Run Value attribute [117](#)

Running Jobs Current Interval attribute [164](#)

Running Query Text attribute [200](#)

S

Safe Auto-Params Per Sec attribute [112](#)

Sample Timestamp attribute [117, 133, 137, 140, 143, 145, 150, 161, 164, 166, 168, 170, 172, 174, 175, 177, 180, 182, 184, 189, 192, 194, 197, 204, 209, 222, 235, 236, 239, 243, 247, 249](#)

SAMPLENO attribute [50, 52, 53, 55, 69, 70, 77, 80, 84, 85, 87, 90, 93, 95, 97, 99, 102, 103, 106, 108, 110, 112, 115, 117, 118, 121, 133, 140, 144, 145, 148, 151, 154, 157, 161, 164, 166, 168, 172, 182, 184, 189, 192, 194, 200, 205, 209, 215, 223, 224, 227, 229, 233, 236, 240, 243, 245, 247](#)

Schema Name attribute [80](#)

Secondary Connection Type attribute [102](#)

Secondary Replicas Count attribute [95](#)

Secondary Suspension Pending attribute [85](#)

Select Into Bulkcopy Enabled attribute [133](#)

Send I/Os Per Sec attribute [233](#)

Send Receive Ack Time attribute [137](#)

Sends Per Sec attribute [137](#)

Sent Messages Per Sec attribute [106](#)

Sequence Number attribute [175](#)

Server Attention Rate attribute [113](#)

Server attribute [50, 52, 54, 55, 69, 70, 77, 80, 84, 85, 87, 90, 93, 95, 97, 99, 102, 104, 106, 107, 110, 113, 115, 117, 119, 121, 129, 137, 141, 144, 145, 148, 150, 151, 155, 157, 161, 164, 166, 168, 170, 172, 174, 175, 177, 180, 182, 184, 189, 192, 194, 197, 200, 205, 209, 214, 217, 224, 227, 229, 233, 235, 236, 240, 243, 245, 247, 249](#)

Server Compilations Per Sec attribute [113](#)

Server CPU Pct IO attribute [217](#)

Server CPU Pct System attribute [189, 209](#)

Server CPU Percent Application attribute [190, 210](#)

Server CPU Percent attribute [209, 217](#)

Server Principal Name attribute [81](#)

Server Re-Compilations Per Sec attribute [113](#)

Server Status attribute [205, 210, 217](#)

Server Type attribute [174, 205, 223](#)

Server User ID attribute [166, 186](#)

Server Version (Superceded) attribute [144, 205, 210, 223](#)

Server Version attribute [144, 205, 210, 215, 217](#)

Service Name attribute [234](#)

Service Start Mode attribute [234](#)

Service State attribute 234
Service Status attribute 234
Service Type attribute 234
Session ID attribute 155, 177, 200
Session Timeout attribute 102
Severity Level attribute 146, 181
Shared Intent Locks attribute 172
Shared Locks attribute 171
Shared Page Locks attribute 172
Shared Table Locks attribute 172
Single User Access attribute 129
Sorting Criteria attribute 148
Source attribute 174
Space Used (MB) attribute 242
SQL Cache Memory(MB) attribute 178
SQL Handle attribute 148, 155, 200
SQL ID attribute 119, 121
SQL Query attribute 119
SQL Receives Per Sec attribute 229
SQL Sends Per Sec attribute 229
SQL State Code attribute 146, 181
SQLServerAgent Failed Jobs attribute 217
Start Step ID attribute 161
Startup Timestamp attribute 203
State attribute 93
Statement attribute 81
Statistic Name attribute 235
Stolen Pages attribute 223
Stolen Pages Growth attribute 223
Stored Procedure Invoked Per sec attribute 227
String Column Name 1 attribute 121
String Column Name 2 attribute 121
String Column Name 3 attribute 122
String Column Name 4 attribute 122
String Column Name 5 attribute 122
String Column Value 1 attribute 122
String Column Value 2 attribute 122
String Column Value 3 attribute 122
String Column Value 4 attribute 122
String Column Value 5 attribute 123
Suboptimal Plans Per Sec attribute 250
Succeeded attribute 81
Successful Jobs attribute 163
Suspect Database attribute 133
Suspend Reason attribute 83
Synchronization Health attribute 83, 94, 107
Synchronization State attribute 83, 138
Synchronized Commit attribute 83

T

Table Count attribute 130
Table Id attribute 151, 242
Table Lock Escalations per Sec attribute 172
Table Locks attribute 171
Table Max Locks (Unicode) attribute 206
Table Max Locks attribute 210
Table Name (Unicode) attribute 168
Table Name attribute 152, 167, 242
Target Memory(MB) attribute 197
Target Server Memory(MB) attribute 178
Task Limit Reached Per sec attribute 227
Tasks Aborted Per sec attribute 227
Tasks Started Per sec attribute 227

Tempdb Free Space (KB) attribute 224
Text (Unicode) attribute 247
Text attribute 247
Threshold attribute 174
Threshold Exceeded attribute 174
thresholds 33
thresholds, using attributes 43
Time Column Name 1 attribute 123
Time Column Name 2 attribute 123
Time Column Name 3 attribute 123
Time Column Name 4 attribute 123
Time Column Value 1 attribute 123
Time Column Value 2 attribute 123
Time Column Value 3 attribute 124
Time Column Value 4 attribute 124
Time Since Startup (Min.) attribute 205, 210, 218
Timestamp attribute 51, 52, 54, 55, 69, 70, 77, 81, 83, 86–88, 93, 94, 97, 98, 102, 104, 106, 108, 109, 113, 115, 117, 119, 124, 130, 138, 143, 146, 148, 150, 152, 155, 157, 160, 163, 165, 167, 170, 171, 174, 176–178, 181, 183, 187, 190, 194, 197, 201, 203, 206, 214, 218, 224, 227, 229, 233, 234, 236, 238, 242, 245, 247, 250
Total Active Transactions attribute 224
Total Autoexist attribute 55
Total Blocking Lock Requests attribute 218
Total Bulk Mode Evaluation Nodes attribute 55
Total Bytes Sent attribute 77
Total Cached Bulk Mode Evaluation Nodes attribute 55
Total Cached Other Evaluation Nodes attribute 56
Total Cached Storage Engine Evaluation Nodes attribute 56
Total Calculation Cache Registered attribute 56
Total Calculations Covered attribute 56
Total Cell By Cell Evaluation nodes attribute 56
Total Cell By Cell Hits In Cache of Evaluation nodes attribute 56
Total Cell By Cell Misses In Cache of Evaluation nodes attribute 57
Total Cells Calculated attribute 57
Total Connected Replicas attribute 110
Total CPU Time (Sec.) attribute 155, 187, 247
Total Databases attribute 90, 108, 139, 210
Total Databases Critical attribute 138
Total Databases Warning attribute 139
Total DBs DBO Only attribute 139
Total DBs in Error attribute 139
Total DBs No Free Space Accounting attribute 139
Total DBs Read Only attribute 139
Total DBs Single User attribute 139
Total Devices attribute 130
Total Dimension Queries attribute 77
Total Disconnected Replicas attribute 111
Total Disk IO attribute 187
Total Elapsed Time (Sec.) attribute 201
Total Elapsed Time(ms) attribute 149
Total Enqueued Local Messages attribute 230
Total Enqueued Messages attribute 230
Total Enqueued Transport Messages attribute 230
Total Error Messages attribute 183
Total Errors Cur Intvl attribute 183
Total Errors High Sev attribute 184, 211
Total Errors Other attribute 183
Total Evaluation Nodes Calculations at Same Granularity attribute 57
Total Evaluation Nodes Covered Single Cell attribute 57

Total Evictions Of Evaluation Nodes attribute 57
Total EXISTING Operators attribute 57
Total Filestore IO Errors attribute 69
Total Filestore IO Errors Per Sec attribute 69
Total FileTables attribute 245
Total Flat Cache Inserts attribute 58
Total Forwarded Messages attribute 230
Total Forwarded Messages Discarded attribute 230
Total Groups attribute 98
Total Hash Index Hits In Cache Of Evaluation Nodes attribute 58
Total Local Failed Replicas attribute 111
Total Local Replicas attribute 111
Total Lock Conflicts attribute 173, 206
Total Locks attribute 206
Total Locks Remaining attribute 218
Total Log Requiring Undo (KB) attribute 88
Total Log Suspend attribute 190
Total Logins (Per Sec.) attribute 238
Total Logons (per Sec.) attribute 240
Total Logouts (per Sec.) attribute 238
Total Measure Group Queries attribute 77
Total Memory Alloc (KB) attribute 187
Total Network Round Trips attribute 77
Total Non Healthy Databases attribute 88
Total Non Healthy Groups attribute 98
Total Non Healthy Replicas attribute 109
Total Non Online Local Replica Groups attribute 99
Total Non Transactional Handles Opened attribute 152
Total Non-Online Databases attribute 90
Total Non-Snapshot Version Transactions attribute 225
Total Nonempty Algorithms attribute 58
Total Nonempty Algorithms For Calculated Members attribute 58
Total Nonempty Unoptimized Algorithms attribute 58
Total Open Connection Count attribute 233
Total OS CPU Pct Busy attribute 240
Total OS CPU Percent attribute 211, 218
Total OS IO Pct Busy attribute 240
Total Other Sleep attribute 190
Total Pages attribute 223
Total Percent CPU Used attribute 155
Total Percent Waiting attribute 155
Total Predictions attribute 54
Total Primary Databases attribute 90
Total Primary Replicas attribute 111
Total Primary Role Group attribute 99
Total Processes attribute 191, 211
Total Processes Bad attribute 190
Total Processes Blocked attribute 190, 211
Total Processes Infected attribute 190
Total Processes Locksleep attribute 190
Total Processes Stopped attribute 191
Total Pyramid Operations attribute 78
Total Queries Answered attribute 78
Total Queries attribute 54
Total Queries From Cache Direct attribute 78
Total Queries From Cache Filtered attribute 78
Total Queries From File attribute 78
Total Recomputes attribute 58
Total Rows attribute 54
Total Rows Converted attribute 70
Total Rows Read attribute 71, 157
Total Rows Sent attribute 78

Total Rows Written attribute 71, 157
Total Secondary Databases attribute 90
Total Secondary Replicas attribute 111
Total Secondary Role Group attribute 99
Total Server Memory (KB) attribute 218
Total Server Memory(MB) attribute 178
Total Since Startup attribute 236
Total Snapshot Transactions attribute 225
Total Sonar Subcubes attribute 58
Total SQL Receives attribute 230
Total SQL Sends attribute 230
Total Storage Engine Evaluation Nodes attribute 59
Total Subcube Hits In Cache Of Evaluation Nodes attribute 59
Total Subcube Misses In Cache Of Evaluation Nodes attribute 59
Total Tables attribute 245
Total Task Limit Reached attribute 227
Total Tasks Started attribute 228
Total Unsynchronized Commit DB Replicas attribute 91
Total Unsynchronized Databases attribute 91
Total Update Snapshot Transactions attribute 225
Total Wait Time (Sec.) attribute 115
Transaction Delay (ms) attribute 138
Transaction Isolation Level of Memory Optimized Tables attribute 52
Transactions per Sec. attribute 130
Type attribute 168

U

Undo Log Remaining (KB) attribute 88
Unsafe Auto-Params Per Sec attribute 113
Update Conflict Ratio attribute 225
Update Extent Locks attribute 173
Update Page Locks attribute 173
Used Memory(MB) attribute 197
User Defined Event ID attribute 81
User Defined Information attribute 81

V

Version Cleanup Rate (KB/sec) attribute 225
Version Generation Rate (KB/sec) attribute 225
Version number attribute 162
Version Store Size attribute 226
Version Store Unit Count attribute 226
Version Store Unit Creation attribute 226
Version Store Unit Truncation attribute 226

W

Wait Timeout (ms) attribute 96, 97
Warning Jobs Current Interval attribute 163
widgets 3
Workload Group Name (Unicode) attribute 250
Write Transactions per sec attribute 130

IBM.[®]