

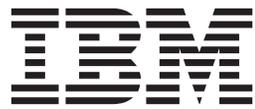
IBM Tivoli Composite Application Manager for Microsoft
Applications: Microsoft SQL Server Agent
Version 6.3

User's Guide



IBM Tivoli Composite Application Manager for Microsoft
Applications: Microsoft SQL Server Agent
Version 6.3

User's Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 405

This edition applies to version 6.3 of IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent (product number 5724-U17) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2005, 2012.**

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

Figures

- | | | | |
|---|----|--|----|
| 1. The Configure Database Agents window | 23 | 2. The Database Server Properties window | 24 |
|---|----|--|----|

Tables

1.	When data is collected and cache duration for attribute groups	6	24.	Trace log files for troubleshooting agents	272
2.	Key aspects of data gathering for attribute groups	9	25.	Problems and solutions for installation and configuration on UNIX and Windows systems	278
3.	Data that is accessed by attribute groups that gather data from operating system or Microsoft SQL Server files.	13	26.	Problems and solutions for installation and configuration of the Microsoft SQL Server agent	280
4.	Additional authorization	21	27.	General problems and solutions for uninstallation	285
5.	Names and descriptions of configuration settings for each interface	26	28.	SQL Server problems and solutions	286
6.	Interfaces for starting and stopping the Microsoft SQL Server agent locally and remotely	36	29.	General agent problems and solutions	287
7.	Cross-linked workspaces within the SQL Server agent	58	30.	Remote deployment problems and solutions	296
8.	Capacity planning for historical data logged by the Microsoft SQL Server agent	223	31.	Workspace problems and solutions	297
9.	User ID and password combinations	240	32.	Specific situation problems and solutions	301
10.	Index and table updates	247	33.	Problems with configuring situations that you solve in the Situation Editor	302
11.	Index and table updates	249	34.	Problems with configuration of situations that you solve in the Workspace area	303
12.	Index and table updates	251	35.	Problems with configuration of situations that you solve in the Manage Tivoli Enterprise Monitoring Services window	304
13.	Top Processes by CPU report	258	36.	Take Action commands problems and solutions	304
14.	Server Transactions Details report.	259	37.	Tivoli Common Reporting for Microsoft SQL Server agent problems and solutions.	305
15.	Wait Time and Requests of Locks report	260	38.	Time periods and suffixes for summary tables and views.	309
16.	Space Utilization by Database report	261	39.	Additional columns to report summarization information	310
17.	Free Space Availability report	262	40.	Overview of attribute groups to event classes and slots	314
18.	Database Status Availability report	263			
19.	CPU Utilization report	264			
20.	Processes Details report	265			
21.	Server Lock Details report	266			
22.	Servers with Most Failed Jobs report	267			
23.	Information to gather before contacting IBM Software Support	269			

Contents

Figures	iii	Availability Replicas and Database Information workspace	46
Tables	v	Availability Replicas Status workspace	46
Chapter 1. Overview of the agent	1	Databases Information workspace	46
IBM Tivoli Monitoring	2	Database Mirror Detail workspace	46
New in this release	2	Device Detail workspace	47
Components of the IBM Tivoli Monitoring environment	4	Enterprise Database Mirror Detail workspace	47
Agent Management Services	5	Enterprise Database Summary workspace	47
User interface options	5	Enterprise Errorlog Alerts workspace	47
Microsoft SQL Server agent data collection	6	Enterprise Log Shipping Database Details workspace	47
Collection processes	6	Enterprise Processes Summary workspace	47
Attribute groups that gather data from Microsoft SQL Server SELECT statements and stored procedures	8	Enterprise Server Locking workspace	47
Attribute groups that gather data from operating system or MS SQL Server files	13	Enterprise Servers Overview workspace	48
Chapter 2. Agent installation and configuration	15	Enterprise Server Statistics workspace	48
Installing language packs	15	Enterprise Server Summary workspace	48
Windows systems	15	Errorlog Alerts workspace	48
UNIX or Linux systems	16	Filegroup Detail workspace	48
Prerequisites checking	16	FileTable Detail workspace	49
Running as a non-administrator user	17	Job Detail workspace	49
Agent-specific installation and configuration	18	Job Summary workspace	49
Granting permissions	19	Lock Resource Type Summary workspace	49
Local configuration	23	Log and Space Information workspace	49
Local environment variables	32	Log Shipping Database Details workspace	49
Remote configuration	34	Log Shipping Error Details workspace	49
Reconfiguration	35	Log Shipping Job Details workspace	50
Starting and stopping the Microsoft SQL Server agent	35	MS SQL Server Historical Summarized Availability workspace	50
Setting up the Microsoft SQL Server agent in a cluster environment	36	MS SQL Server Historical Summarized Availability Daily workspace	50
Chapter 3. Workspaces reference	41	MS SQL Server Historical Summarized Availability Hourly workspace	51
Predefined workspaces	42	MS SQL Server Historical Summarized Availability Weekly workspace	51
Additional Availability Database Details workspace	44	MS SQL Server Historical Summarized Capacity workspace	52
Additional Database Details workspace	44	MS SQL Server Historical Summarized Capacity Daily workspace	52
Analysis Services Data Mining Statistics workspace	44	MS SQL Server Historical Summarized Capacity Hourly workspace	52
Analysis Services Memory Statistics workspace	44	MS SQL Server Historical Summarized Capacity Weekly workspace	53
Analysis Services Storage Engine Query Processing workspace	44	MS SQL Server Historical Summarized Performance workspace	53
Analysis Services Rows Processing workspace	44	MS SQL Server Historical Summarized Performance Daily workspace	53
Availability Database Information workspace	45	MS SQL Server Historical Summarized Performance Hourly workspace	53
Availability Groups Information workspace	45	MS SQL Server Historical Summarized Performance Weekly workspace	54
Availability Groups Information In Cluster workspace	45	MS SQL Server Overview	54
Availability Replicas Information workspace	46	Process Holding Lock SQL Text workspace	54
		Processes workspace	54
		Server Batch Statistics workspace	55
		Server Configuration workspace	55
		Server Detail workspace	55

Services Detail workspace	55
Server Locking workspace	55
Server Statistics workspace	56
Server Summary workspace	56
SQL Server CPU Bottleneck workspace	56
SQL Server Disk Space and Logon Statistics	56
SQL Server Integration Service Details workspace	56
SQL Server Logical Disk Bottleneck workspace	56
SQL Server Memory Bottleneck workspace	57
SQL Server Physical Disk Bottleneck workspace	57
SQL Server Service Broker Information workspace	57
SQL Server Transactions Summary workspace	57
SQL Text for Process ID workspace	58
Table Detail workspace	58
Table Information workspace	58
Cross-linked workspaces	58

Chapter 4. Attributes reference 59

Attributes in each attribute group	59
MS SQL Additional Database Details attributes	60
MS SQL Additional Availability Database Details attributes	62
MS SQL Analysis Services Data Mining attributes	63
MS SQL Analysis Services MDX Processing attributes	65
MS SQL Analysis Services Memory Statistics attributes	69
MS SQL Analysis Services Rows Processing attributes	77
MS SQL Analysis Services Storage Engine Query attributes	78
MS SQL Availability Database Details attributes	85
MS SQL Availability Database Details In Cluster attributes	90
MS SQL Availability Database Statistics attributes	91
MS SQL Availability Database Summary attributes	93
MS SQL Availability Groups Details attributes	96
MS SQL Availability Groups Details In Cluster attributes	99
MS SQL Availability Group Listener Details attributes	101
MS SQL Availability Groups Summary attributes	103
MS SQL Availability Replicas Details attributes	105
MS SQL Availability Replicas Details In Cluster attributes	109
MS SQL Availability Replicas Statistics attributes	110
MS SQL Availability Replicas Status attributes	112
MS SQL Availability Replicas Status Summary attributes	115
MS SQL Batch Stats attributes	118
MS SQL Blocked Queries Details attributes	119
MS SQL Configuration attributes	121
MS SQL Database Detail attributes	122
MS SQL Database Mirroring attributes	131
MS SQL Database Summary attributes	135
MS SQL Device Detail attributes	136
MS SQL Enterprise Database Mirroring attributes	140

MS SQL Enterprise Log Shipping DB Details attributes	144
MS SQL Error Event Details attributes	145
MS SQL Filegroup Detail attributes	146
MS SQL FileTable Detail attributes	147
MS SQL Individual Queries Details attributes	148
MS SQL Integration Service Details attributes	151
MS SQL Job Detail attributes	152
MS SQL Job Summary attributes	156
MS SQL Lock Conflict Detail attributes	158
MS SQL Lock Detail attributes	159
MS SQL Lock Resource Type Summary attributes	162
MS SQL Lock Summary attributes	163
MS SQL Log Shipping DB Details attributes	165
MS SQL Log Shipping Errors attributes	166
MS SQL Log Shipping Jobs Detail attributes	166
MS SQL Memory Manager attributes	167
MS SQL Problem Detail attributes	170
MS SQL Problem Summary attributes	171
MS SQL Process Detail attributes	172
MS SQL Process Summary attributes	176
MS SQL Remote Servers attributes	178
MS SQL Resource Pool Stats attributes	179
MS SQL Running Queries Details attributes	181
MS SQL Service Broker Activation attributes	185
MS SQL Service Broker Statistics attributes	186
MS SQL Service Broker Transport attributes	188
MS SQL Services Detail attributes	190
MS SQL Server Detail attributes	192
MS SQL Server Properties attributes	195
MS SQL Server Summary attributes	198
MS SQL Server Transactions Summary attributes	203
MS SQL Statistics Detail attributes	205
MS SQL Statistics Summary attributes	206
MS SQL Table Detail attributes	208
MS SQL Table Summary attributes	210
MS SQL Text attributes	212
MS SQL Workload Group Stats attributes	213
VMS SQL Remote Servers attributes	215
VMS SQL Server Enterprise View attributes	217

Chapter 5. Situations reference. 227

Predefined situations	228
Situation descriptions	230
MS_SQL_Block_Critical	230
MS_SQL_Block_Warning	230
MS_SQL_Cache_Ave_FreePage_Crit	230
MS_SQL_Cache_Ave_FreePage_Warn	230
MS_SQL_Cache_Hit_Ratio_Crit	230
MS_SQL_Cache_Hit_Ratio_Warn	230
MS_SQL_Cache_Max_FreePage_Crit	230
MS_SQL_Cache_Max_FreePage_Warn	230
MS_SQL_Client_Cnt_Pct_Used_Crit	230
MS_SQL_Client_Cnt_Pct_Used_Warn	230
MS_SQL_Collection_Status_Warning	230
MS_SQL_CPU_Critical	231
MS_SQL_CPU_Warning	231
MS_SQL_DB_Error_Status	231
MS_SQL_DB_Freespace_Critical	231
MS_SQL_DB_Join_State_Warn	231

MS_SQL_DB_Space_Pct_Used_Warn.	231
MS_SQL_DB_Space_Pct_Used_Crit	231
MS_SQL_DB_Num_Errors_Warning	231
MS_SQL_DB_Freespace_Warning	231
MS_SQL_DB_Status_Crit	231
MS_SQL_DB_Suspect_Crit	231
MS_SQL_DB_Sync_Health_Crit	231
MS_SQL_DB_Sync_State_Warn	231
MS_SQL_Device_Free_PCT_Warning	231
MS_SQL_Error_Event_Warning	232
MS_SQL_Error_Warning	232
MS_SQL_ErrorLog_Size_Warning	232
MS_SQL_Failed_Jobs_Crit	232
MS_SQL_Fragmentation_Crit	232
MS_SQL_Fragmentation_Warn	232
MS_SQL_Group_Role_Crit	232
MS_SQL_IOERR_Startup_Warning	232
MS_SQL_IO_Disk_Errors_Crit	232
MS_SQL_IOError_Curintvl_Warning	232
MS_SQL_Log_Freespace_Critical	232
MS_SQL_Log_Freespace_Warning	232
MS_SQL_Log_Shipping_Status_Err	233
MS_SQL_Log_Shipping_Warning	233
MS_SQL_Log_Space_Pct_Used_Crit	233
MS_SQL_Log_Space_Pct_Used_Warn	233
MS_SQL_Log_Suspend_Warning	233
MS_SQL_LogonPct_Crit	233
MS_SQL_LogonPct_Warning	233
MS_SQL_Long_Running_Transaction_Warning	233
MS_SQL_Mirror_Failover_Critical	233
MS_SQL_Mirror_Lagging_Warning	233
MS_SQL_Mirror_Suspended_Warning	233
MS_SQL_Network_Read_Rate_Crit	234
MS_SQL_Network_Read_Rate_Warn	234
MS_SQL_Network_Write_Rate_Crit	234
MS_SQL_Network_Write_Rate_Warn	234
MS_SQL_Num_Process_Blocked_Crit	234
MS_SQL_Num_Process_Blocked_Warn	234
MS_SQL_Oldest_Transaction_Crit	234
MS_SQL_Proc_Buffs_Used_Crit	235
MS_SQL_Proc_Buffs_Used_Warn	235
MS_SQL_Proc_Cache_Active_Crit	235
MS_SQL_Proc_Cache_Active_Warn	235
MS_SQL_Proc_Cache_Used_Crit	235
MS_SQL_Proc_Cache_Used_Warn	236
MS_SQL_Process_Blocked_Warning	236
MS_SQL_Process_Infected_Warning	236
MS_SQL_Processes_Bad_Warning	236
MS_SQL_Processes_Stop_Warning	236
MS_SQL_ProcessLockSleep_Warning	236
MS_SQL_ProcessOthSleep_Warning	236
MS_SQL_Rem_Serv_Stat_Critical	236
MS_SQL_Repl_Latency_Crit	236
MS_SQL_Repl_Latency_Warn	236
MS_SQL_Replica_Join_State_Warn	236
MS_SQL_Replica_Role_Health_Crit	236
MS_SQL_Replica_Sync_Health_Warn	237
MS_SQL_Replicas_Conn_State_Crit	237
MS_SQL_Status_Critical	237
MS_SQL_Status_Inactive	237
MS_SQL_Total_Locks_Critical	237

MS_SQL_Total_Locks_Warning	237
----------------------------	-----

Chapter 6. Take Action commands

reference 239

Predefined Take Action commands	239
Database Check-up and Tune-up action	240
Dump Database action	241
Dump Transaction Log action	244
Rebuild Fragmented Indexes action	246
Refresh Query Optimizer Statistics action	248
Start SQL Server action	249
Stop SQL Server action	250
Update Space Usage Information action	250

Chapter 7. Policies reference. 253

Predefined policies	253
Policy descriptions	253
MS_SQL_DB_free_space_critical	253

Chapter 8. Tivoli Common Reporting

for the monitoring agent 255

Cognos-based report packages.	255
Prerequisites.	255
Importing Cognos report packages	257
Cognos data models and reports	258

Chapter 9. Troubleshooting 269

Gathering product information for IBM Software Support	269
Built-in troubleshooting features	269
Problem classification.	270
Trace logging	270
Overview of log file management	270
Examples of trace logging	270
Principal trace log files	271
Enabling detailed tracing in the collector trace log	274
Option one: modify the settings file	274
Option two: modify environment variables	274
RAS trace parameters	275
Setting RAS trace parameters by using the GUI	275
Manually setting RAS trace parameters.	276
Problems and workarounds	277
Installation and configuration troubleshooting	277
Troubleshooting for SQL Server	286
Agent troubleshooting	287
Troubleshooting for remote deployment	296
Workspace troubleshooting.	297
Situation troubleshooting	301
Take Action command troubleshooting	304
Tivoli Common Reporting troubleshooting	305
Support information	307

Appendix A. Upgrading for warehouse

summarization 309

Tables in the warehouse.	309
Effects on summarized attributes	309
Upgrading your warehouse with limited user permissions	310

Appendix B. Event mapping 313

**Appendix C. Discovery Library
Adapter for the Microsoft SQL Server
agent 385**

DLA data model class types represented in CDM 385
DLA data model classes for Microsoft SQL Server
agent 385
 SqlServer class 386
 SqlServerDatabase class 387
 SqlServerTable class 388
 SqlServerProcess class 390
 TMSAgent class 391

**Appendix D. Integration with Tivoli
Business Service Manager 395**

Components for integrating with Tivoli Business
Service Manager 395
Tasks to integrate the agent with Tivoli Business
Service Manager 396
 Installing the Discovery Library Toolkit on the
 Tivoli Business Service Manager 396

Configuring the Tivoli Event Integration Facility
(EIF) probe to enrich events 396
Creating a service in Tivoli Business Service
Manager 397
Creating a data source mapping for each data
source. 397
Configuring additional IBM Tivoli Monitoring
web services. 398
Viewing data in the Tivoli Enterprise Portal . . 398

Appendix E. Documentation library 399

Prerequisite publications. 399
Related publications 400
Other sources of documentation 400

Accessibility. 403

Notices 405

Trademarks 407

Index 409

Chapter 1. Overview of the agent

The IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent provides you with the capability to monitor Microsoft SQL Server. You can also use the agent to take basic actions with the Microsoft SQL Server.

IBM® Tivoli® Monitoring is the base software for the Microsoft SQL Server agent.

The Microsoft SQL Server agent monitors the following functions:

- Availability and resources
- Performance
- Error and event log
- Historical data

Features of the monitoring agent

The Microsoft SQL Server agent offers a central point of management for distributed databases. The monitoring agent provides a comprehensive means for gathering exactly the information you need to detect problems early and prevent them. Information is standardized across all systems so you can monitor hundreds of servers from a single workstation. You can easily collect and then analyze specific information using the Tivoli Enterprise Portal interface.

The Microsoft SQL Server agent is an intelligent, remote monitoring agent that resides on managed systems. It assists you in anticipating trouble and warns system administrators when critical events take place on systems. With the Microsoft SQL Server agent, database and system administrators can set threshold levels as desired and can set flags to alert them when the system reaches these thresholds.

This monitoring agent provides the following benefits:

- Simplifies application and system management by managing applications, operating systems, and resources across your environment.
- Increases profits by providing you with real-time access to reliable, up-to-the-minute data that allows you to make faster, better-informed operating decisions.
- Scales and ports to new Windows operating systems.
- Improves system performance by letting you integrate, monitor, and manage your system, network, console, and mission-critical applications. The monitoring agent sends an alert when conditions on the system network meet threshold-based conditions. These alerts notify your system administrator to limit and control database usage. You can view data gathered by monitoring agents in reports and charts for the status of your distributed database systems.
- Enhances efficiency by monitoring different Microsoft SQL Server versions on separate systems and networks from a single PC screen. Depending on your configuration, you can collect and monitor data across systems. The monitoring agent gathers and filters status information at the managed system rather than at the hub, eliminating unnecessary data transmission and sending only data that is relevant to changes in status conditions. The Microsoft SQL Server agent helps you to monitor and to gather the consistent, accurate, and timely information you need to effectively perform your job.

IBM Tivoli Monitoring

IBM Tivoli Monitoring provides a way to monitor the availability and performance of all the systems in your enterprise from one or several designated workstations. It also provides useful historical data that you can use to track trends and to troubleshoot system problems.

You can use IBM Tivoli Monitoring to achieve the following tasks:

- Monitor for alerts on the systems that you are managing by using predefined situations or custom situations.
- Establish your own performance thresholds.
- Trace the causes leading to an alert.
- Gather comprehensive data about system conditions.
- Use policies to take actions, schedule work, and automate manual tasks.

The Tivoli Enterprise Portal is the interface for IBM Tivoli Monitoring products. You can use the consolidated view of your environment as seen in the Tivoli Enterprise Portal to monitor and resolve performance issues throughout the enterprise.

See the IBM Tivoli Monitoring publications listed in “Prerequisite publications” on page 399 for complete information about IBM Tivoli Monitoring and the Tivoli Enterprise Portal.

New in this release

For version 6.3 of the Microsoft SQL Server agent, the following enhancements have been made since version 6.2.3, including the fix packs:

- Additional supported operating systems as listed in the Prerequisites topic for the Microsoft SQL Server agent in the IBM Tivoli Composite Application Manager for Microsoft Applications Information Center
- New attribute groups:
 - MS SQL Analysis Services Data Mining
 - MS SQL Analysis Services MDX Processing
 - MS SQL Analysis Services Memory Statistics
 - MS SQL Analysis Services Rows Processing
 - MS SQL Analysis Services Storage Engine Query
 - MS SQL Availability Database Details
 - MS SQL Availability Database Details In Cluster
 - MS SQL Availability Database Statistics
 - MS SQL Availability Database Summary
 - MS SQL Availability Groups Details attributes
 - MS SQL Availability Groups Details In Cluster
 - MS SQL Availability Group Listener Details
 - MS SQL Availability Groups Summary
 - MS SQL Availability Replicas Details
 - MS SQL Availability Replicas Details In Cluster
 - MS SQL Availability Replicas Statistics
 - MS SQL Availability Replicas Status
 - MS SQL Availability Replicas Status Summary
 - MS SQL Blocked Queries Details
 - MS SQL Running Queries Details

- MS SQL FileTable Detail
- MS SQL Individual Queries Details
- MS SQL Running Queries Details
- MS SQL Service Broker Activation
- MS SQL Service Broker Statistics
- MS SQL Service Broker Transport
- MS SQL Server Transactions Summary
- MS SQL Table Summary
- MS SQL Additional Database Details
- MS SQL Additional Availability Database Details
- MS SQL Server Properties
- Modified attribute groups:
 - MS SQL Server Summary
 - MS SQL Statistics Summary
 - MS SQL Job details
 - MS SQL Table Detail
 - MS SQL Lock Detail
 - MS SQL Server Detail
 - MS SQL Device Detail
 - MS SQL Server Properties
 - MS SQL Table Detail
 - MS SQL Filegroup Detail
- New navigator items:
 - AlwaysOn
 - Services
- New workspaces:
 - SQL Server Transactions Summary
 - Availability Groups Information
 - Availability Replicas Status
 - Availability Database Information
 - Availability Replicas Information
 - Additional Availability Database Details
 - Availability Replicas and Database Information
 - SQL Server Service Broker Information
 - SQL Server Integration Service Details
 - Services Detail
 - Table Information
 - Analysis Services Data Mining Statistics
 - Analysis Services Memory Statistics
 - Analysis Services Storage Engine Query Processing
 - Analysis Services Rows Processing
 - Query Details
 - Individual Queries Details
 - FileTable Detail
 - Additional Database Detail

- Availability Groups Information In Cluster
- New predefined situations:
 - MS_SQL_Replica_Join_State_Warn
 - MS_SQL_DB_Join_State_Warn
 - MS_SQL_Group_Role_Crit
 - MS_SQL_DB_Sync_Health_Crit
 - MS_SQL_Replicas_Conn_State_Crit
 - MS_SQL_Replica_Role_Health_Crit
 - MS_SQL_Replica_Sync_Health_Warn
 - MS_SQL_DB_Sync_State_Warn
- New environment variables:
 - COLL_JOB_DISABLED
 - COLL_KOQFGRPD_MAX_ROW
 - COLL_KOQDEVD_MAX_ROW
 - COLL_KOQDBD_INTERVAL
 - COLL_DBD_FRENAME_RETRY_CNT
- Support for handling multilingual error log file in the MS SQL Problem Detail attribute group
- New koqErrConfig.ini file, which is saved at *candle_home*\TMAITM6 directory for a 32-bit agent, and at *candle_home*\TMAITM6_x64 for a 64-bit agent
- New Cognos® data models and Cognos-based reports
- Added the self-describing agent support for the monitoring agent to generate agent support files as part of the agent image
- Added the prerequisites checking utility to verify prerequisites before the agent is installed or upgraded
- Integration with Tivoli Business Service Manager V6.1 or later
- Updated koq.baroc file to support Tivoli Enterprise Console® event mapping

Components of the IBM Tivoli Monitoring environment

After you install and set up the Microsoft SQL Server agent, you have an environment that contains the client, server, and monitoring agent implementation for Tivoli Monitoring.

This Tivoli Monitoring environment contains the following components:

Tivoli Enterprise Portal client

The portal has a user interface based on Java™ for viewing and monitoring your enterprise.

Tivoli Enterprise Portal Server

The portal server is placed between the client and the Tivoli Enterprise Monitoring Server and enables retrieval, manipulation, and analysis of data from the monitoring agents. The Tivoli Enterprise Portal Server is the central repository for all user data.

Tivoli Enterprise Monitoring Server

The monitoring server acts as a collection and control point for alerts received from the monitoring agents, and collects their performance and availability data. The Tivoli Enterprise Monitoring Server is also a repository for historical data.

Tivoli Enterprise Monitoring Agent, Microsoft SQL Server agent

This monitoring agent collects data and distributes the data to the Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Tivoli Enterprise Portal, Tivoli Data Warehouse, and Tivoli Integrated Portal.

Multiple copies of this agent can run on the same system.

IBM Tivoli Netcool/OMNIBus

Tivoli Netcool/OMNIBus is an optional component and the recommended event management component. The Netcool/OMNIBus software is a service level management (SLM) system that delivers real-time, centralized monitoring of complex networks and IT domain events. Event information is tracked in a high-performance, in-memory database and presented to specific users through individually configurable filters and views. The software includes automation functions that you can use to perform intelligent processing on managed events. You can use this software to forward events for Tivoli Monitoring situations to Tivoli Netcool/OMNIBus.

IBM Tivoli Enterprise Console

The Tivoli Enterprise Console is an optional component that acts as a central collection point for events from various sources, including events from other Tivoli software applications, Tivoli partner applications, custom applications, network management platforms, and relational database systems. You can view these events through the Tivoli Enterprise Portal (by using the event viewer), and you can forward events from Tivoli Monitoring situations to the Tivoli Enterprise Console component. If you do not already use Tivoli Enterprise Console and need an event management component, you can choose to use IBM Tivoli Netcool/OMNIBus.

IBM Tivoli Common Reporting

Tivoli Common Reporting is a separately installable feature available to users of Tivoli software that provides a consistent approach to generating and customizing reports. Some individual products provide reports that are designed for use with Tivoli Common Reporting and have a consistent look and feel.

IBM Tivoli Application Dependency Discovery Manager (TADDM)

TADDM delivers automated discovery and configuration tracking capabilities to build application maps that provide real-time visibility into application complexity.

IBM Tivoli Business Service Manager

The Tivoli Business Service Manager component delivers real-time information to help you respond to alerts effectively based on business requirements. Optionally, you can use this component to meet service-level agreements (SLAs). Use the Tivoli Business Service Manager tools to help build a service model that you can integrate with Tivoli Netcool/OMNIBus alerts or optionally integrate with data from an SQL data source. Optional components provide access to data from other IBM Tivoli applications such as Tivoli Monitoring and TADDM.

Agent Management Services

You can use IBM Tivoli Monitoring Agent Management Services to manage the Microsoft SQL Server agent. These services are available in the following IBM Tivoli Monitoring OS agents: Windows, Linux, and UNIX. The services are designed to keep the Microsoft Active Directory agent available, and to provide information about the status of the product to the Tivoli Enterprise Portal. IBM Tivoli Monitoring V6.2.2, Fix Pack 2 or later provides support for Agent Management Services. For more information about Agent Management Services, see the *IBM Tivoli Monitoring Administrator's Guide*, "Agent Management Services" chapter.

User interface options

Installation of the base software and other integrated applications provides the following interfaces that you can use to work with your resources and data:

Tivoli Enterprise Portal user interface

The client interface is a graphical user interface (GUI) based on Java on a Windows or Linux workstation. You can run the Tivoli Enterprise Portal as a desktop application or a browser application. The browser application is automatically installed with the Tivoli Enterprise Portal Server. The desktop application is installed using the Tivoli Monitoring installation media or with

Java Web Start. Tivoli Enterprise Portal Server. To start the Tivoli Enterprise Portal browser client in your Internet browser, enter the URL for a specific Tivoli Enterprise Portal browser client installed on your Web server.

Command line interface

You can use IBM Tivoli Monitoring commands to manage the Tivoli Monitoring components and their configuration. There are also commands that you can run at the Tivoli Enterprise Console event server or the Tivoli Netcool/OMNIBus ObjectServer to configure event synchronization for enterprise situations.

Manage Tivoli Enterprise Monitoring Services window

The window for the Manage Tivoli Enterprise Monitoring Services utility is used for configuring the agent and starting Tivoli services not already designated to start automatically.

IBM Tivoli Enterprise Console

An event management application that integrates system, network, database, and application management to help ensure the optimal availability of an IT service for an organization.

Microsoft SQL Server agent data collection

This section provides information about how the monitoring agent collects data, including the following topics:

- “Collection processes”
- “Attribute groups that gather data from Microsoft SQL Server SELECT statements and stored procedures” on page 8
- “Attribute groups that gather data from operating system or MS SQL Server files” on page 13

Collection processes

The Microsoft SQL Server agent gathers data when it is requested by a workspace refresh, a situation sampling of attributes, or historical data collection. At that time, the monitoring agent gathers all attributes in the attribute groups that make up a workspace or situation. The default refresh or sampling intervals do not put a significant load on the system or the Microsoft SQL Server as it gathers the data. Most of the attributes gathered by the Microsoft SQL Server agent come from monitoring data provided by Windows Performance Monitor (Perfmon). This data is acquired through the use of requests for Perfmon data and SQL database SELECT statements. The Perfmon and SQL data are gathered on a defined cycle while most other attributes are gathered on demand when requested through a workspace refresh, a situation sampling of attributes, or historical data collection. Other attribute groups contain data that is collected from system information or external Microsoft SQL Server data such as the database instance log. This data is primarily gathered on demand, as it is requested. The Microsoft SQL Server agent attempts to lessen the possibility of over sampling by maintaining a cache of attribute data for a set length of time. If the data is sampled within a shorter duration than the interval set for the cache, the cache data is considered fresh and the collector does not acquire new data.

Table 1 provides information about when the attribute data is gathered (on demand or cyclically) and the cache duration for each Microsoft SQL Server agent attribute group.

Table 1. When data is collected and cache duration for attribute groups

Attribute group	When data is collected (cycle time or on demand)	Cache duration in seconds
MS SQL Batch Stats	Every 180 seconds	No caching
MS SQL Configuration	On demand	90
MS SQL Database Detail	On demand	10
MS SQL Database Mirroring	On demand	No caching
MS SQL Database Summary	Every 180 seconds	No caching

Table 1. When data is collected and cache duration for attribute groups (continued)

Attribute group	When data is collected (cycle time or on demand)	Cache duration in seconds
MS SQL Device Detail	On demand	No caching
MS SQL Enterprise Log Shipping DB Details	On demand	No caching
MS SQL Filegroup	On demand	No caching
MS SQL Job Detail	Every 180 seconds	No caching
MS SQL Job Summary	On demand	No caching
MS SQL Lock Conflict Detail	On demand	No caching
MS SQL Lock Detail	On demand	No caching
MS SQL Lock Resource Type Summary	Every 300 seconds	No caching
MS SQL Lock Summary	On demand	No caching
MS SQL Log Shipping DB Details	On demand	No caching
MS SQL Log Shipping Errors	On demand	No caching
MS SQL Log Shipping Jobs Detail	On demand	No caching
MS SQL Memory Manager	On demand	No caching
MS SQL Problem Detail	Every 180 seconds	No caching
MS SQL Process Detail	On demand	10
MS SQL Process Summary	Every 180 seconds	No caching
MS SQL Problem Summary	Every 180 seconds	No caching
MS SQL Remote Servers	Every 60 seconds	No caching
MS SQL Resource Pool Stats	On demand	No caching
MS SQL Server Detail	On demand	60
MS SQL Server Enterprise View	On demand	No caching
MS SQL Server Summary	Every 180 seconds	No caching
MS SQL Services Detail	On demand	No caching
MS SQL Statistics Detail	On demand	300
MS SQL Statistics Summary	Every 60 seconds	No caching
MS SQL Table Detail	On demand	120
MS SQL Text	On demand	No caching
MS SQL Workload Group Stats	On demand	No caching
VMS SQL Remote Servers	On demand	No caching
VMS SQL Server Enterprise View	On demand	No caching
MS SQL Server Transaction Summary	On demand	10
MS SQL Service Broker Activation	On demand	10
MS SQL Service Broker Statistics	On demand	10
MS SQL Service Broker Transport	On demand	10
MS SQL Table Summary	On demand	120
MS SQL Integration Service Details	On demand	No caching
MS SQL Analysis Services Data Mining	On demand	No caching
MS SQL Analysis Services Memory statistics	On demand	No caching

Table 1. When data is collected and cache duration for attribute groups (continued)

Attribute group	When data is collected (cycle time or on demand)	Cache duration in seconds
MS SQL Analysis Services Storage Engine Query	On demand	No caching
MS SQL Analysis Services MDX Processing	On demand	No caching
MS SQL Analysis Services Rows Processing	On demand	No caching
MS SQL Running Queries Details	On demand	30
MS SQL Individual Queries Details	On demand	30
MS SQL Blocked Queries Details	On demand	10
MS SQL FileTable Detail	On demand	120
MS SQL Additional Database Details	On demand	60
MS SQL Availability Groups Details	On demand	60
MS SQL Availability Replicas Details	On demand	60
MS SQL Availability Replicas Status	On demand	60
MS SQL Availability Replicas Status Summary	On demand	60
MS SQL Availability Groups Summary	On demand	60
MS SQL Availability Replicas Statistics	On demand	60
MS SQL Availability Database Details	On demand	60
MS SQL Availability Database Statistics	On demand	60
MS SQL Availability Database Summary	On demand	60
MS SQL Availability Group Listener Details	On demand	60
MS SQL Availability Groups Details In Cluster	On demand	60
MS SQL Availability Replicas Details In Cluster	On demand	60
MS SQL Availability Database Details In Cluster	On demand	60
MS SQL Additional Availability Details	On demand	60
MS SQL Server Properties	On demand	10

The Microsoft SQL Server agent maintains long-running processes for the monitoring agent that communicates with the Tivoli Enterprise Monitoring Server and the collector that drives data collection. Short-running processes are used to access system data, access database instance log file data, issue and process the output of Microsoft SQL Server stored procedures, and perform other database instance interactions.

Attribute groups that gather data from Microsoft SQL Server SELECT statements and stored procedures

Some of the attribute groups gather data from the data that is returned by Microsoft SQL Server SELECT statements, stored procedures, or both. Table 2 on page 9 provides information about how each attribute group collects data (SELECT statements or stored procedures). The table also provides the following information:

- For attribute groups that use SELECT statements to gather data, the number of individual SELECT statements issued for the group and the tables that are accessed are listed.
- For attribute groups that use stored procedures to gather data, the names of the stored procedures are listed.

Table 2. Key aspects of data gathering for attribute groups

Attribute group	Primary collection mechanism	Number of SELECT statements issued	Tables accessed	Name of stored procedures
MSSQL Batch Stats	Perfmon	0	None	None
MSSQL Server Summary	SELECT and Perfmon	5	sysprocesses, sysjobhistory	serverproperty ('processid'), DBCC
MSSQL Server Detail	SELECT and Perfmon	5	sysprocesses, sysjobhistory	serverproperty ('processid'), DBCC
MSSQL Database Summary	SELECT and Perfmon	4	sysdatabases, sysfiles	sp_replcounters
MSSQL Database Detail	SELECT and Perfmon	4	sysdatabases, sysfiles	sp_replcounters
MSSQL Database Mirroring	SELECT and Perfmon	2	sys.database_mirroring	None
MSSQL Device Detail	SELECT and Perfmon	4	sysfiles, sysdevices, sysusages, spt_values	None
MSSQL Enterprise Log Shipping DB Details	SELECT	1	log_shipping_monitor_primary, log_shipping_monitor_history_detail, log_shipping_monitor_secondary	None
MSSQL Lock Detail	SELECT	2	syslockinfo, sysobjects	None
MSSQL Lock Conflict Detail	SELECT	1	sysprocesses	None
MSSQL Log Shipping DB Details	SELECT	1	log_shipping_monitor_primary, log_shipping_monitor_history_detail, log_shipping_monitor_secondary	None
MSSQL Log Shipping Errors	SELECT	1	log_shipping_monitor_error_detail	None
MSSQL Log Shipping Jobs Detail	SELECT	1	log_shipping_monitor_history_detail, sysjobhistory	None
MSSQL Text	SELECT	1	sysprocesses	None
MSSQL Process Summary	SELECT	2	sysprocesses	None
MSSQL Process Detail	SELECT	3	sysprocesses	None
MSSQL Statistics Summary	SELECT	1	@@cpu_busy, @@io_busy, @@idle, @@connections, @@pack_sent, @@pack_received, @@packet_errors, @@total_read, @@total_write, @@total_errors,	None

Table 2. Key aspects of data gathering for attribute groups (continued)

Attribute group	Primary collection mechanism	Number of SELECT statements issued	Tables accessed	Name of stored procedures
MSSQL Statistics Detail	SELECT	1	@@cpu_busy, @@io_busy, @@idle, @@connections, @@pack_sent, @@pack_received, @@packet_errors, @@total_read, @@total_write, @@total_errors,	None
MSSQL Remote Servers	SELECT	1	sys.servers	None
MSSQL Resource Pool Stats	Perfmon	0	None	None
MSSQL Configuration	SELECT	1	sys.configurations, sys.scurconfigs, sys.configures	None
MSSQL Server Enterprise View	SELECT	1	sys.processes, sys.jobhistory, sys.lockinfo	None
MSSQL Lock Summary	SELECT and Perfmon	16	sys.lockinfo	None
MSSQL Memory Manager	Perfmon	0	None	None
MSSQL Job Summary	SELECT	3	sys.categories, sys.jobs, sys.jobhistory	
MSSQL Job Detail	SELECT	3	sys.categories, sys.jobs, sys.jobhistory	
MSSQL Filegroup Detail	SELECT and Perfmon	2	sys.databases, sys.files	None
MSSQL Lock Resource Type Summary	Perfmon	0	None	None
MSSQL Workload Group Stats	Perfmon	0	None	None
MS SQL Server Transaction Summary	Perfmon	0	None	None
MS SQL Service Broker Activation	Perfmon	0	None	None
MS SQL Service Broker Statistics	Perfmon	0	None	None
MS SQL Service Broker Transport	Perfmon	0	None	None
MS SQL Table Summary	SELECT	3	sys.tables, sys.indexes, sys.dm_db_index_physical_stats, sys.dm_db_partition_stats	None
MS SQL Integration Service Details	Perfmon	0	None	None
MS SQL Analysis Services Data Mining	Perfmon	0	None	None

Table 2. Key aspects of data gathering for attribute groups (continued)

Attribute group	Primary collection mechanism	Number of SELECT statements issued	Tables accessed	Name of stored procedures
MS SQL Analysis Services Memory Statistics	Perfmon	0	None	None
MS SQL Analysis Services Storage Engine Query	Perfmon	0	None	None
MS SQL Analysis Services MDX Processing	Perfmon	0	None	None
MS SQL Analysis Services Rows Processing	Perfmon	0	None	None
MS SQL Running Queries Details	SELECT	1	sys.dm_exec_requests, sys.dm_exec_sessions, sys.dm_exec_sql_text	None
MS SQL Individual Queries Details	SELECT	1	sys.dm_exec_connections, sys.dm_exec_requests, sys.dm_exec_sessions, sys.dm_exec_sql_text	None
MS SQL Blocked Queries Details	SELECT	3	sys.dm_os_sys_info, sys.configurations, sys.dm_exec_query_stats, sys.dm_exec_requests, sys.dm_exec_sql_text, sys.dm_exec_plan_attributes	None
MS SQL FileTable Detail	SELECT	3	sys.databases, sys.filetables, sysindexes, sys.dm_filestream_non_transacted_handles	None
MS SQL Additional Database Details	SELECT	1	sys.databases, sys.database_filestream_options	None
MS SQL Availability Groups Details	SELECT	3	sys.availability_groups, sys.dm_hadr_availability_group_states, sys.dm_hadr_database_replica_states, sys.dm_hadr_availability_replica_states, sys.availability_replicas	None
MS SQL Availability Replicas Details	SELECT	2	sys.availability_replicas, sys.availability_groups, sys.dm_hadr_availability_group_states	None

Table 2. Key aspects of data gathering for attribute groups (continued)

Attribute group	Primary collection mechanism	Number of SELECT statements issued	Tables accessed	Name of stored procedures
MS SQL Availability Replicas Status	SELECT	2	sys.dm_hadr_availability_replica_states, sys.availability_groups, sys.dm_hadr_database_replica_states	None
MS SQL Availability Replicas Status Summary	SELECT	2	sys.dm_hadr_availability_replica_states, sys.availability_groups, sys.dm_hadr_database_replica_states	None
MS SQL Availability Groups Summary	SELECT	3	sys.availability_groups, sys.dm_hadr_availability_group_states, sys.dm_hadr_database_replica_states, sys.dm_hadr_availability_replica_states, sys.availability_replicas	None
MS SQL Availability Replicas Statistics	Perfmon	0	None	None
MS SQL Availability Database Details	SELECT	1	dm_hadr_database_replica_states, sys.availability_groups	None
MS SQL Availability Database Statistics	Perfmon	0	None	None
MS SQL Availability Database Summary	SELECT	1	dm_hadr_database_replica_states, sys.availability_groups	None
MS SQL Availability Group Listener Details	SELECT	1	sys.availability_group_listener_ip_addresses, sys.availability_group_listenersavg_list, sys.availability_groups	None
MS SQL Availability Groups Details In Cluster	SELECT	1	sys.availability_groups_cluster	None
MS SQL Availability Replicas Details In Cluster	SELECT	1	sys.dm_hadr_availability_replica_cluster_states, sys.availability_groups_cluster, sys.dm_hadr_availability_replica_cluster_nodes,	None
MS SQL Availability Database Details In Cluster	SELECT	1	sys.dm_hadr_database_replica_cluster_states, sys.availability_databases_cluster, sys.availability_groups_cluster	None

Table 2. Key aspects of data gathering for attribute groups (continued)

Attribute group	Primary collection mechanism	Number of SELECT statements issued	Tables accessed	Name of stored procedures
MS SQL Additional Availability Database Details	SELECT	1	sys.dm_hadr_database_replica_states, sys.dm_hadr_availability_replica_states	None
MS SQL Server Properties	SELECT	1	SERVERPROPERTY	None

Attribute groups that gather data from operating system or MS SQL Server files

Table 3 contains information about the attribute groups that gather data from the operating system or Microsoft SQL Server files.

Table 3. Data that is accessed by attribute groups that gather data from operating system or Microsoft SQL Server files

Attribute group	Data accessed
MS SQL Error Event Details	MS SQL server error log file
MS SQL Problem Summary	MS SQL server error log file
MS SQL Problem Detail	MS SQL server error log file
MS SQL Services Detail	Windows Services API
MS SQL Integration Details	Windows Services API
MS SQL Analysis service Data Mining	Windows Services API
MS SQL Analysis Service Memory Statistics	Windows Services API
MS SQL Analysis Service Storage Engine Query	Windows Services API
MS SQL Analysis Service MDX Processing	Windows Services API
MS SQL Analysis Service Rows Processing	Windows Services API

Chapter 2. Agent installation and configuration

Agent installation and configuration requires the use of the *IBM Tivoli Monitoring Installation and Setup Guide* and agent-specific installation and configuration information.

To install and configure Microsoft SQL Server agent, use the procedures for installing monitoring agents in the *IBM Tivoli Monitoring Installation and Setup Guide* along with the agent-specific installation and configuration information.

If you are installing silently by using a response file, see "Performing a silent installation of IBM Tivoli Monitoring" in the *IBM Tivoli Monitoring Installation and Setup Guide*.

With the self-describing agent capability, new or updated IBM Tivoli Monitoring agents using IBM Tivoli Monitoring V6.2.3 or later can become operational after installation without having to perform additional product support installation steps. To take advantage of this capability, see "Enabling self-describing agent capability at the hub monitoring server" in the *IBM Tivoli Monitoring Installation and Setup Guide*. Also, see "Self-describing monitoring agents" in the *IBM Tivoli Monitoring Administrator's Guide*.

Installing language packs

You can install language packs for your agent support files on the Tivoli Enterprise Monitoring agent and the Tivoli Enterprise Portal Server. There are separate procedures for installing language packs on Windows and on UNIX or Linux systems.

Before you begin

Ensure that you have installed the product in English before you install a language pack for the agent support files on the Tivoli Enterprise Monitoring Server, and the Tivoli Enterprise Portal Server.

Windows systems

This section contains the procedure for installing a language pack on Windows operating systems.

Procedure

1. Double-click `lpinstaller.bat` in the language pack CD to start the installation program.
2. Select the language of the installer and click **OK**.
3. Click **Next** on the Introduction panel.
4. Click **Add/Update** and click **Next**.
5. Select the folder in which the National Language Support package (NLSPackage) files are located.

Note: Usually the NLSPackage files are located in the `nlspackage` folder where the executable installer is located.

6. Select the language support for the agent of your choice and click **Next**.

Note: You can select multiple languages by pressing the Ctrl key.

7. Select the languages that you want to install and click **Next**.
8. Examine the installation summary page and click **Next** to start the installation.
9. Click **Finish** after the installation completes.
10. Restart the Tivoli Enterprise Portal, Tivoli Enterprise Portal Server, and Eclipse Help Server if any of these components are installed.

UNIX or Linux systems

This section contains the procedure for installing a language pack on UNIX or Linux operating systems.

Procedure

1. Run the following command to create a temporary directory on the computer. Ensure that the full path of the directory does not contain any spaces:

```
mkdir dir_name
```

2. Mount the language pack CD to the temporary directory that you have created.
3. Run the following command to start the installation program:

```
cd dir_name  
lpinstall.sh ITM Home Directory
```

where *ITM Home Directory* is the location where you have installed IBM Tivoli Monitoring. Usually it is `/opt/IBM/ITM` for AIX® and Linux systems.

4. Select the language of the installer and click **OK**.
5. Click **Next** on the Introduction panel.
6. Click **Add/Update** and click **Next**.
7. Select the folder in which the National Language Support package (NLSPackage) files are located.

Note: Usually, the NLSPackage files are located in the `nlspackage` folder where the installer executable is located.

8. Select the language support for the agent of your choice and click **Next**.

Note: You can select multiple languages by pressing the Ctrl key.

9. Select the languages that you want to install and click **Next**.
10. Examine the installation summary page and click **Next** to start the installation.
11. Click **Finish** after the installation completes.
12. Restart the Tivoli Enterprise Portal, Tivoli Enterprise Portal Server, and Eclipse Help Server if any of these components are installed.

Prerequisites checking

The prerequisite checker utility verifies whether all the prerequisites that are required for the agent installation are met. The prerequisite checker creates a log file that contains a report of all the prerequisites checks when the prerequisite checker was run.

For the Microsoft SQL Server agent, the prerequisite checker verifies the following requirements:

- Memory
- Disk
- Operating systems

Additionally, the prerequisite checker verifies whether the user, who installs the agent, is a member of the Administrators group.

For detailed information about installation prerequisites, see the Prerequisites topic for the agent in the IBM Tivoli Composite Application Manager for Microsoft Applications Information Center (http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/topic/com.ibm.itcamms.doc_6.3/welcome_msapps63.html).

You can run the prerequisite checker in stand-alone mode or remotely. For more information about the prerequisite checker, see "Prerequisite Checking for IBM Tivoli Monitoring Agents" in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Running as a non-administrator user

This section provides information about running the Microsoft SQL Server agent as a non-administrator user.

About this task

The Microsoft SQL Server agent can run as non-administrator. However, some functionality might not be available. This process supports Domain Users only.

Procedure

1. Start Windows application Active Directory Users and Computers and create a domain user.
 - Make sure that the new user is a member of the *Domain Users* group.
 - Make sure that the SQL Server that you monitor is a member of *Domain Computers*.
2. Add the newly created domain user in the *SQL Server Login* user group. The domain user must have sysadmin SQL Server role permissions on the SQL Server.
3. Log on to the SQL Server with a domain administrator account.
4. Use **File Systems**, to give the Modify permission to every drive that the Microsoft SQL Server agent accesses and propagate permissions to all subdirectories. Complete the following steps to propagate permissions:
 - a. Go to **My Computer**.
 - b. Right-click the *drive*.
 - c. Click the **Security** tab.
 - d. Add the newly created user.
 - e. Give modify permissions to the newly created user.
 - f. Click **OK**. This procedure takes a few minutes to apply permissions throughout the file system.
5. By using the Windows Registry, grant read access to HKEY_LOCAL_MACHINE, and propagate settings. Complete the following steps to propagate settings:
 - a. Right-click the HKEY_LOCAL_MACHINE directory and select **Permissions**.
 - b. Add the newly created user.
 - c. Select the newly created user.
 - d. Select the **Allow Read** check box.
 - e. Click **OK**. This procedure takes a few minutes to propagate through the entire HKEY_LOCAL_MACHINE tree.
6. By using the Windows Registry, grant the agent-specific registry permissions according to the following list:
 - If you installed a 32-bit agent on a 32-bit operating system, grant full access to the HKEY_LOCAL_MACHINE\SOFTWARE\Candle directory, and then propagate the settings.
 - If you installed a 32-bit agent on a 64-bit operating system, grant full access to the HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Candle directory, and then propagate the settings.
 - If you installed a 64-bit agent on a 64-bit operating system, grant full access to the HKEY_LOCAL_MACHINE\SOFTWARE\Candle directory, and then propagate the settings.

Complete the following steps to propagate settings:

- a. Right-click the directory for which you have granted full access and select **Permissions**.
- b. Add the newly created user.
- c. Select the newly created user.
- d. Select the **Allow Full Control** check box.
- e. Click **OK**. This procedure takes a few minutes to propagate through the entire HKEY_LOCAL_MACHINE\SOFTWARE\Candle tree.

7. Add a new Domain User to the *Performance Monitor Users* group.
8. Verify that Domain Users are members of the *Users* group.
9. Grant the following permissions to the Windows directory while running as a non-administrator user:
 - If a 32-bit agent is installed on a 32-bit operating system, grant read and write access to the `OS_installation_drive:\Windows\system32` directory
 - If a 32-bit agent is installed on a 64-bit operating system, grant read and write access to the `OS_installation_drive:\Windows\SysWOW64` directory
 - If a 64-bit agent is installed on a 32-bit or a 64-bit operating system, grant read and write access to the `OS_installation_drive:\Windows\system32` directory

Note: You do not need to grant the permissions to the Windows directory if you are using Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012.

10. Grant the following permissions to the SQL Server data file and log file:
 - The default data file path is `SQLServer_root_dir\DATA`, where `SQLServer_root_dir` is the root directory of the SQL Server instance. For example, if the root directory of the SQL Server instance is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL`, the data file path is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA`.
 - The default log file path is `SQLServer_root_dir\LOG`, where `SQLServer_root_dir` is the root directory of the SQL Server instance. For example, if the root directory of the SQL Server instance is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL`, the log file path is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG`.
11. Grant full permissions to the `Candle_Home` directory. The default path is `C:\IBM\ITM`.
12. Restart the SQL Server to ensure that local security permissions are applied.
13. Change the logon settings for the SQL Server agent services to the non-administrator user by completing the following steps:
 - a. Click **Start > Administrative Tools > Services**.
 - b. Right-click the **Monitoring Agent For SQL Server *instance_name***, and click **Properties**. The SQL Service Properties window opens.
 - c. Click the **Log On** tab.
 - d. Click **This account** and type the user name.
 - e. In the **Password** and **Confirm Password** fields, enter the password, and click **OK**.
 - f. Repeat steps b - e for the **Monitoring Agent For SQL Server Collector *instance_name***, where *instance_name* is the Microsoft SQL Server instance name.

Agent-specific installation and configuration

When performing the steps to install and configure the Microsoft SQL Server agent as described in the *IBM Tivoli Monitoring Installation and Setup Guide*, "Installing monitoring agents," use the agent-specific configuration information provided in this chapter.

You can install and configure the Microsoft SQL Server agent locally or remotely by using a GUI or command line.

Agent-specific information is provided for the following procedures:

- "Granting permissions" on page 19
- "Local configuration" on page 23
 - "Using the Manage Tivoli Enterprise Monitoring Services window" on page 23
 - "Silent installation" on page 32
- "Local environment variables" on page 32

- “Remote configuration” on page 34
- “Reconfiguration” on page 35
- “Configuration settings” on page 25
- “Starting and stopping the Microsoft SQL Server agent” on page 35
- “Setting up the Microsoft SQL Server agent in a cluster environment” on page 36

Never attempt to start the monitoring agent until you have completed the configuration steps appropriate to the installation of the monitoring agent.

Granting permissions

If the Microsoft SQL Server agent is configured to Windows authentication, the monitoring agent will access the Microsoft SQL Server using the run-as user ID, otherwise the monitoring agent will use the configured Microsoft SQL Server user ID. Further reference in this section to Microsoft SQL Server user ID or SQL Server ID refers to either the run-as user ID or the Microsoft SQL Server user ID.

The Microsoft SQL Server user ID must have access to the Microsoft SQL Server agent. The procedure described in this section includes creating a Microsoft SQL Server user ID and granting permission to the new user ID, which is the minimum authority required.

The SQL Server ID used to configure this monitoring agent must have the following SQL Server authorities:

- Required authorization roles
 - Database roles: Public access is required for each database that is being monitored.
 - Server roles: No special Server Role is required.
- Optional authorization roles

Each Take Action command has a separate set of authorization roles that are required for the SQL Server credentials to pass to the Take Action command. These authorities are listed for each Take Action command in Chapter 6, “Take Action commands reference,” on page 239.

The monitored attributes in Table 4 on page 20 require additional authorization for the SQL Server ID used to configure the Microsoft SQL Server agent. If you want to monitor any of these attributes, configure the SQL Server ID used by the monitoring agent with the corresponding authorization.

Table 4. Additional authorization

Attribute group	Navigation item, workspace, view	Situation	Authorization
MS SQL Job Detail	Navigator item: Jobs	None predefined	<p>Server Role: Server AdministratorsThis authorization enables data collection for all SQL Server jobs.</p> <p>—OR—</p> <p>Database Role (msdb database): publicAn SQL Server Agent proxy account must be defined. The agent can collect job data only on the jobs that this SQL Server ID owns.</p> <p>—OR—</p> <p>Database Role (msdb database): SQLAgentUserRole (SQL Server 2005 only)The agent can collect job data only on the jobs that this SQL Server ID owns.</p> <p>—OR—</p> <p>Database Role (msdb database): SQLAgentReaderRole (SQL Server 2005 only)The agent can collect job data on all SQL Server jobs.</p> <p>—OR—</p> <p>Database Role (msdb database): SQLAgentOperatorRole (SQL Server 2005 and 2008)The agent can collect job data on all SQL Server jobs.</p>
	Workspace: Job Detail		
	View: Job Status, Job Detail		
MS SQL Job Summary	Navigator item: Jobs		
	Workspace: Job Summary		
	View: Running Job Current Interval, Job Summary		
MS SQL Lock Detail	Navigator item: Server Locking	None predefined	VIEW SERVER STATE permission is required on the SQL Server
MS SQL Lock Summary	Workspace: Server Locking, Lock Detail		
	View: Log Detail, Lock Detail		

MS SQL Server Detail	Navigator item: Servers	<ul style="list-style-type: none"> • MS_SQL_Proc_Buffs_Active_Warn • MS_SQL_Proc_Buffs_Active_Crit • MS_SQL_Proc_Buffs_Used_Warn • MS_SQL_Proc_Buffs_Used_Crit • MS_SQL_Proc_Cache_Active_Warn • MS_SQL_Proc_Cache_Active_Crit • MS_SQL_Proc_Cache_Used_Warn • MS_SQL_Proc_Cache_Used_Crit 	Database Role: db_owner (on default database associated with the SQL Server ID) —OR— Server Role: System Administrator
	Workspace: Server Detail		
	View: Server Detail-Status, Server Detail-Statistics		
MS SQL Server Summary	Navigator item: Servers	None predefined	VIEW SERVER STATE permission is required on the SQL Server
	Workspace: Server Summary		
	View: CPU Utilization, Server Summary		
MS SQL Statistics Summary	Navigator item: Servers	<ul style="list-style-type: none"> • MS_SQL_Client_Cnt_Pct_Used_Crit • MS_SQL_Client_Cnt_Pct_Used_Warn • MS_SQL_IOERR_Startup_Warning • MS_SQL_IOError_Curintvl_Warning • MS_SQL_LogonPct_Crit • MS_SQL_LogonPct_Warning • MS_SQL_Network_Read_Rate_Crit • MS_SQL_Network_Read_Rate_Warn • MS_SQL_Network_Write_Rate_Crit • MS_SQL_Network_Write_Rate_Warn • MS_SQL_Pct_IO_Warning 	VIEW SERVER STATE permission is required on the SQL Server
	Workspace: Server Statistics		
	View: Statistics Summary Chart view, IO Summary, Statistic Summary table view		

MS SQL Table Detail	Navigator item: Databases	<ul style="list-style-type: none"> • MS_SQL_Fragmentation_Warn • MS_SQL_Fragmentation_Crit • MS_SQL_Opt_Stats_Age_Warn • MS_SQL_Opt_Stats_Age_Crit 	Database Role: db_owner (for each database)
	Workspace: Table Detail		—OR—
	View: Table Detail, Fragmentation (greater than 30%) by Index Name		Server Role: System Administrator —OR— VIEW DATABASE STATE and VIEW SERVER STATE permissions are required on the SQL Server

Granting permissions for SQL Server 2005, 2008, and 2012

Before installing the Microsoft SQL Server agent, you must grant permissions for the Microsoft SQL Server.

Before you begin

Ensure that you have the Database administrator authorization role to grant permissions. If you are configuring the Microsoft SQL Server agent locally, complete this procedure between the steps for installing and configuring the monitoring agent. If you are configuring the Microsoft SQL Server agent remotely, complete this procedure after installing and configuring the monitoring agent.

Perform the appropriate installation procedures in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Procedure

1. Click **Start > Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**.
2. In the **Tree** tab, select **Logins** in the Security folder (**Console Root > WindowsName > Security > Logins**).
3. Right-click **Logins** and select **New Login**.
4. Select the **General Folder** option and type the SQL Server user ID if you are using SQL Server authentication. Type the Windows user ID if you are using Windows authentication.
5. Select one of the following authentication methods:
 - SQL Server
 - Windows
6. Type a password in the **Password** field if you are using the SQL Server authentication method.
7. Select the **User Mapping** option.
8. In the **Specify which databases can be accessed by this login** area, select the check box for each database that you currently have to give permission to each selected database.
9. Select the **Server Roles** option.
10. If you are running a Take Action command, you might need additional authority. See the Take Action command descriptions in "Predefined Take Action commands" on page 239.
11. Click **OK** to display the Confirm Password window.
12. Retype the password that you typed for the user ID.
13. Click **OK** to display the new user ID in the **Logins** list.

What to do next

Configure the Microsoft SQL Server agent to start the monitoring agent and begin monitoring your Microsoft SQL Server application.

Local configuration

If you are installing and configuring locally, use the steps in the *IBM Tivoli Monitoring Installation and Setup Guide*, "Installing monitoring agents." Also, use the agent-specific configuration information in this section and in Table 5 on page 26 for the Manage Tivoli Enterprise Monitoring Services window.

Using the Manage Tivoli Enterprise Monitoring Services window

Use the Configure Database Agents window to select the database that you want to monitor.




Figure 1. The Configure Database Agents window

In the Configure Database Agents window, select the server that you want to monitor from the **Database Servers Available** list, and move the server to the **Server to Monitor** list. The Database Server Properties window opens.




Figure 2. The Database Server Properties window

The following fields are populated in the Database Server Properties window:

- Server Name

- Database Version
- Home Directory
- Error Log File

The following fields in the Database Server Properties window are optional:

- Windows Authentication
- Support Long Lived Database Connections
- Extended Parns
- Include
- Day(s) Frequency
- Weekly Frequency
- Monthly Frequency
- Collection Start Time

If you do not select the **Windows Authentication** field, you must enter your login and password in the **Login** and **Password** fields by using only ASCII characters. You can use the **Extended Parns** field to enable or disable the data collection for the Table Detail attribute group. To disable data collection for the Table Detail attribute group, type koqtbld in the **Extended Parns** field. Click **OK** to see the Configure Database Agents window. Click **OK** in the Configure Database Agents window to complete the configuration process.

You can specify a daily, weekly, or monthly frequency for the collection of the MS SQL Table Detail attribute group.

You can select the **Table Detail Continuous Collection** checkbox to enable continuous collection of the MS SQL Table Detail attribute group. If you select the **Table Detail Continuous Collection** check box, you must enter a value in the **Interval Between Two Continuous Collection (in min.)** field.

Configuration settings: Table 5 on page 26 contains a list of the configuration settings for each of the interfaces where you can specify these settings and a description of each setting.

Table 5. Names and descriptions of configuration settings for each interface

Interfaces where configuration settings are specified				
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line	Description	Examples
Server Name	Database Server Instance Name¹	INSTANCE= <i>InstanceName</i>	<p>Name of the SQL Server instance that is to be monitored.</p> <p>Use the host name if the SQL Server being monitored is the default instance; otherwise use the instance name.</p> <p>The name must be short enough to fit within the total managed system name, which must be between 2 and 32 characters in length.</p>	<p>If the SQL Server instance being monitored is the default SQL Server instance and the host name is popcorn, enter popcorn in this field.</p> <p>If the SQL Server instance being monitored is a named instance with the instance name of mysqlserver and the host name is popcorn, enter mysqlserver in this field.</p>
Login	Database Server User Id¹	DBSETTINGS.db_login= <i>UserId</i>	<p>SQL Server user ID to be used to connect to the SQL Server.</p> <p>See “Granting permissions” on page 19 for more information.</p> <p>Login is required only when Windows Authentication parameter is set to False.</p> <p>Use only ASCII characters.</p>	
Password	Password¹	DBSETTINGS.db_password= <i>Password</i>	<p>Password for the SQL Server user ID</p> <p>Password is required only when Windows Authentication parameter is set to False.</p> <p>Use only ASCII characters.</p>	

Table 5. Names and descriptions of configuration settings for each interface (continued)

Interfaces where configuration settings are specified			Description	Examples
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line		
Home Directory	Database Server Home Directory Path ¹	DBSETTINGS.db _home= <i>HomeDirPath</i>	Install directory Server instance monitored.	The default home directory path for the default SQL Server 2005 instance is C:\Program Files\Microsoft SQL Server\MSSQL. A named SQL Server 2005 instance has a default home directory path in the format C:\Program Files\Microsoft SQL Server\MSSQL\$ <i>instance_name</i> , where <i>instance_name</i> is the SQL Server instance name.
Error Log File	Database Server Error Log File ¹	DBSETTINGS.db _errorlog= <i>ErrorlogPath</i>	Fully qualified location and name of the SQL Server Error Log	The default error log path for the default SQL Server 2005 instance is C:\Program Files\Microsoft SQL Server\MSSQL\LOG\ERRORLOG. A named SQL Server 2005 instance has a default error log path in the format C:\Program Files\Microsoft SQL Server\MSSQL\$ <i>instance_name</i> \LOG\ERRORLOG, where <i>instance_name</i> is the SQL Server instance name.
Right-click the agent instance, and select Change Startup... to set the ID under which the agent instance is to run	<p>Use this account² If selecting this option, complete the following fields as described:</p> <ul style="list-style-type: none"> • Account: ID under which the agent instance is to run • Password and Confirm password: password 	<p>_WIN32_STARTUP _.<i>Username</i> and _WIN32_STARTUP _.<i>Password</i>—or— _WIN32_STARTUP _.<i>LocalSystem</i></p>	<p>The ID and passwords must have Windows Administrator authority for the system on which the monitoring agent is to run.</p> <p>See “Granting permissions” on page 19 for information about the required user ID permissions.</p> <p>The LocalSystem account is the default ID used by the monitoring agent.</p>	<p>The following tacmd command property, explicitly sets the user ID for the monitoring agent to the LocalSystem account: _WIN32_STARTUP_.<i>LocalSystem</i>=1.</p> <p>When setting the _WIN32_STARTUP_.<i>LocalSystem</i> property, the _WIN32_STARTUP_.<i>InteractWithDesktop</i> property must also be set.</p>

Table 5. Names and descriptions of configuration settings for each interface (continued)

Interfaces where configuration settings are specified				
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line	Description	Examples
Windows Authentication	Windows Authentication	DBSETTINGS.db _winauth= 1/0	Enables or disables Windows Authentication	DBSETTINGS.db_winauth=1 —OR— DBSETTINGS.db_winauth=0
Support Long Lived Database Connections	Support Long Lived Database Connections	DBSETTINGS.db _11dbconn= 1/0	Enables or disables long lived database connections The following Take Action commands do not use Long Lived Database Connections: <ul style="list-style-type: none"> • Database Check-up and Tune-up • Dump Database • Dump Transaction Log • Rebuild Fragmented Indexes • Refresh Query Optimizer Statistics • Update Space Usage Information The following attribute groups do not use Long Lived Database Connections: <ul style="list-style-type: none"> • MS SQL Text • MS SQL Filegroup Detail • MS SQL Server Summary 	DBSETTINGS.db_11dbconn=1 —OR— DBSETTINGS.db_11dbconn=0
Extended Parms	Extended Parms	DBSETTINGS.db _extparms="KOQTBLD"	Disables capture of Table Detail attributes	

Table 5. Names and descriptions of configuration settings for each interface (continued)

Interfaces where configuration settings are specified				
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line	Description	Examples
Include	Database names to include	DBSETTINGS.db_include	<p>Use this feature to filter databases that you want to monitor.</p> <p>To specify database filter, you must first select a separator. A separator is a character that distinguishes a database name or database expression from the other database name or database expression.</p> <p>While selecting a separator, ensure that database names and database expression do not contain the character that you choose as a separator. You must not use the wildcard characters that are typically used in the T-SQL query (for example, %, _ [], ^, -) if they are used in the database names or database expression.(Continued on the next page.)</p>	<p>Examples of filters:</p> <p>Case 1: % usage</p> <p>Example: @@%m%</p> <p>Output: All the databases that have the character m in their names are filtered.</p> <p>Case 2: _ usage</p> <p>Example: @@_</p> <p>Output: All the databases that are of length four characters are filtered.</p> <p>Case 3: [] usage</p> <p>Example: @@[m]_</p> <p>Output: All the databases of length four characters and whose names start with the character m are filtered.(Continued on the next page.)</p>

Table 5. Names and descriptions of configuration settings for each interface (continued)

Interfaces where configuration settings are specified				
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line	Description	Examples
Include (Continued)	Database names to include	DBSETTINGS.db_include	<p>When specifying database filter:</p> <ul style="list-style-type: none"> Database names must start with a separator. Database expression must start with two separators. <p>Note: Database expression is a valid expression that can be used in the LIKE part of the T-SQL query. However, you cannot use the T-SQL ESCAPE clause while specifying the database expression. The following workspaces are affected by database filter:</p> <ul style="list-style-type: none"> Database Information Enterprise Database Summary Device Detail Table Detail Filegroup Detail 	<p>Case 4: [^] usage</p> <p>Example: @@[^m]%</p> <p>Output: All the databases (of any length) except those whose names start with the character m are filtered.</p> <p>Case 5: Wrong input</p> <p>Example: @%m%</p> <p>Output: None of the databases are filtered.</p> <p>Case 6: Default</p> <p>Example: Field left blank (No query is typed)</p> <p>Output: All the databases are filtered.</p> <p>Case 7: Mixed patterns</p> <p>Example: @@[m-t]_d%</p> <p>Output: All the databases (of any length) whose names start with the characters m, n, o, p, q, r, s, t, followed by any character, with the character d in the third place are filtered.</p>
Day(s) Frequency	Table Detail Collection Day(s) Frequency ³	DBSETTINGS.db_tbl_d_daily	Use this feature to define the frequency of collecting data of Table Detail attributes. The values can be from zero to 31.	DBSETTINGS.db_tbl_d_daily=1
Weekly Frequency	Table Detail Collection Weekly Frequency ³	DBSETTINGS.db_tbl_d_weekly	Use this feature to specify a particular day for collecting data for Table Detail attributes. The values can be from zero to seven.	DBSETTINGS.db_tbl_d_weekly=1

Table 5. Names and descriptions of configuration settings for each interface (continued)

Interfaces where configuration settings are specified			Description	Examples
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line		
Monthly Frequency	Table Detail Collection Monthly Frequency ³	DBSETTINGS.db_tbl_d_monthly	Use this feature to define the data collection of Table Detail attributes on a particular day of the month. The possible values are 1, 2, 3, and so on.	DBSETTINGS.db_tbl_d_monthly=5
Collection Start Time	Collection Start Time Hour	DBSETTINGS.db_tbl_d_coll_stime_hour	The possible values are zero to 23. The default value is zero.	DBSETTINGS.db_tbl_d_coll_stime_hour=1
	Collection Start Time Minute	DBSETTINGS.db_tbl_d_coll_stime_minute	The possible values are from zero to 59. The default value is zero.	DBSETTINGS.db_tbl_d_coll_stime_minute=59
Table Detail Continuous Collection	Enable/Disable Continuous Collection for Table Detail	DBSETTINGS.db_tbl_dconcoll1	Use this feature for the continuous background collection of Table Detail data.	DBSETTINGS.db_tbl_dconcoll=1
Interval Between Two Continuous Collection (in min.)	Interval time between two collection in minutes	DBSETTINGS.db_tbl_dcoll_time	Specify the time for the interval between two collections in minutes. The minimum interval time is 3 minutes. You can select the Interval Between Two Continuous Collection (in min.) check box or you can use Scheduling to specify continuous collection of the Table Detail attribute group. If you use the Between Two Continuous Collection (in min.) check box, you must specify the time interval for collection. If you use Scheduling to specify the collection of the Table Detail attribute group, the minimum time interval is 1 day.	DBSETTINGS.db_tbl_dcoll_time=10

Table 5. Names and descriptions of configuration settings for each interface (continued)

Interfaces where configuration settings are specified			Description	Examples
Manage Tivoli Enterprise Monitoring Services window	Tivoli Enterprise Portal	tacmd command line		
<ul style="list-style-type: none"> • ¹ Database Server Properties tab • ² Agent tab Run as area • ³ The agent collects the data at the time interval for which data collection occurs frequently. For example, if you specify all three Table Details Collection configuration settings (Day, Weekly, and Monthly), the agent starts the data collection according to the following conditions: <ul style="list-style-type: none"> – If Day(s) Frequency ≤ 7, the Day(s) Frequency settings are selected, and the Weekly and Monthly Frequency settings are ignored. – If Day(s) Frequency > 7, the Weekly Frequency settings are selected, and the Day(s) and Monthly Frequency settings are ignored. 				

Configuring local environment variables

You can configure local environment variables to change the behavior of the Microsoft SQL Server agent.

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, from the **Actions** menu, click **Advanced > Edit Variables**.
2. In the Tivoli Enterprise Monitoring Server: Override Local Variable Settings window, click **Add**.
3. In the Add Environment Setting Override window, enter the variable and the corresponding value.

Note: See “Local environment variables” to view the list of environment variables that you can configure.

Silent installation

If you are performing a silent installation using a response file, see the *IBM Tivoli Monitoring Installation and Setup Guide*, “Performing a silent installation of IBM Tivoli Monitoring.”

Local environment variables

You can change the behavior of the Microsoft SQL Server agent by configuring local environment variables.

Variables for checking SQL Server service availability

To check the availability of the SQL Server service, add the following environment variables:

- **COLL_MSSQL_RETRY_INTERVAL:** This variable provides the retry interval (in minutes) to check the SQL Server service status. If value is less than or equal to zero, then the variable takes the default value of 1 minute.
- **COLL_MSSQL_RETRY_CNT:** This variable provides the number of retries that the SQL Server agent makes to check whether SQL Server service is started or not. If the SQL Server service is not started after the number of retries then collector stops working. If value is less than or equal to zero, then the variable takes the default value of infinity.

Variables for monitoring the SQL Server error log file

To monitor the MS SQL Error Event Details attribute group, you can add the following environment variables:

- **COLL_ERRORLOG_STARTUP_MAX_TIME:** This variable provides the time interval (T) for error collection before agent startup. The default value is 0 minutes. This variable can take the following values:

- T = 0** The agent starts monitoring the error log file from the time the agent starts or is restarted. The agent does not read the errors that were logged in the error log file before the agent was started.
- T = 1** The agent monitors the error log file according to the following values that are set for the **COLL_ERRORLOG_STARTUP_MAX_EVENT_ROW** variable, which is represented by R.
 - If $R < 0$, the agent starts monitoring the error log file from the time the agent starts or is restarted.
 - If $R = 1$, the agent monitors all the errors that are logged in the error log file.
 - If $R > 1$ and the agent is installed for the first time, the agent monitors the error log file until R errors are monitored. If $R > 1$ and the agent is restarted, the agent monitors all the previously missed R errors.
- T > 1** The agent monitors all previous errors that were logged up to T minutes from the time the agent starts or is restarted. The agent monitoring also depends on the following values that you set for the **COLL_ERRORLOG_STARTUP_MAX_EVENT_ROW** variable:
 - If $R \leq 0$, the agent starts monitoring the error log file from the time the agent starts or is restarted.
 - If $R = 1$, the agent monitors the error log file for all the errors that are logged up to T minutes.
 - If $R > 1$, the agent monitors not more than R errors that are logged in last T minutes.
- **COLL_ERRORLOG_STARTUP_MAX_EVENT_ROW**: This variable provides the maximum number of errors that must be processed at agent startup. The default value is 0. This variable can take the following values:
 - R = 0** The agent starts monitoring the error log file from the time the agent starts or restarts. The agent does not read errors that were created in the error log file before the agent was started.
 - R = 1** The agent monitors the errors that were logged in the last T minutes from the time the agent starts or restarts.
 - R > 1** The agent monitors R errors logged in the last T minutes.
- **COLL_ERRORLOG_MAX_EVENT_ROW**: This variable provides the number of error rows to be displayed on the Tivoli Enterprise Portal. The default value is 50. This variable can take the following values:
 - X = 0** The agent does not display error logs on the Tivoli Enterprise Portal.
 - X > 0** The agent displays X error rows on the Tivoli Enterprise Portal.

Variable for query timeout interval

To set the query timeout interval for the SQL Server agent, use the **QUERY_TIMEOUT** variable. This environment variable defines the maximum amount of time (in seconds) that the SQL Server agent waits to receive a response for a query that is sent to the SQL Server. The value for this variable must be less than 45 seconds. However, if you set the value for this variable as 0 seconds, the SQL Server agent waits indefinitely to receive a response from the SQL Server. If the SQL Server agent accesses a large number of locked databases, the value for this variable must preferably be in the range of 10 - 20 seconds. If the query is not processed within the set timeout interval, the SQL Server agent skips the timed out query and moves to the next query in the queue. The agent does not display data on the Tivoli Enterprise Portal for the query that has timed out.

Variable for viewing information about enabled jobs

To view information about enabled jobs, use the **COLL_JOB_DISABLED** environment variable. If you set the value of this variable as 1, the Microsoft SQL Server agent does not display information about disabled jobs. If you do not specify this variable, you can view information about enabled and disabled jobs.

Variable for limiting rows in MS SQL Filegroup Detail attribute group

To limit the number of rows that the collector service fetches for the MS SQL Filegroup Detail attribute group, use the **COLL_KOQFGRP_MAX_ROW** environment variable. This environment variable defines the maximum number of rows that the collector service fetches for the Filegroup Detail attribute group.

You can use this environment variable to modify the default limit of maximum rows in the `koqcoll.ct1` file. Complete the following steps to modify the default limit:

1. Specify the maximum number of rows for `KOQFGRPD` in the `koqcoll.ct1` file.
2. Add the `COLL_KOQFGRPD_MAX_ROW` environment variable, and ensure that the value of this variable is the same as the value that you have specified in the `koqcoll.ct1` file.

If the value in the `koqcoll.ct1` file is less than the value that is specified in the `COLL_KOQFGRPD_MAX_ROW` environment variable, the value in the `koqcoll.ct1` file is treated as the value for the maximum number of rows.

If the value in the `koqcoll.ct1` file is greater than the value that is specified in the `COLL_KOQFGRPD_MAX_ROW` environment variable, the value in the `COLL_KOQFGRPD_MAX_ROW` environment variable is treated as the value for the maximum number of rows.

Variable for limiting rows in MS SQL Device Detail attribute group

To limit the number of rows that the collector service fetches for the MS SQL Device Detail attribute group, use the `COLL_KOQDEVD_MAX_ROW` environment variable. This environment variable defines the maximum number of rows that the collector service fetches for the Device Detail attribute group. You can use this environment variable to modify the default limit of maximum rows in the `koqcoll.ct1` file. Complete the following steps to modify the default limit:

1. Specify the maximum number of rows for `KOQDEVD` in the `koqcoll.ct1` file.
2. Add the `COLL_KOQDEVD_MAX_ROW` environment variable, and ensure that the value of this variable is the same as the value that you have specified in the `koqcoll.ct1` file.

If the value in the `koqcoll.ct1` file is less than the value that is specified in the `COLL_KOQDEVD_MAX_ROW` environment variable, the value in the `koqcoll.ct1` file is treated as the value for the maximum number of rows.

If the value in the `koqcoll.ct1` file is greater than the value that is specified in the `COLL_KOQDEVD_MAX_ROW` environment variable, the value in the `COLL_KOQDEVD_MAX_ROW` environment variable is treated as the value for the maximum number of rows.

Variables for enhancing collection for MS SQL Database Detail attribute group

To enhance the MS SQL Database Detail attribute group collection, you can use the following environment variables:

- **COLL_KOQDBD_INTERVAL:** This environment variable enables you to specify a time interval (in minutes) between two consecutive collections of the MS SQL Database Detail attribute group.
- **COLL_DBF_RENAME_RETRY_CNT:** This environment variable enables you to specify the number of attempts that can be made to move the `%COLL_HOME%_tmp_%COLL_VERSION%_%COLL_SERVERID%_%COLL_SERVERID%_DBD_TEMP` file to the `%COLL_HOME%_tmp_%COLL_VERSION%_%COLL_SERVERID%_%COLL_SERVERID%_DBD_PREV` file.

Remote configuration

If you are installing and configuring remotely, use the steps in the *IBM Tivoli Monitoring Installation and Setup Guide*, "Deploying non-OS agents." Also, use the agent-specific configuration information in this section and in Table 5 on page 26 for the following interfaces:

- Tivoli Enterprise Portal
- `tacmd` command line

Using the Tivoli Enterprise Portal

To deploy this monitoring agent remotely using the command line, use the procedure, "Deploying through the portal," in the *IBM Tivoli Monitoring Installation and Setup Guide*.

In the New Managed System Configuration window, use the settings in Table 5 on page 26 for the **Tivoli Enterprise Portal Database Server Properties** tab and the **Agent** tab **Run as** information.

Using the tacmd command line

To deploy this monitoring agent remotely using the command line, use the procedure, “Deploying through the command line,” in the *IBM Tivoli Monitoring Installation and Setup Guide*. Also, use the agent-specific configuration information in Table 5 on page 26 for the tacmd addSystem command. The *IBM Tivoli Monitoring Command Reference* has complete information about the **tacmd addSystem** command.

Use the **-t** or **--type TYPE** parameter to specify the Microsoft SQL Server agent that you are configuring: OQ.

Specify the properties with the **-p** or **-property** option.

For example:

```
tacmd addSystem -t OQ -n Primary:myhostname:NT -p
DBSETTINGS.db_login=sa
DBSETTINGS.db_password=sapwd
DBSETTINGS.db_home="C:\Program Files\Microsoft SQL Server\MSSQL"
DBSETTINGS.db_errorlog="C:\Program Files\Microsoft SQL Server\MSSQL\LOG\ERRORLOG"
DBSETTINGS.db_winauth=1
DBSETTINGS.db_1ldbconn=0
DBSETTINGS.db_include=#master##temp%
INSTANCE=MyServer
DBSETTINGS.db_winauth=1
DBSETTINGS.db_tbl_d_weekly=1
DBSETTINGS.db_tbl_d_coll_stime_hour=23
DBSETTINGS.db_tbl_d_coll_stime_minute=59
```

The parameters in the example are shown on separate lines for clarity. When typing the command, type all of the parameters on one line.

Reconfiguration

If you need to reconfigure the Microsoft SQL Server agent, ensure that the steps for installing the monitoring agent in the *IBM Tivoli Monitoring Installation and Setup Guide* are completed.

Local reconfiguration

If you are reconfiguring an instance locally, use the Manage Tivoli Enterprise Monitoring Services window. See Table 5 on page 26 for the configuration settings.

Remote reconfiguration

If you are reconfiguring remotely, use one of the following interfaces:

- tacmd command line

Use the **configureSystem** command. See the *IBM Tivoli Monitoring Command Reference Guide* for complete information about this command.

When reconfiguring, enter the information for the property that you are changing as well as the INSTANCE property.

- Tivoli Enterprise Portal

See Table 5 on page 26 for the configuration settings.

Starting and stopping the Microsoft SQL Server agent

Table 6 on page 36 shows which interfaces you can use on Windows operating systems locally or remotely to start the monitoring agent.

Table 6. Interfaces for starting and stopping the Microsoft SQL Server agent locally and remotely

Local	Remote
<ul style="list-style-type: none"> • Manage Tivoli Enterprise Monitoring Services • tacmd startAgent • tacmd stopAgent • tacmd restartAgent 	<ul style="list-style-type: none"> • Tivoli Enterprise Portal • tacmd startAgent • tacmd stopAgent • tacmd restartAgent

Manage Tivoli Enterprise Monitoring Services

Enter only ASCII characters in the fields for the Manage Tivoli Enterprise Monitoring Services window.

Tivoli Enterprise Portal

See the "Working with monitoring agents," "Starting and stopping a monitoring agent" in the *IBM Tivoli Monitoring Administrator's Guide* for information about using the Tivoli Enterprise Portal to start or stop the monitoring agent.

tacmd command line

In the following examples, the tacmd command is used to start, stop, or restart the Microsoft SQL Server agent instances that are configured for the Microsoft SQL Server on Windows operating systems, where *myhostname* is the host name of the computer where the monitoring agent is running.

- Local
 - tacmd startAgent -t oq
 - tacmd stopAgent -t oq
 - tacmd restartAgent -t oq
- Remote
 - tacmd stopagent -t oq -m InstanceName:hostname:MSS
 - tacmd startagent -t oq -m InstanceName:hostname:MSS
 - tacmd restartagent -t oq -m InstanceName:hostname:MSS

The command for restarting the monitoring agent remotely can be tacmd restartagent -t oq -n *INST1:ESX1-PG3-2:MSS*, where *INST1* is the Microsoft SQL Server agent instance name, and *ESX1-PG3-2* is the host name of the computer where the monitoring agent is running.

For information about using the tacmd commands, see the *IBM Tivoli Monitoring Command Reference*.

Setting up the Microsoft SQL Server agent in a cluster environment

To use this monitoring agent in a Microsoft Cluster Server environment requires special configuration.

The *IBM Tivoli Monitoring Installation and Setup Guide* contains an overview of clustering. The information provided here is specifically for installing and setting up the Microsoft SQL Server agent in a Microsoft Cluster Server environment.

The SQL Server and the Microsoft SQL Server agent both support multiple instances running on the same node. One agent does not interfere with another while running on the same node.

When the SQL Resource Group is moved from node to node, the server down situation event fires. This event is caused by the Microsoft SQL Server agent coming online faster than the SQL Server. When the SQL Server comes online the event clears.

Do not use the Start and Stop Take Action commands for the Microsoft SQL Server agent because they conflict with the actions taken by the cluster server.

Requirements

You can set up a cluster environment for the following versions of the Microsoft SQL Server:

- Microsoft SQL Server 2005
- Microsoft SQL Server 2008
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2012

In addition to installing and setting up the Microsoft SQL Server agent, the following additional steps are required for the cluster environment:

- Setting **CTIRA_HOSTNAME** to a common value for all monitoring agents (usually the cluster name)
- Setting **CTIRA_HIST_DIR** to a common disk location if history is stored at the monitoring agents (if history for the Microsoft SQL Server agent is configured to be stored at the monitoring agent)
- Creating a monitoring agent cluster resource in the Resource Group of the Virtual Server 3

On Windows systems, IBM Tivoli Monitoring requires that monitoring agents are installed in the same directory path as the OS agent. Therefore, each node in a cluster must have installed all monitoring agents (on the nodes system disk) that are required to support the cluster applications that can run on that cluster node.

Installing and configuring the monitoring agent

This section contains the procedure for installing the Microsoft SQL Server agent on each node of a cluster where it is possible the Microsoft Virtual Servers to run.

Procedure

1. Setting **CTIRA_HOSTNAME**

- Because there can be multiple instances of the Microsoft SQL Server agent, each instance must be configured with a **CTIRA_HOSTNAME**. Set the **CTIRA_HOSTNAME** environment variable to the name of the Microsoft Cluster Server cluster for all monitoring agents running in that cluster. By setting the **CTIRA_HOSTNAME** for all agents in the cluster to the same name, you can navigate to all of the monitoring agents for that cluster in the Tivoli Enterprise Portal.
- When deciding on the value for **CTIRA_HOSTNAME**, consider that the managed system name is comprised of three parts: **CTIRA_SUBSYSTEM_ID**, **CTIRA_HOSTNAME**, and **CTIRA_NODETYPE**. Also, the name is limited to 31 characters. By default for the Microsoft SQL Server Agent, **CTIRA_NODETYPE** is set to MSS, and **CTIRA_SUBSYSTEM_ID** is set to the *Microsoft SQL Virtual Server* name. The **CTIRA_SUBSYSTEM_ID** is used to distinguish the multiple instances of the Microsoft SQL Server agent.

2. Setting **CTIRA_HIST_DIR**

- If history for the Microsoft SQL Server agent is configured to be stored at the monitoring agent, each instance of the monitoring agent must be configured with common **CTIRA_HIST_DIR** that points to a shared disk directory.
- If history is stored at the Tivoli Enterprise Management Server, setting **CTIRA_HIST_DIR** is not required. Storing history at the Tivoli Enterprise Management Server puts a higher burden on that server.

3. Setting **COLL_HOME**

If you want to collect data for attribute groups that use configuration files at a shared location, and these files are used by multiple cluster nodes, set the **COLL_HOME** variable to *X:\shared-location*, where *X* is a shared drive that is accessible to the cluster nodes. For example, you can set the value for the **COLL_HOME** variable when you are defining the configuration settings for the MS SQL Table Detail attribute group or MS SQL Error Event Details attribute group.

4. Creating a monitoring agent cluster resource

- Set each Microsoft SQL Server agent startup parameter to manual so the cluster resource can control the starting and stopping of the monitoring agent.

- After these parameters are set for each Microsoft SQL Server agent, instance cluster resources to control the monitoring agents must be created. Each Microsoft SQL Server agent is comprised of two Windows Services: KOQAGENT x and KOQCOLL x , where x is the agent instance number.

Example

Examples of Windows Services names:

- Monitoring Agent for Microsoft SQL Server – SQLTEST
- Monitoring Agent for Microsoft SQL Server – Collector SQLTEST
- Monitoring Agent for Microsoft SQL Server – SQLTEST2
- Monitoring Agent for Microsoft SQL Server – Collector SQLTEST2

Creating a resource

This section contains the procedure for creating a resource for the Windows Server 2003 operating system.

Procedure

1. Click **Start > Administrative Tools > Cluster Administrator**.
2. Select the group for the **Instance that is being worked > SQLTEST**.
3. Right-click the group, and then click **New > Resource**.
4. In the New Resource window, complete the fields as follows:
 - Name:** KOQAGENT0
 - Description:** *resource_description*
 - Resource type:** Generic Service
 - Group:** SQLTEST
5. In the Possible Owners window, accept the default of **all Available Nodes**.
6. In the Dependencies window, do not add any dependencies on history disk.
7. In the Generic Service Parameters window, complete the fields as follows:
 - Service name:** KOQAGENT0
 - Start Parameters:** -Hkey "KOQ\610\SQLVS1"
8. Click **Finish**.
9. Click **Advanced**.
10. Clear the **Affect the group** check box.
11. Right-click the group, and then click **New > Resource**.
12. Complete the following information:
 - Name:** KOQCOLL_*instance_name*
 - Resource Type:** Generic Service
 - Group:** SQLTEST
13. In the Possible Owners window, accept the default of **all Available Nodes**.
14. In the Dependencies window, do not add any dependencies on history disk.
15. In the Generic Service Parameters window, complete the fields as follows:
 - Name:** KOQCOLL*instance_name*
 - Start Parameters:** -Hkey "KOQ\610\SQLVS1"
16. Click **Finish**.
17. Click **Advanced**.
18. Clear the **Affect the group** check box.
19. Bring the two agent resources online.
20. Repeat these steps for the other instances of the monitoring agent in the cluster environment.

Creating a service

This section contains the procedure for creating a service for Windows 2008 systems.

Procedure

1. Click **Start > Administrative Tools > Failover Cluster Management**.
2. Expand **Failover Cluster Management**.
3. Expand **Services and Applications** and right-click the SQL instance that you want to configure.
4. Click **Add a resource > Generic Service**. The New Resource Wizard opens.
5. On the Select Service page, select **Monitoring Agent for Microsoft SQL Server**, and click **Next**.

Note: To add the Microsoft SQL Server- Collector service, select Monitoring Agent for **Microsoft SQL Server- Collector service** on the Select Service page, and click **Next**.

6. On the Confirmation page, check the details, and click **Next**.
7. On the Summary page, click **Finish**.
8. Right-click the service that you have added, and click **Properties**.
9. In the **Dependencies** tab, select **SQL Server** or **Monitoring Agent for Microsoft SQL Sever** from the **Resources** list.
10. Click **Apply > OK**.
11. Right-click the service that you created, and click **Bring this resource online**.

Making configuration changes

After control of the agent cluster resource is given to the cluster server, you must take the agent cluster resource offline to make configuration changes or edit the agent variables on the node from which the agent cluster resource runs. If the agent cluster resource is not offline when the agent config utility attempts to take the agent offline, the cluster server notices that the monitoring agent went offline and attempts to bring the monitoring agent back online. When finished with the configuration changes for the monitoring agent, bring the agent cluster resource back online.

Chapter 3. Workspaces reference

A workspace is the working area of the Tivoli Enterprise Portal application window. The Navigator tree contains a list of the workspaces provided by the agent.

About workspaces

Use the Navigator tree to select the workspace you want to see. As part of the application window, the status bar shows the Tivoli Enterprise Portal Server name and port number to which the displayed information applies and the ID of the current user.

When you select an item in the Navigator tree, a default workspace is displayed. When you right-click a Navigator item, a menu that includes a Workspace item is displayed. The Workspace item contains a list of workspaces for that Navigator item. Each workspace has at least one view. Some views have links to other workspaces. You can also use the Workspace Gallery tool as described in the *Tivoli Enterprise Portal User's Guide* to open workspaces.

The workspaces in the Navigator are displayed in a Physical view that shows your enterprise as a physical mapping or a dynamically populated logical view that is agent-specific. You can also create a Logical view. The Physical view is the default view.

This monitoring agent provides predefined workspaces. You cannot modify or delete the predefined workspaces, but you can create new workspaces by editing them and saving the changes with a different name.

When multiple instances of the monitoring agent are defined on a system, the top-level node becomes Microsoft Active Directory. The Microsoft Active Directory workspace is undefined at this node. A subnode for each instance is created called *Instance:Hostname:OQ*. A workspace that is called *Instance:Hostname:OQ* is associated with the instance node. This workspace is comparable to the Microsoft Active Directory workspace.

A table view within a workspace corresponds to a group of attributes; the columns in the table view show some or all the attributes available in the attribute group.

Additional information about workspaces

For more information about creating, customizing, and working with workspaces, see "Using workspaces" in the *Tivoli Enterprise Portal User's Guide*.

For a list of the predefined workspaces for this monitoring agent and a description of each workspace, see "Predefined workspaces" on page 42 and the information about each individual workspace.

Some attribute groups for this monitoring agent might not be represented in the predefined workspaces or views for this agent. For a full list of the attribute groups, see "Attributes in each attribute group" on page 59.

Predefined workspaces

When you install the Microsoft SQL Server agent, you receive a collection of predefined workspaces. The predefined workspaces display data collected on an interval basis.

You can use the predefined workspaces in the Microsoft SQL Server agent for the following purposes:

- Analyze the performance of your system.
- Identify problems and bottlenecks.
- Obtain the information you need to tune your system.
- Evaluate how your tuning decisions affect the performance of your system.

In most cases, a predefined workspace has an attribute group that contains attributes similar to the columns in the predefined workspace. You can use these values to build your own workspaces and to create your own situations.

The Microsoft SQL Server agent provides predefined workspaces, which are organized by Navigator item.

- Microsoft SQL Server Navigator item
This Navigator item is available only when more than one Microsoft SQL Server agent instance is configured and running on the same system.
- *SQL Server instance name* Navigator item
This Navigator item is available only when more than one Microsoft SQL Server agent instance is configured and running on the same system.
 - MS SQL Server Overview workspace
 - Enterprise Servers Overview workspace
- Microsoft SQL Server - *SQL Server instance name* Navigator item
This Navigator item is available only when exactly one Microsoft SQL Server agent instance is configured and running on the system.
 - MS SQL Server Overview workspace
 - Enterprise Servers Overview workspace
- AlwaysOn Navigator item
 - Availability Database Information workspace
 - Availability Groups Information workspace
 - Availability Groups Information In Cluster workspace
 - Availability Replicas and Database Information workspace
 - Availability Replicas Information workspace
 - Availability Replicas Status workspace
 - Additional Availability Database Details workspace
- Databases Navigator item
 - Additional Database Details workspace
 - Databases Information workspace
 - Log and Space Information workspace
 - Enterprise Database Summary workspace
 - FileTable Detail workspace
 - Device Detail workspace
 - Table Detail workspace
 - Filegroup Detail workspace
 - Database Mirror Detail workspace
 - Enterprise Database Mirror Detail workspace

- Log Shipping Error Details
- Log Shipping Job Details
- Log Shipping Database Details
- Enterprise Log Shipping Database Details
- Errorlog Alerts Navigator item
 - Enterprise Errorlog Alerts workspace
 - Errorlog Alerts workspace
- Jobs Navigator item
 - Job Summary workspace
 - Job Detail workspace
- Server Locking Navigator item
 - Enterprise Server Locking workspace
 - Server Locking workspace
 - Process Holding Lock SQL Text workspace
 - SQL Text for Process ID workspace
 - Lock Resource Type Summary workspace
- Processes Navigator item
 - Enterprise Processes Summary workspace
 - Processes Information workspace
 - SQL Text for Process ID workspace
- Servers Navigator item
 - Enterprise Server Statistics workspace
 - Enterprise Server Summary workspace
 - Server Statistics workspace
 - Server Summary workspace
 - Server Configuration workspace
 - Server Detail workspace
 - Services Detail workspace
 - Server Batch Statistics workspace
- Services Navigator item
 - Services Details workspace
 - SQL Server Integration Service Details workspace
 - Analysis Services Data Mining Statistics workspace
 - Analysis Services Memory Statistics workspace
 - Analysis Services Storage Engine Query Processing workspace
 - Analysis Services Rows Processing workspace

The Microsoft SQL Server agent has following composite workspaces:

- SQL Server CPU Bottleneck workspace
- SQL Server Disk Space and Logon Statistics
- SQL Server Logical Disk Bottleneck workspace
- SQL Server Memory Bottleneck workspace
- SQL Server Physical Disk Bottleneck workspace

Additional Availability Database Details workspace

Additional Availability Database Details is a predefined workspace that contains the Additional Availability Database Details table view. This view displays additional information about the availability databases that are available on an SQL Server instance.

Additional Database Details workspace

Additional Database Details workspace is a predefined workspace that displays information about the databases that are available in the SQL Server version 2012 or later. The predefined workspace contains the Additional Database Details table view that displays information about the databases that are available in the SQL Server version 2012 or later.

Analysis Services Data Mining Statistics workspace

Analysis Services Data Mining Statistics is a predefined workspace that displays statistical information about the data mining feature of the SQL Server Analysis Services. The predefined workspace contains the following views:

- Total Data Mining Processing bar chart that displays statistics about the total predictions, total queries, and total rows.
- Data Mining Processing bar chart that displays statistics about the number of predictions per second, number of queries per second, and number of rows per second.
- Data Mining Statistics table view that displays statistics about the data mining queries.

Analysis Services Memory Statistics workspace

Analysis Services Memory Statistics is a predefined workspace that displays statistical information about the memory that is used by the SQL Server Analysis Services. The predefined workspace contains the following views:

- Memory Limits Statistics bar chart that displays the memory that is currently used by the SQL Server Analysis Services. This view also displays the upper and lower limits of the amount of memory that can be used by the SQL Server Analysis Services.
- Background Cleaner Memory Statistics view that displays the amount of background cleaner memory that is currently used, and the amount of shrinkable and non-shrinkable memory that is known to the background cleaner.
- Analysis Services Memory Statistics table view that displays statistics about the memory that is used by the SQL Server Analysis Services.

Analysis Services Storage Engine Query Processing workspace

Analysis Services Storage Engine Query Processing is a predefined workspace that displays statistical information about the dimension queries and the storage engine queries that are processed by the SQL Server Analysis Services. The predefined workspace contains the following views:

- Dimension Queries Statistics bar chart that displays the following information:
 - The number of dimension queries that are currently processed
 - The total number of dimension queries that have been processed since the server was started
- Analysis Services Storage Engine Query Processing table view that displays the statistical information about storage engine query processing.
- Analysis Services MDX Processing table view that displays the metrics of MDX processing.

Analysis Services Rows Processing workspace

Analysis Services Rows Processing is a predefined workspace that displays the summary of the rows that are processed by the SQL Server Analysis Services. The predefined workspace contains the following views:

- Rows Processing Per Second bar chart that displays the following information:

- The rate at which rows are converted per second
- The rate at which rows are read per second
- The rate at which rows are written per second
- Total Rows Processing bar chart that displays the following information:
 - The total number of rows that are converted per second
 - The total number of rows that are read per second
 - The total number of rows that are written per second

Availability Database Information workspace

Availability Database Information is a predefined workspace that provides information about the availability databases that are available on an SQL Server instance. The predefined workspace contains the following views:

- Availability Database Summary table view that displays a summary of all the databases in the availability groups that are available on an SQL Server instance.
- Availability Database Details table view that displays the details of all the availability databases that are available on an SQL Server instance.
- Availability Database Statistics table view that displays the detailed information about the statistics of all the availability databases that are available on an SQL Server instance.

This workspace provides a cross-link to the Additional Availability Database Details workspace.

Availability Groups Information workspace

Availability Groups Information is a predefined workspace that provides information about the availability groups that are available on an SQL Server instance. This workspace is supported for SQL Server 2012, or later. The predefined workspace contains the following views:

- Availability Groups Details table view that displays information about each availability group that is available on an SQL Server instance.
- Availability Groups Summary table view that displays a summary of the availability groups that are available on an SQL Server instance.
- Availability Groups Listener table view that displays information about the availability group listener of all the availability groups that are hosted by the local SQL Server instance.

The Availability Groups Details view provides a cross-link to the Availability Replicas and Database Information workspace.

Availability Groups Information In Cluster workspace

Availability Groups Information In Cluster is a predefined workspace that provides information about the availability groups that are available in a Windows Service Failover Clustering (WSFC) cluster. This workspace is supported for SQL Server 2012 Enterprise Edition, or later. The predefined workspace contains the following views:

- Availability Groups Details In Cluster table view that displays information about all the availability groups that are available in a WSFC cluster.
- Availability Replicas Details In Cluster table view that displays information about all the availability replicas on all the nodes of a WSFC cluster.
- Availability Database Details In Cluster table view that displays information about all the availability databases that are available in a WSFC cluster.

The Availability Groups Details view provides a cross-link to the Availability Replicas and Database Information workspace.

Availability Replicas Information workspace

Availability Replicas Information is a predefined workspace that provides information about the availability replicas that are available on an SQL Server instance. The predefined workspace contains the following views:

- Availability Replicas Details table view that displays the details of all the availability replicas that are available in all the availability groups in an SQL Server instance..
- Availability Replicas Statistics table view that displays the statistics of all the availability replicas that are available in all the availability groups in an SQL Server instance, except the local replica.

Availability Replicas and Database Information workspace

Availability Replicas and Database Information is a predefined workspace that provides information about the availability replicas and databases that are available on an SQL Server instance. The predefined workspace contains the following views:

- Availability Replicas Details table view that displays the details of all the availability replicas that are available on a local server instance.
- Availability Replicas Status table view that displays the status of all the availability replicas that are available on a local server instance.
- Availability Database Details table view that displays the details of all the availability databases that are available on an SQL Server instance.

Availability Replicas Status workspace

Availability Replicas Status is a predefined workspace that provides a summary of the status of the availability replicas that are available on an SQL Server instance. The predefined workspace contains the following views:

- Availability Replicas Status Summary table view that displays a summary of the status of all the availability replicas that are available on an SQL Server instance.
- Availability Replicas Status table view that displays the details of all the availability replicas that are available on an SQL Server instance.

Databases Information workspace

Databases Information is a predefined workspace that displays information about the databases for the server. The predefined workspace contains the following views:

- Databases Summary table view that displays summary information for the databases (such as the error status and the options being used in your environment)
- Databases Summary bar chart that displays data from selected columns from the Databases Summary table view (such as the minimum percentages of freespace in the database and the error log)
- Databases Detail table view that displays specific information for a database (such as the name of the database and the owner)
- Database Detail bar chart that displays data from selected columns in the Database Detail table view (such as the percentages of freespace in the database and log)

This workspace provides a cross-link to the Availability Replicas Information workspace.

Database Mirror Detail workspace

Database Mirror Detail is a predefined workspace that provides details about the database mirrors. The predefined workspace contains the following views:

- Database Mirrors view that displays the database mirrors that are created on this instance
- Database Mirror Detail view that displays statistics related to each database mirror

Device Detail workspace

Device Detail is a predefined workspace that provides details about the SQL Server devices. The predefined workspace contains the following views:

- Device Free Space Percent view that displays a bar chart of the percent free space on a device
- Device Detail view that displays statistics related to each device

Enterprise Database Mirror Detail workspace

Enterprise Database Mirror Detail is a predefined workspace that provides details about the database mirrors for all the servers. The predefined workspace contains the following views:

- Transaction Delay view that shows a bar chart of the transaction delays of mirrors
- Enterprise Database Mirrors view that shows information about all database mirrors

Enterprise Database Summary workspace

Enterprise Database Summary is a predefined workspace that displays an overview of SQL server database activity. The predefined workspace contains the following views:

- Databases Summary table view that displays summary information for the databases (such as the error status and the options being used in your environment)
- Percent Data Freespace bar chart showing the percentage of freespace in the database from the column in the Database Summary table view

Enterprise Errorlog Alerts workspace

Enterprise Errorlog Alerts is a predefined workspace that displays an overview of information about the error log for each SQL server in your environment. The predefined workspace contains the following views:

- Alert Summary table view with overview information (such as the size of the error log)
- Alert Summary bar chart showing the total number of errors from the column in the Alert Summary table view

Enterprise Log Shipping Database Details workspace

Enterprise Log Shipping Database Details is a predefined workspace that provides detailed information about Log Shipping Databases for the enterprise. The predefined workspace contains the following views:

- Enterprise Threshold Exceeded bar chart that displays information about the number of times the threshold exceeded
- Enterprise Log Shipping Database Details table that displays the database details

Enterprise Processes Summary workspace

Enterprise Process Summary is a predefined workspace that displays overview information about the processes for the SQL server. The predefined workspace contains the following views:

- Process Summary table view that displays information about the processes (such as the number of processes and the statuses)
- CPU Percent that displays information from selected columns in the Server Summary table view
- Processes Blocked bar chart that displays the number of processes blocked by locks

Enterprise Server Locking workspace

Enterprise Server Locking is a predefined workspace that displays overview information about locks for all the SQL servers in your environment. The predefined workspace contains the following views:

- Server Locking table view with overview information (such as the total number of locks)

- Server Locking bar chart that displays information about selected columns in the Server Locking table view (such as the total number of locks and the total number of processes being blocked by those locks)

Enterprise Servers Overview workspace

Enterprise Servers Overview is a predefined workspace that displays an overview of the health of every Microsoft SQL server. The predefined workspace contains the following views:

- CPU Percent bar chart view that displays the Microsoft SQL server CPU percentage usage for every monitored Microsoft SQL server
- Lock Analysis bar chart view that displays the total number of locks and the total number of processes blocked (waiting for a lock) for every monitored Microsoft SQL server
- Statistics Summary bar chart view that displays the number of physical reads and writes for every monitored Microsoft SQL server

Enterprise Server Statistics workspace

Enterprise Server Statistics is a predefined workspace that displays statistical information about the SQL servers in your environment. The predefined workspace contains the following views:

- Server Statistics table that displays the statistical information (such as the physical reads per second)
- Server Statistics bar chart that displays information from selected columns in the Server Statistics table view (such as the physical reads and writes per second)

Enterprise Server Summary workspace

Enterprise Server Summary is a predefined workspace that displays an overview of activity for your SQL servers. The predefined workspace contains the following views:

- Server Summary table view that displays summary information about the SQL servers (such as the name and status)
- Remote Servers table view that displays summary information about the remote SQL servers (such as the name and status)
- CPU Utilization bar chart that displays information from selected columns in the Server Summary table view

Errorlog Alerts workspace

Errorlog Alerts is a predefined workspace that displays information about the error log alerts. The predefined workspace contains the following views:

- Alerts Summary table view that displays summary information for the error logs (such as the size of the error log and the number of errors during the interval)
- Alert Summary bar chart that displays information about selected columns in the table view that displays summary information
- Alerts Detail table view that displays specific information for an error (such as the severity level of the error)

Filegroup Detail workspace

Filegroup Detail is a predefined workspace that provides details about the SQL Server database file groups. The predefined workspace contains the following views:

- Filegroup Sizes view that displays a bar chart of the sizes of the file groups
- Filegroup Detail view that displays statistics for each file group

FileTable Detail workspace

FileTable Detail workspace is a predefined workspace that displays information about the tables that are configured as FileTables in the SQL Server. The predefined workspace contains the FileTable Detail table view that displays information about each FileTable that is available in the SQL Server.

Job Detail workspace

Job Detail is a predefined workspace that provides details about the SQL Server jobs. The predefined workspace contains the following views:

- Job Status view that displays the status of all jobs
- Job Detail view that displays statistic information about all jobs

Job Summary workspace

Job Summary is a predefined workspace that provides summary details about the SQL Server jobs. The predefined workspace contains the following views:

- Running Jobs Current[®] Interval view that displays a bar chart of the number of jobs that are currently running
- Job Summary view that displays job statistics

Lock Resource Type Summary workspace

Lock Resource Type Summary is a predefined workspace that provides details about the SQL Server lock summary by resource type. The predefined workspace contains the following views:

- Lock Requests per Second by Resource Type view that displays a bar chart for lock requests per second
- Lock Resource Type Summary view that displays statistics of lock resource types

Log and Space Information workspace

Log and Space Information is a predefined workspace that displays information about the transaction log, space being used, and devices for the selected database. The predefined workspace contains the following views:

- Log and Space Information table view that displays information about the transaction log and space usage (such as the size of the log and the number of devices)
- Data and Log Size bar chart that displays data from selected columns in the Log and Space Information table view (such as the size a freespace for the database and log)
- Devices table view that displays information for the devices (such as the name, type, and size of the device)
- Device Size bar chart that displays data from selected columns in the Devices table view (such as the size and percentage of free space for the device)

Log Shipping Database Details workspace

Log Shipping Database Details is a predefined workspace that provides details about log shipping databases. The predefined workspace contains the following views:

- Threshold Exceeded bar chart that displays information about the number of times threshold exceeded
- Log Shipping Database Details table that displays information about log shipping databases

Log Shipping Error Details workspace

Log Shipping Error Details is a predefined workspace that provides details about log shipping job errors. The predefined workspace contains the following views:

- Job Completion Status bar chart that displays the number of error records
- Log Shipping Errors table that displays log shipping error details

Log Shipping Job Details workspace

Log Shipping Job Details is a predefined workspace that provides details about log shipping jobs. The predefined workspace contains the following views:

- Job Completion Status bar chart that displays information about the number of jobs that ended in the error state
- Log Shipping Job Details table that displays log shipping job details

MS SQL Server Historical Summarized Availability workspace

The MS SQL Server Historical Summarized Availability workspace shows the availability information by server averaged over the number of months that you specify in the Time Span dialog. This workspace consists of the following three views:

- Availability by Server (average over months), which shows the average percentage of time that the system has been available since startup, grouped by server instance
- Problem Summary Over Time by Server (average over months), which shows individual averages that include the following information, high and total, per server:
 - Error messages
 - Error messages with a severity of 17 or higher
 - Error messages with a severity level of less than 17 that have occurred
- Client Usage by Server (average over months), which shows the following information per server:
 - Average percentage of maximum user connections that are in use
 - Average percentage of client workstation connections in use

MS SQL Server Historical Summarized Availability Daily workspace

The MS SQL Server Historical Summarized Availability Daily workspace shows the availability information, client usage, problem summary, database summary, and process summary information by day for a managed server. This workspace consists of the following five views:

- Availability (daily), which shows the percentage of time that the system has been available since startup, summarized by day
- Client Usage (daily), which shows the following information, summarized by day:
 - Average percentage of maximum user connections in use
 - Average percentage of client workstation connections in use
- Problem Summary (daily), which shows the following information, high and total, summarized by day:
 - Error messages
 - Error messages with a severity of 17 or higher
 - Error messages with a severity level of less than 17 that have occurred
- Database Summary (daily), which shows information about databases on the managed server, summarized by day. The information includes the maximum and average statistics information for the following types of databases:
 - Databases in read-only status
 - Databases in DBO-only status
 - Databases that can be accessed by only one user at a time
 - Databases that have errors
- Process Summary (daily), which shows the following maximum and average statistics information, summarized by day:
 - Percentages of blocked processes
 - Percentages of stopped processes
 - Percentages of processes in lock sleep

MS SQL Server Historical Summarized Availability Hourly workspace

The MS SQL Server Historical Summarized Availability Hourly workspace shows the availability information, client usage, problem summary, database summary, and process summary information by hour for a managed server. This workspace consists of the following five views:

- Availability (hourly), which shows the percentage of time that the system has been available since startup, summarized by hour
- Client Usage (hourly), which shows the following information, summarized by hour:
 - Average percentage of maximum user connections in use
 - Average percentage of client workstation connections in use
- Problem Summary (hourly), which shows the following information, high and total, summarized by hour:
 - Error messages
 - Error messages with a severity of 17 or higher
 - Error messages with a severity level of less than 17 that have occurred
- Database Summary (hourly), which shows information about databases on the managed server, summarized by hour. The information includes the maximum and average statistics information for the following types of databases:
 - Databases in read-only status
 - Databases in DBO-only status
 - Databases that can be accessed by only one user at a time
 - Databases that have errors
- Process Summary (hourly), which shows the following maximum and average statistics information, summarized by hour:
 - Percentages of blocked processes
 - Percentages of stopped processes
 - Percentages of processes in lock sleep

MS SQL Server Historical Summarized Availability Weekly workspace

The MS SQL Server Historical Summarized Availability Weekly workspace shows the availability information, client usage, problem summary, database summary, and process summary information by week for a managed server. This workspace consists of the following five views:

- Availability (weekly), which shows the percentage of time that the system has been available since startup, summarized by week
- Client Usage (weekly), which shows the following information, summarized by week:
 - Average percentage of maximum user connections in use
 - Average percentage of client workstation connections
- Problem Summary (weekly), which shows the following information, high and total, summarized by week:
 - Error messages
 - Error messages with a severity of 17 or higher
 - Error messages with a severity level of less than 17 that have occurred
- Database Summary (weekly), which shows information about databases on the managed server, summarized by week. The information includes the maximum and average statistics information for the following types of databases:
 - Databases in read-only status
 - Databases in DBO-only status
 - Databases that can be accessed by only one user at a time

- Databases that have errors
- Process Summary (weekly), which shows the following maximum and average statistics information, summarized by week:
 - Percentages of blocked processes
 - Percentages of stopped processes
 - Percentages of processes in lock sleep

MS SQL Server Historical Summarized Capacity workspace

The MS SQL Server Historical Summarized Capacity workspace shows the percentage of system resources used for the time span that you specify in the Time Span dialog. This workspace consists of the following three views:

- Freespace by Server (average over months), which shows the average of the lowest percentage of free space in the database and the average of the lowest percentage of free space in the transaction log during the time span that you specify in the Time Span dialog
- CPU Usage by Server (average over months), which shows the percentage of CPU usage for the following during the time span that you specify in the Time Span dialog:
 - SQL server
 - I/O operations
 - All processes on the host
- Lock Capacity by Server (average over months), which shows the total number of locks that can still be taken out and the number of locks that are blocking other processes during the time span that you specify in the Time Span dialog

MS SQL Server Historical Summarized Capacity Daily workspace

The MS SQL Server Historical Summarized Capacity Daily workspace shows the percentage of system resources used, summarized by day. This workspace consists of the following four views:

- Freespace (daily), which shows the average of the lowest percentage of free space in the database and the average of the lowest percentage of free space in the transaction log, summarized by day
- Freespace by Database (daily), which shows the minimum percentage of free space in the database and the minimum percentage of free space in the transaction log for the database, summarized by day
- CPU Usage (daily), which shows the percentage of CPU usage for the following, summarized by day:
 - SQL server
 - I/O operations
 - All processes on the host
- Lock Capacity (daily), which shows the total number of locks that can still be taken out and the number of locks that are blocking other processes, summarized by day

MS SQL Server Historical Summarized Capacity Hourly workspace

The MS SQL Server Historical Summarized Capacity Hourly workspace shows the percentage of system resources used, summarized by hour. This workspace consists of the following four views:

- Freespace (hourly), which shows the average of the lowest percentage of free space in the database and the average of the lowest percentage of free space in the transaction log, summarized by hour
- Freespace by Database (hourly), which shows the minimum percentage of free space in the database and the minimum percentage of free space in the transaction log for the database, summarized by hour
- CPU Usage (hourly), which shows the percentage of CPU usage for the following, summarized by hour:
 - SQL server
 - I/O operations
 - All processes on the host

- Lock Capacity (hourly), which shows the total number of locks that can still be taken out and the number of locks that are blocking other processes, summarized by hour

MS SQL Server Historical Summarized Capacity Weekly workspace

The MS SQL Server Historical Summarized Capacity Weekly workspace shows the percentage of system resources used, summarized by week. This workspace consists of the following four views:

- Freespace (weekly), which shows the average of the lowest percentage of free space in the database and the average of the lowest percentage of free space in the transaction log, summarized by week
- Freespace by Database (weekly), which shows the minimum percentage of free space in the database and the minimum percentage of free space in the transaction log for the database, summarized by week
- CPU Usage (weekly), which shows the percentage of CPU usage for the following, summarized by week:
 - SQL server
 - I/O operations
 - All processes on the host
- Lock Capacity (weekly), which shows the total number of locks that can still be taken out and the number of locks that are blocking other processes, summarized by week

MS SQL Server Historical Summarized Performance workspace

The MS SQL Server Historical Summarized Performance workspace shows the average cache information and resources for the time span that you specify in the Time Span dialog. This workspace consists of the following three views:

- Cache Hit Ratio by Server (average over months), which shows the ratio of data cache hits to total requests, indicating the effectiveness of the data cache during the time span that you specify in the Time Span dialog
- Cache Buffers by Server (average over months), which shows the number of cache buffers in the free buffer pool and the average number of buffers that the LazyWriter scanned when searching the data cache for an unused buffer to replenish the free buffer pool during the specified time period
- Cache Size (KB) by server (average over months), which shows the number of kilobytes (KB) allocated for the data cache as well as the number of kilobytes (KB) allocated for the procedure cache during the specified time period

MS SQL Server Historical Summarized Performance Daily workspace

The MS SQL Server Historical Summarized Performance Daily workspace shows the cache information and resources, summarized by day. This workspace consists of the following three views:

- Cache Hit Ratio (daily), which shows the average, minimum, and maximum ratios of data cache hits to total requests, indicating the effectiveness of the data cache, summarized by day
- Cache Buffers by Server (daily), which shows the average, minimum, and maximum numbers of cache buffers in the free buffer pool and average, minimum, and maximum numbers of buffers that the LazyWriter scanned when searching the data cache for an unused buffer to replenish the free buffer pool, summarized by day
- Cache Size (KB) by server (daily), which shows the average, minimum, and maximum number of kilobytes (KB) allocated for the data cache as well as the average, minimum, and maximum number of kilobytes (KB) allocated for the procedure cache, summarized by day

MS SQL Server Historical Summarized Performance Hourly workspace

The MS SQL Server Historical Summarized Performance Hourly workspace shows the cache information and resources, summarized by hour. This workspace consists of the following three views:

- Cache Hit Ratio (hourly), which shows the average, minimum, and maximum ratios of data cache hits to total requests, indicating the effectiveness of the data cache, summarized by hour
- Cache Buffers by Server (hourly), which shows the average, minimum, and maximum numbers of cache buffers in the free buffer pool and average, minimum, and maximum numbers of buffers that the LazyWriter scanned when searching the data cache for an unused buffer to replenish the free buffer pool, summarized by hour
- Cache Size (KB) by server (hourly), which shows the average, minimum, and maximum number of kilobytes (KB) allocated for the data cache as well as the average, minimum, and maximum number of kilobytes (KB) allocated for the procedure cache, summarized by hour

MS SQL Server Historical Summarized Performance Weekly workspace

The MS SQL Server Historical Summarized Performance Weekly workspace shows the cache information and resources, summarized weekly. This workspace consists of the following three views:

- Cache Hit Ratio (weekly), which shows the average, minimum, and maximum ratios of data cache hits to total requests, indicating the effectiveness of the data cache, summarized by week
- Cache Buffers by Server (weekly), which shows the average, minimum, and maximum numbers of cache buffers in the free buffer pool and average, minimum, and maximum numbers of buffers that the LazyWriter scanned when searching the data cache for an unused buffer to replenish the free buffer pool, summarized by week
- Cache Size (KB) by server (weekly), which shows the average, minimum, and maximum number of kilobytes (KB) allocated for the data cache as well as the average, minimum, and maximum number of kilobytes (KB) allocated for the procedure cache, summarized by week

MS SQL Server Overview

MS SQL Server Overview is a predefined workspace that displays an overview of the health of the Microsoft SQL server. The predefined workspace contains the following views:

- CPU Percent bar chart view that displays the Microsoft SQL server CPU percentage usage
- Cache Utilization bar chart view that displays the data cache and procedure cache sizes
- Errorlog Alerts bar chart view that displays error message counts, including the total number of error messages and a count of the high severity error messages
- Percent Freespace bar chart view that displays the minimum database data free space percentage and minimum database log free space percentage for the Microsoft SQL server
- Processes Summary bar chart view that displays counts of various Microsoft SQL server processes by type, including the total number of Microsoft SQL server processes, blocked Microsoft SQL server processes, stopped Microsoft SQL server processes, and infected Microsoft SQL server processes

Process Holding Lock SQL Text workspace

Process Holding Lock SQL Text workspace is a predefined workspace that displays information about a process that has a lock on the database. Part of the information displayed is the SQL command that the process is executing. You access this workspace from the Server Locking workspace. The predefined workspace contains the following views:

- Event Console view
- Process Holding Lock SQL Text table view that displays detailed information about the process that has a lock on the database such as the Client Process ID, Process Status, and Total CPU Time (Sec.).

Processes workspace

Processes is a predefined workspace that displays information about the processes for SQL servers. The predefined workspace contains the following views:

- Process Summary table view that displays information about the processes (such as the number of processes and the statuses)

- CPU Percent bar chart that displays data from selected columns in the Process Summary table view (such as the percentages of CPU being used by the system and the application)
- Process Summary bar chart that displays data from selected columns in the Process Summary table view (such as the total processes and total processes blocked)
- Process Detail table view that displays information about specific processes (such as the usage and time)

Server Batch Statistics workspace

Server Batch Statistics is a predefined workspace that displays information about batch requests that are sent to the SQL Server. The predefined workspace contains the following views:

- Server Batch Statistics table view that displays information about batch requests, such as batch requests sent to the SQL Server per second and the number of failed auto-parameterization attempts per second
- Rate of Batch Request plot chart view that displays the status of batch requests that are sent to the SQL Server per second

Server Configuration workspace

Server Configuration is a predefined workspace that displays information about the configuration of the SQL servers. The predefined workspace contains the following views:

- Server Summary table view that displays summary information about the SQL servers (such as the name and status)
- Cache Utilization bar chart that displays data from selected columns in the Server Summary table view (such as the size of the caches for data and procedures)
- Server Options table view that displays performance information about the selected server (such as the status and the size of the data and procedure caches)
- Server Configuration table view that displays information about specific configuration values (such as the name and value for the parameter)

Server Detail workspace

Server Detail is a predefined workspace that gives the status of the SQL Server. The predefined workspace contains the following views:

- Server Detail – Status view that displays the status of the SQL Server
- Server Detail – Statistics view that displays statistics of the SQL Server

Services Detail workspace

Services Detail is a predefined workspace that provides details and the status of the SQL Server Windows services. The predefined workspace contains the following views:

- Situation Event Console view that displays all events on the server
- Services Details view that displays information about each SQL service

Server Locking workspace

Server Locking is a predefined workspace that displays information about the locks for SQL servers. The predefined workspace contains the following views:

- Lock Conflict table view that displays information about the locks conflicts (such as the type of lock and the ID of the process holding the lock)
- Lock Detail table view that displays information about specific locks (such as the type of lock and the cursor associated with the lock)

Server Statistics workspace

Server Statistics is a predefined workspace that displays statistical information for SQL servers. The predefined workspace contains the following views:

- Statistics Summary table view that displays summary statistics for the SQL servers (such as the number of active connections)
- Statistics Summary bar chart and IO Summary bar chart that display data from selected columns in the Statistics Summary table view
- Statistics Detail table view that displays each statistic for the selected SQL server (such as the name of the statistic and the average, current, minimum, and maximum value for the statistic)

Server Summary workspace

Server Summary is a predefined workspace that displays an overview of the SQL server activity for your environment. The predefined workspace contains the following views:

- Server Summary table view that displays summary information about the SQL Server (such as the name and status)
- CPU Utilization bar chart view that displays data from selected columns in the Server Summary table (such as the percentage of CPU being used by the server)
- Server Properties table view that displays the properties of the SQL Server

SQL Server CPU Bottleneck workspace

SQL Server CPU Bottleneck is a composite workspace that displays real-time data about the SQL Server process CPU utilization for all the SQL Server nodes. This workspace contains the following views:

- SQL Server Process User Time Percent view that displays the user time and the privileged time (in percentage) used by the SQL Server process.
- Total CPU Time Percent view that displays the total user time, the privileged time, and the processor time (in percentage) of the CPU.

SQL Server Disk Space and Logon Statistics

SQL Server Disk Space and Logon Statistics is a composite workspace that displays information about the amount of available free space and active logons for all the SQL Server instances. This workspace contains the following views:

- Minimum Percentage Data and Log Freespace Availability view that displays the minimum percentage of free space that is available on the disk for all the SQL Server instances.
- Server Logon Details view that displays the current active logons and average number of active logons for all the SQL Server instances.

SQL Server Integration Service Details workspace

SQL Server Integration Service Details is a predefined workspace that displays information about the integration service pipeline for the current instance of the SQL Server. This workspace contains the following views:

- Integration Service Pipeline Buffer Statistics view that displays a bar chart of the flat buffer memory and the private buffer memory statistics about the integration service pipeline.
- Integration Service Pipeline Details view that displays a table with detailed information about the integration service pipeline.

SQL Server Logical Disk Bottleneck workspace

SQL Server Logical Disk Bottleneck workspace is a composite workspace that displays the current logical disk activities. This workspace contains the following views:

- Logical Disk Read/Write Percent view that displays the details about the percentage of time that is spent by the logical disk in servicing read and write requests.

- Disk Queue Length view that displays the details about the total number of pending requests on the logical disk.
- Disk Space view that displays the details about the disk volume.

SQL Server Memory Bottleneck workspace

SQL Server Memory Bottleneck is a composite workspace that displays the memory and the processor utilization of the SQL Server. This workspace contains the following views:

- SQL Server Process Memory Utilization view that displays the memory used by the SQL Server.
- SQL Server Processor Utilization Percent view that displays the percentage of CPU time that the SQL Server process uses, and the percentage of CPU time that the other processes on the host use.
- Memory Pages to be Written per Sec view that displays the number of memory buffers that are written by the lazy writer per second, and the number of pages that are flushed to the disk per second by the checkpoint.
- Cache Hit Ratio view that displays the current ratio of the number of data cache hits to the number of total requests.
- Stolen Pages view that displays the details about the number of pages that are used for miscellaneous server purposes.

SQL Server Physical Disk Bottleneck workspace

SQL Server Physical Disk Bottleneck is a composite workspace that displays the current physical disk activities. This workspace contains the following views:

- Physical Disk Read/Write Percent view that displays details about the percentage of time that is spent by the physical disk in servicing read and write requests.
- Disk Queue Length view that displays the details about the number of pending requests on a physical disk including the requests that are currently in service.
- Physical Disk Average Read/Write per Sec view that displays the average number of read and write requests that are serviced on the physical disk per second.

SQL Server Service Broker Information workspace

SQL Server Service Broker Information is a predefined workspace that displays information about Service Broker activation, Service Broker statistics, and Service Broker transport. This workspace contains the following views:

- Service Broker Activation Details view that displays the details of the Service Broker activation of each database in an SQL Server instance.
- Service Broker Statistics view that displays a summary of the Service Broker statistics.
- Service Broker Transport view that displays the summary of the activities that are related to the Server Broker transport.

SQL Server Transactions Summary workspace

SQL Server Transactions Summary is a predefined workspace that displays the summary of all the transactions on the SQL Server since the server was started. The predefined workspace contains the following views:

- Total Active Transactions Statistics view that displays a bar chart of the number of the total active transactions, total non-snapshot transactions, and total snapshot transactions in the SQL Server.
- Version Store Unit Statistics view that displays a bar chart of the details about the version store unit count, the version store unit creation, and the version store unit truncation in the SQL Server.
- Server Transaction Summary view that displays the summary of the transaction statistics in the SQL Server.

SQL Text for Process ID workspace

SQL Text for Process ID is a predefined workspace that displays information about the SQL text string for the selected process. The predefined workspace contains the following views:

- SQL text table view that displays detailed information about the SQL text strings (such as the total CPU time)
- Event Console view

Table Detail workspace

Table Detail is a predefined workspace that provides details about the SQL Server database tables. The predefined workspace contains the following views:

- Fragmentation (greater than 30%) by Index Name view that displays a bar chart of the fragmentation for each index
- Table Detail view that displays statistics for each table

Table Information workspace

Table Information is a predefined workspace that provides details about the SQL Server database tables. The predefined workspace contains the following views:

- Fragmentation (greater than 30%) by Index Name view that displays a bar chart of the fragmentation for each index.
- Table Detail view that displays the detailed information for each table in the current instance of the SQL Server.
- Table Summary view that displays the summarized information of all the tables in the current instance of the SQL Server.

Cross-linked workspaces

This section contains cross-links within the SQL Server agent navigator items and workspaces. The following table contains source, target, and filtering criteria for the cross-linked workspaces.

Table 7. Cross-linked workspaces within the SQL Server agent

Source navigator item	Source workspace name	Source view name	Target navigator item	Target workspace name	Filter
AlwaysOn	Availability Database Information	Availability Database Details	AlwaysOn	Additional Availability Database Details	N/A
AlwaysOn	Availability Groups Information	Availability Groups Details	AlwaysOn	Availability Replicas and Database Information	N/A

Chapter 4. Attributes reference

Attributes are organized into attribute groups. Attributes in an attribute group relate to a single object such as an application, or to a single kind of data such as status information.

Attributes in a group can be used in queries, query-based views, situations, policy workflows, take action definitions, and launch application definitions. Chart or table views and situations are two examples of how attributes in a group can be used:

- Chart or table views

Attributes are displayed in chart and table views. The chart and table views use queries to specify which attribute values to request from a monitoring agent. You use the Properties editor to apply filters and set styles to define the content and appearance of a view based on an existing query.

- Situations

You use attributes to create situations that monitor the state of your operating system, database, or application. A situation describes a condition you want to test. When you start a situation, the values you assign to the situation attributes are compared with the values collected by the Microsoft SQL Server agent and registers an *event* if the condition is met. You are alerted to events by indicator icons that are displayed in the Navigator.

Additional information about attributes

For more information about using attributes and attribute groups, see the *Tivoli Enterprise Portal User's Guide*.

For a list of the attributes groups, a list of the attributes in each attribute group, and descriptions of the attributes for this monitoring agent, refer to the Attributes in each attribute group section in this chapter.

Attributes in each attribute group

This monitoring agent contains the following attribute groups:

- MS SQL Batch Stats
- MS SQL Configuration
- MS SQL Database Detail
- MS SQL Database Mirroring
- MS SQL Database Summary
- MS SQL Device Detail
- MS SQL Enterprise Database Mirroring
- MS SQL Enterprise Log Shipping DB Details
- MS SQL Filegroup Detail
- MS SQL Job Detail
- MS SQL Job Summary
- MS SQL Lock Conflict Detail
- MS SQL Lock Detail
- MS SQL Lock Resource Type Summary
- MS SQL Lock Summary
- MS SQL Log Shipping DB Details
- MS SQL Log Shipping Errors

- MS SQL Log Shipping Jobs Detail
- MS SQL Memory Manager
- MS SQL Problem Detail
- MS SQL Problem Summary
- MS SQL Process Detail
- MS SQL Process Summary
- MS SQL Remote Servers
- MS SQL Resource Pool Stats
- MS SQL Server Detail
- MS SQL Server Summary
- MS SQL Services Detail
- MS SQL Statistics Detail
- MS SQL Statistics Summary
- MS SQL Table Detail
- MS SQL Text
- MS SQL Workload Group Stats
- VMS SQL Remote Servers
- VMS SQL Server Enterprise View

IBM Tivoli Monitoring provides other attribute groups that are available to all monitoring agents, for example Universal Time and Local Time. The attributes in these common attribute groups are documented in the Tivoli Enterprise Portal Help.

The following sections contain descriptions of the Microsoft SQL Server agent attribute groups, which are listed alphabetically. Each description contains a list of attributes in the attribute group.

Note: If all attributes in an attribute group contain the values Not Collected, Not Applicable or No Data, during the execution of the historical data collection (HDC), the .hdr file is not created for the attribute group, and HDC data is not visible for the attribute group.

MS SQL Additional Database Details attributes

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

Containment Indicates whether containment is enabled for the database. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.
No Containment	1	The database is not contained.
Partial Containment	2	The database is partially contained.
Full Containment	3	The database is fully contained.

Database ID The database ID. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.

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 following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_Applicable	The value is not supported for SQL Server 2008 or earlier.

Filestream Directory Name The name of the filestream directory. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_Applicable	The value is not supported for SQL Server 2008 or earlier.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Non Transactional Access Level The level of non-transactional access for the database. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.
ReadOnly	1	The level of non-transactional access for the database is read-only access.
Full Access	2	The level of non-transactional access for the database is full access.
In Transition to ReadOnly	3	The level of non-transactional access for the database is in transition to read-only access.
In Transition to Off	4	The level of non-transactional access for the database is in transition to off.

Originnode The key for the table view in the format serverid:hostname:MSS.

Server The name of the SQL Server.

MS SQL Additional Availability Database Details attributes

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

Database Name The name of the availability database. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not applicable.

Filestream Sent Rate (KB Per Sec) The rate (in KB per second) at which filestream data is sent to the secondary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not applicable.
No Data	-3	No data is collected.

Group Database ID The ID of the availability group database. The following values are valid:

External value	Internal value	Description
Not Applicable	Not_applicable	The value is not applicable.
Not Collected	Not_collected	The value cannot be collected.
No Data	nodata	No data is collected.

Group ID The ID of the availability group to which the availability database belongs. The following values are valid:

External value	Internal value	Description
Not Applicable	Not_applicable	The value is not applicable.
Not Collected	Not_collected	The value cannot be collected.
No Data	nodata	No data is collected.

Hub Timestamp The time when data is inserted at the hub.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not applicable.
No Data	-3	No data is collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not applicable.
No Data	-3	No data is collected.

Originnode The key for the table view in the format serverid:hostname:MSS.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not applicable.
No Data	-3	No data is collected.

Redo Rate (KB Per Sec) The rate (in KB per second) at which log records are being redone. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not applicable.
No Data	-3	No data is collected.

Role Indicates whether a replica is currently a primary or secondary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not applicable.
No Data	-3	No data is collected.
Resolving	1	The replica is recovering from a failover.
Primary	2	The replica is a primary replica.
Secondary	3	The replica is a secondary replica.
Invalid	4	The availability group that is associated with the database is offline.

Server The name of the SQL Server.

MS SQL Analysis Services Data Mining attributes

The MS SQL Analysis Services Data Mining attribute group provides detailed information about the data mining of the SQL Server Analysis Services.

Cases Per Sec The rate at which cases are currently processed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Concurrent Data Mining Queries The number of data mining queries that are currently being processed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Current Models Processing The number of SQL Server Analysis Services Data Mining models that are currently being processed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Hub Timestamp The time when data is inserted at the hub.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Predictions Per Sec The rate at which predictions that are generated in data mining queries. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Queries Per Sec The rate at which data mining queries are generated. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Rows Per Sec The number of rows that are currently processed per second for a data mining prediction query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server The name of the SQL Server. The value format is an alphanumeric string with a maximum of 30 characters.

Total Predictions The number of data mining prediction queries that are currently received by the server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Queries The number of data mining queries that are currently received by the server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Rows The number of rows that are currently returned by data mining queries. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Analysis Services MDX Processing attributes

The MS SQL Analysis Services MDX Processing attribute group provides a summary of data that is processed by Multidimensional Expressions (MDX).

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Hub Timestamp The time when data is inserted at the hub.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Server The name of the SQL Server.

Total Autoexist The number of times that the Autoexist operation is currently performed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Total Calculation Cache Registered The number of calculation caches that are currently registered. This attribute is supported only on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Total Cells Calculated The number of cell properties that are currently calculated. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Total EXISTING Operators The number of times that the EXISTING set operation is currently performed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Flat Cache Inserts The number of cell values that are currently inserted into the flat calculation cache. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Total NonEmpty Algorithms The number of times a NonEmpty algorithm is currently used. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total NonEmpty Algorithms For Calculated Members The number of times that the NonEmpty algorithm currently looped through the calculated members. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total NonEmpty Unoptimized Algorithms The number of times that a NonEmpty algorithm that is not optimized is currently used. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Recomputes The number of cells that are currently computed again because of an error. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Sonar Subcubes The number of subcubes that are currently generated by the query optimizer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Storage Engine Evaluation Nodes The number of storage engine evaluation nodes that are currently built by the MDX execution plans. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

MS SQL Analysis Services Memory Statistics attributes

The MS SQL Analysis Services Memory Statistics attribute group provides information about memory that is used by the SQL Server Analysis Services.

Aggregation Map Files The current number of aggregation map files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Aggregate Cache (KB) The amount of memory that is currently allocated to the file cache. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Cleaner Balance Per Sec The rate at which balance and shrink operations are performed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Cleaner Current Price The current price of memory in dollars per byte and per unit of time. This price is normalized to 1000. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Cleaner Memory Non-Shrinkable (KB) The amount of memory that cannot be purged by the background cleaner. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Cleaner Memory Shrinkable (KB) The amount of memory that can be purged by the background cleaner. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Cleaner Memory (KB) The sum of shrinkable memory and the amount of non-shrinkable memory that is currently known to the background cleaner. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Cleaner Memory Shrunk KB Sec The rate at which the cleaner memory is shrunk. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Dimension Index (Hash) Files The current number of dimension index (hash) files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Dimension Property Files The current number of dimension property files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Dimension String Files The current number of dimension string files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Fact Aggregation Files The current number of fact aggregation files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Fact Data Files The current number of fact data files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Fact String Files The current number of fact string files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore Clock Pages Valid Per Sec The rate at which the background cleaner examines the filestore clock pages that are valid for eviction. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore Page Faults Per Sec The rate at which filestore page faults are generated. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore (KB) The amount of memory that is currently allocated to the filestore. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore KB Reads Per Sec The rate at which filestore operations are read. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore KB Writes Per Sec The rate at which filestore operations are written. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore Memory Pinned (KB) The amount of filestore memory that is currently pinned. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore Reads Per Sec The rate at which filestore pages are read. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Filestore Writes Per Sec The rate at which filestore pages are written. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, it is available for reports, queries, and workspaces.

In-Memory Map File (KB) The current size of the in-memory map file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Map File (KB) Per Sec The rate at which write operations occur in the in-memory map file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Other File (KB) The current size of the in-memory other file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Other File (KB) Per Sec The rate at which write operations occur in the in-memory other file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Property File (KB) The current size of the in-memory property file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Property File KB Per Sec The rate at which write operations occur in the in-memory property file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory String File (KB) The current size of the in-memory string file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory String File KB Per Sec The rate at which write operations occur in the in-memory string file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Aggregation Map File (KB) The current size of the in-memory aggregation map file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Aggregation Map File KB Per Sec The rate at which write operations occur in the in-memory aggregation map file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Dimension Index (Hash) File (KB) The current size of the in-memory dimension index (hash) file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Fact Aggregation File (KB) The current size of the in-memory fact aggregation file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Fact Aggregation File KB Per Sec The rate at which write operations occur in the in-memory fact aggregation file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Fact Data File (KB) The current size of the in-memory fact data file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Fact Data File KB Per Sec The rate at which write operations occur in the in-memory fact data file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Fact String File (KB) The current size of the in-memory fact string file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

In-Memory Fact String File (KB) Per Sec The rate at which write operations occur in the in-memory fact string file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Map Files The current number of map files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Memory Limit High (KB) The highest memory limit that is specified in the configuration file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Memory Limit Low (KB) The lowest memory limit that is specified in the configuration file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Memory Usage (KB) The amount of memory that is used by the server to calculate the cleaner memory price. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Other Files The current number of other files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Page Pool 1 Lookaside (KB) The amount of memory that is currently available in the 1 KB lookaside list. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Page Pool 64 Lookaside (KB) The amount of memory that is currently available in the 64 KB lookaside list. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Page Pool 8 Lookaside (KB) The amount of memory that is currently available in the 8 KB lookaside list. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Aggregation Map File (KB) The potential size of the in-memory aggregation map files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Dimension Index (Hash) File (KB) The potential size of the in-memory dimension index (hash) files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Dimension Property File (KB) The potential size of the in-memory dimension property files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Dimension String File (KB) The potential size of the in-memory dimension string files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Fact Aggregation File (KB) The potential size of the in-memory fact aggregation files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Fact Data File (KB) The potential size of the in-memory fact data files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Fact String File (KB) The potential size of the in-memory fact string files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Map File (KB) The potential size of the in-memory map files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Potential In-Memory Other File (KB) The potential size of the in-memory other files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Quota (KB) The current amount of memory quota. A memory quota is also referred to as a memory grant or memory reservation. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Quota Blocked The number of requests for memory quota that are currently blocked until some memory quotas are freed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server The name of the SQL Server.

Total Filestore I/O Errors The current number of filestore I/O errors. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Filestore I/O Errors Per Sec The rate at which filestore I/O errors occur. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Analysis Services Rows Processing attributes

The MS SQL Analysis Services Rows Processing attribute group provides a summary of the rows that are processed by the SQL Server Analysis Services.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Hub Timestamp The time when data is inserted at the hub.

Rows Converted Per Sec The rate at which rows are converted during processing per second. The following value is valid.

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Rows Read Per Sec The rate at which rows are read from all the relational databases per second. The following value is valid.

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Rows Written Per Sec The rate at which rows are written to the database during processing. The following value is valid.

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server The name of the SQL Server.

Total Rows Converted The number of rows that are currently converted during processing. The following value is valid.

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Rows Read The number of rows that are currently read from all the relational databases. The following value is valid.

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Rows Written The number of rows that are currently written to the database during processing. The following value is valid.

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Analysis Services Storage Engine Query attributes

The MS SQL Analysis Services Storage Engine Query attribute group provides the summary of statistics that are related to the storage engine query processing feature of the SQL Server Analysis Services.

Aggregation Hits Per Sec The rate of aggregation hits. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

External value	Internal value	Description
Not Applicable	-2	The attribute is not supported.

Aggregation Lookups Per Sec The rate of aggregation lookups. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Current Dimension Queries The number of dimension queries that are currently being processed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Current Measure Group Queries The number of measure group queries that are currently being worked on. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Current Pyramid Operations The number of pyramid operations that are currently in progress. This attribute is supported only on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Data Bytes Per Sec The number of bytes that are currently read per second from the data file. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Data Reads Per Sec The number of logical read operations that are currently using the data file per second. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Dimension Cache Hits Per Sec The number of dimension cache hits per second. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Dimension Cache Lookups Per Sec The rate of dimension cache lookups. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Dimension Queries Per Sec The rate at which dimension queries are processed. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

External value	Internal value	Description
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Filtered Rows Per Sec The rate at which rows are currently filtered. This attribute is supported only on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Hub Timestamp The time when data is inserted at the hub.

Map Bytes Per Sec The number of bytes that are currently read from the map file per second. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Map Reads Per Sec The number of logical read operations that are currently using the map file per second. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Measure Group Cache Hits Per Sec The rate of group cache hits. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Measure Group Lookups Per Sec The rate of measure group cache lookups. This attribute is not supported on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Measure Group Queries Per Sec The rate of measure group queries that are sent to the server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Queries Answered Per Sec The rate at which queries are currently answered. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Pyramid Operations Per Sec The rate at which pyramid operations are currently started. This attribute is supported only on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Queries From Cache Direct Per Sec The rate at which queries are answered directly from the cache. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Queries From Cache Filtered Per Sec The rate at which queries are answered by filtering the existing cache entry. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Queries From Files Per Sec The rate at which queries are answered from files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Rows Sent Per Sec The rate at which rows are sent by the server to the clients. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server The name of the SQL Server.

Total Dimension Queries The current number of dimension queries. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Measure Group Queries The current number of measure group queries. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Pyramid Operations The current number of pyramid operations. This attribute is supported only on SQL Server 2005. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The attribute is not supported.

Total Queries Answered The current number of queries that are answered. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Queries From Cache Direct The number of queries that are currently derived from the cache for every partition. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Queries From Cache Filtered The number of queries that are currently answered by filtering the existing cache entries. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Queries From File The number of queries that are currently answered from files. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Rows Sent The number of rows that are currently sent by the server to the clients. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Bytes Sent The number of bytes that are currently sent by the server to the clients in response to queries. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Availability Database Details attributes

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

Database ID The ID of the availability database that is hosted by the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Database Name The name of the availability database that is hosted by the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Database State Indicates the current state of the availability database. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

External value	Internal value	Description
No Data	-3	No data is collected.
Online	1	The database is available for access.
Restoring	2	The database is currently unavailable because of one of the following reasons: <ul style="list-style-type: none"> • One or more files of the primary group are being restored. • One or more secondary files are being restored offline.
Recovering	3	The database is currently unavailable and is being recovered. If the database recovery is successful, the database state changes to Online. If the database recovery fails, the database state changes to Suspect.
Recovery Pending	4	The database is currently unavailable because of one of the following reasons: <ul style="list-style-type: none"> • Some files are missing. • The database cannot start because of system resource limitations.
Suspect	5	The database is currently unavailable because of the following reasons: <ul style="list-style-type: none"> • The primary file group is damaged. • Database cannot be recovered during the SQL Server startup.
Emergency	6	The database is currently in the single-user mode, and is being repaired or restored. The following actions are completed when the database is in the Emergency mode: <ul style="list-style-type: none"> • The database is marked read-only. • Logging is disabled. • Access is limited to the members of the sysadmin fixed server role.
Offline	7	The database is not available for access.

Group Database ID The ID of the availability group database that is hosted by the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group ID The ID of the availability group to which the availability database belongs. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The name of the availability group to which the availability database belongs. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Hub Timestamp The time when data is inserted at the hub.

Local Indicates whether the availability replica is hosted by the local SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
No	1	The replica is not hosted by the local instance.
Yes	2	The replica is hosted by the local instance.

Originnode The key for the table view in the format serverid:hostname:MSS.

Replica ID The ID of the availability replica. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Role Indicates whether the availability database is a primary or a secondary database. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Resolving	1	The system is trying to gather information about the role of the database.
Primary	2	The database is the primary database.
Secondary	3	The database is the secondary database.
Invalid	4	The availability group that is associated with the database is offline.

Server The name of the SQL Server.

Suspend Reason The reason for a database to enter the suspended state. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected. This enum is shown if the database is not suspended.
Suspend From User	1	The reason for the database to be suspended is Suspend From User.
Suspend From Partner	2	The reason for the database to be suspended is Suspend From Partner.
Suspend From Redo	3	The reason for the database to be suspended is Suspend From Redo.
Suspend From Apply	4	The reason for the database to be suspended is Suspend From Apply.
Suspend From Capture	5	The reason for the database to be suspended is Suspend From Capture.
Suspend From Restart	6	The reason for the database to be suspended is Suspend From Restart.
Suspend From Undo	7	The reason for the database to be suspended is Suspend From Undo.
Suspend From Revalidation	8	The reason for the database to be suspended is Suspend From Revalidation.

Synchronization Health Indicates the health of the availability database replica. The replica is healthy if the synchronized and the asynchroninzed replicas are in the Synchronized state. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Not Healthy	1	The synchronization health of the database replica that is hosted on the SQL Server instance is not healthy.
Partially Healthy	2	The synchronization health of the database replica that is hosted on the SQL Server instance is partially healthy.
Healthy	3	The synchronization health of the database replica that is hosted on the SQL Server instance is healthy.

Synchronization State The current synchronization state of the availability database replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

External value	Internal value	Description
Not Synchronizing	1	On a primary database, this value indicates that the database is not ready to synchronize the transaction log with the secondary database. On a secondary database, this value indicates that the database is not synchronizing because of one of the following reasons: <ul style="list-style-type: none"> • The database has not started log synchronization. • The database is being suspended. • The database is going through the transition state during the startup. • The database role is changing.
Synchronizing	2	On a primary replica, this value indicates that the database is ready to accept a scan request from the secondary database. On a secondary replica, this value indicates that an active data movement is in progress.
Synchronized	3	On a primary replica, this value indicates that at least one secondary database is synchronized, and is available as a target for a failover without data loss. On a secondary replica, this value indicates that all secondary databases are synchronized with the corresponding primary databases.
Reverting	4	Indicates that the availability database replica is in the Reverting state. Note: If the availability database replica is forced to shut down when the replica is in the Reverting state, you may not be able to start the replica again.
Initializing	5	Indicates that the availability database replica is in the Initializing state. Note: If the availability database replica is forced to shut down when the replica is in the Initializing state, you may not be able to start the replica again.

Synchronized Commit Indicates whether the transaction commit is synchronized with the database replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Not Synchronized	1	The transaction commit is not synchronized with the database replica.
Synchronized	2	The transaction commit is synchronized with the database replica.

MS SQL Availability Database Details In Cluster attributes

The MS SQL Availability Database Details In Cluster attribute group 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 attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Database Name The name of the availability database. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Failover Ready Indicates whether the secondary database is synchronized with the primary database. The following values are valid:

Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Not Ready	1	The secondary database is not synchronized with the primary database.
Ready	2	The secondary database is synchronized with the primary database.

Group Database ID The availability group database ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group ID The availability group ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The availability group name. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Hub Timestamp The time when data is inserted at the hub.

Join State Indicates whether the availability database on the availability replica is joined to the availability group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.
No Data	-3	No data is collected.
Not Joined	1	The availability database is not joined to the availability group.
Joined	2	The availability database is joined to the availability group.

Originnode The key for the table view in the format serverid:hostname:MSS.

Replica ID The replica ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Secondary Suspension Pending Indicates whether the database is pending suspension after a forced failover. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
No	1	The secondary database is not pending suspension.
Yes	2	The secondary database is pending suspension.

Server The name of the SQL Server.

MS SQL Availability Database Statistics attributes

MS SQL Availability Database Statistics attribute group provides information about the availability database statistics. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Database Name The name of the availability database. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

File Bytes Received Per Sec The amount of filestream data that is currently being received per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Hub Timestamp The time when data is inserted at the hub.

Log Bytes Received Per Sec The amount of log bytes that is currently being received per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Originnode The key for the table view in the format serverid:hostname:MSS.

Recovery Queue The number of log records that are in the recovery queue because they are not redone. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Redo Blocked Per Sec The number of times that the redo operation is currently being blocked per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Redo Bytes Remaining (KB) The amount of log bytes (in KB) that has not completed the redo phase. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Redone Bytes Per Sec The amount of log bytes that was redone on the secondary database in the last second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Server The name of the SQL Server.

Total Log Requiring Undo (KB) The amount of log data (in KB) that must be undone. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Undo Log Remaining (KB) The amount of log data (in KB) that remains to be undone. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

MS SQL Availability Database Summary attributes

The MS SQL Availability Database Summary attribute group provides a summary of the information about availability databases that are hosted by the SQL Server instance. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Originnode The key for the table view in the format serverid:hostname:MSS.

Max Unsynchronized Databases The maximum number of the availability databases that have been unsynchronized since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Hub Timestamp The time when data is inserted at the hub.

Max Non Healthy Databases The maximum number of the availability databases that have been non-healthy since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Max Non Synchronized Commit Database Replicas The maximum number of availability database replicas on which the transaction commit operation has not been synchronized since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Max Non-Online Databases The maximum number of databases that have not been online since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Min Non Healthy Databases The minimum number of the availability databases that have been non-healthy since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Min Non Synchronized Commit Database Replicas The minimum number of availability database replicas on which the transaction commit operation has not been synchronized since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Min Non-Online Databases The minimum number of databases that have not been online since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Min Unsynchronized Databases The minimum number of the availability databases that have been unsynchronized since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Server The name of the SQL Server.

Total Non Healthy Databases The total number of the availability databases that have been non-healthy since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Non Synchronized Commit Database Replicas The total number of availability database replicas on which the transaction commit operation has not been synchronized since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Non-Online Databases The total number of databases that have not been online since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Databases The total number of availability databases that are hosted by the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Primary Databases The total number of primary availability databases that are hosted by the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Secondary Databases The total number of secondary availability databases that are hosted by the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Unsynchronized Databases The total number of availability databases that have been unsynchronized since the agent startup. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

MS SQL Availability Groups Details attributes

The MS SQL Availability Groups attribute group provides details about the availability groups that are hosted by the local SQL Server instance. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Note: In this attribute group, the Primary Replica ID is blank if the SQL query returns the NULL value.

Backup Preference The preferred location to perform backup operations on the availability databases that are available in the availability group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Primary Replica	1	The backup operation is performed on the primary replica.
Secondary Replica	2	The backup operation is performed on the secondary replica.
Primary Replica acceptable if no Secondary Replica Available	3	The backup operation is performed on the primary replica if a secondary replica is not available.
No Preference	4	No preference is specified for performing backup operations.

Database Count The total number of databases that are connected to the availability group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

Failure Condition The condition level that is specified for triggering an automatic failure of the availability group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
SQL Service Down or Group Lease Expired	1	An automatic failure is triggered if the SQL service is down, or the group lease has expired.

External value	Internal value	Description
Health check timeout Threshold Group Exceeded	2	An automatic failure is triggered when the availability group exceeds the health check timeout threshold.
Critical SQL Server Internal Errors	3	An automatic failure is triggered when the system encounters critical SQL Server internal errors.
SQL Server Internal Errors	4	An automatic failure is triggered when the system encounters SQL Server internal errors.
Unsolvable Deadlock or Exhaustion of SQL Engine	5	An automatic failure is triggered when the system encounters an unsolvable deadlock or the exhaustion of the SQL engine.

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. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

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 following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Group Role Indicates whether the local server instance is a primary or a secondary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Resolving	1	The system is trying to gather information about the role of the replica.
Primary	2	The replica is a primary replica.
Secondary	3	The replica is a secondary replica.
Invalid	4	The availability group that is associated with the replica is offline.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Resolving	1	The replica is recovering from a failover.
Primary	2	The replica is a primary replica.
Secondary	3	The replica is a secondary replica.
Invalid	4	The availability group that is associated with the replica is offline.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
In Progress	1	The replica is recovering from a failover.
Online	2	The replica is online.

Orignode The key for the table view in the format serverid:hostname:MSS.

Primary Replica The name of the primary replica. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Primary Replica ID The ID of the primary replica. The format is an alphanumeric string with a maximum of 128 characters. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Secondary Replicas Count The number of secondary replicas in the availability group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

Server The name of the SQL Server.

Synchronization Health A summary of the synchronization health of all the replicas in the availability group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Not Healthy	1	All the replicas in the availability group are not healthy.
Partially Healthy	2	At least one replica in the availability group is not healthy.
Healthy	3	All the replicas in the availability group are healthy.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

MS SQL Availability Groups Details In Cluster attributes

The MS SQL Availability Groups Details In Cluster attribute group 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 attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Backup Preference The preferred location that is specified for performing backup operations on the availability database. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Primary Replica	1	The primary replica is specified as the preferred location for performing backup operations on the availability database.
Secondary Replica	2	The secondary replica is specified as the preferred location for performing backup operations on the availability database.
Primary Replica acceptable if no Secondary Replica Available	3	The primary replica is specified as the location for performing backup operations on the availability database if no secondary replica is available.

External value	Internal value	Description
No Preference	4	No preferred location is specified for performing backup operations on the availability database.

Cluster Resource ID The WSFC cluster resource ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Failure Condition The value that indicates when an automatic failover occurs. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
SQL Service Down or Group Lease Expired	1	An automatic failover occurs when the SQL service is down, or the group lease has expired.
Health Check Timeout Threshold Group Exceeded	2	An automatic failover occurs when the health check timeout threshold of the availability group has been exceeded.
Critical SQL Server Internal Errors	3	An automatic failover occurs because critical SQL Server internal errors are detected.
SQL Server Internal Errors	4	An automatic failover occurs because SQL Server internal errors are detected.
Unsolvable Deadlock or Exhaustion of SQL Engine	5	An automatic failover occurs because one of the following conditions exists: <ul style="list-style-type: none"> The worker threads of the SQL engine are exhausted An unsolvable deadlock exists

Group ID The availability group ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The availability group name. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Hub Timestamp The time when data is inserted at the hub.

Originnode The key for the table view in the format serverid:hostname:MSS.

Resource Group ID The WSFC cluster resource group ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_Applicable	The value is not supported.
No Data	nodata	No data is collected.

Server The name of the SQL Server.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

MS SQL Availability Group Listener Details attributes

The MS SQL Availability Group Listener Details attribute group provides information about the availability group listeners of all the availability groups that are hosted by the local SQL Server instance. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Originnode The key for the table view in the format serverid:hostname:MSS.

Group ID The availability group ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The availability group name. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Hub Timestamp The time when data is inserted at the hub.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.
No Data	-3	No data is collected.
No	1	The IP address of the availability groups listener is not conformant.
Yes	2	The IP address of the availability groups listener is conformant.

Is DHCP Indicates whether the IP address of the availability groups listener is configured by the Dynamic Host Configuration Protocol (DHCP). The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.
No Data	-3	No data is collected.
No	1	The IP address of the availability groups listener is not configured by the DHCP.
Yes	2	The IP address of the availability groups listener is configured by the DHCP.

Listener IP Address The IP address of the availability group listener. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Listener Name The name of the availability group listener. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Listener Subnet IP The IP address of the subnet to which the availability group listener belongs. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Listener ID The resource ID of the availability group listener in the WSFC cluster. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Port The TCP port number of the availability group listener. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected. This value indicates that the listener is configured outside the SQL Server and the port number of the listener is not added to the availability group

Server The name of the SQL Server.

State The state of the availability group listener in the WSFC cluster. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported for SQL Server 2008 or earlier.
No Data	-3	No data is collected.
Offline	1	The availability group listener is offline.
Online	2	The availability group listener is online.
Online Pending	3	The availability group listener is being brought online.
Failed	4	The availability group listener could not be brought online and has failed.

MS SQL Availability Groups Summary attributes

The MS SQL Availability Groups Summary attribute group provides a summary of the availability groups that are hosted by the local SQL Server instance. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Originnode The key for the table view in the format serverid:hostname:MSS.

Server The name of the SQL Server.

Total Groups The total number of availability groups that are available on the SQL Server instance in the current interval. The following values are valid:

External value	Internal value	Description
Not Applicable	Not_applicable	The value is not applicable.
Not Collected	Not_collected	The value is not supported.

Total Non Healthy Groups The total number of availability groups that are currently not healthy in the current interval. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

MS SQL Availability Replicas Details attributes

The MS SQL Availability Replicas Details attribute group provides details about the availability replicas in each availability group. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Availability Mode Indicates whether the availability replica is running in the synchronous-commit mode or the asynchronous-commit mode. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Asynchronous Commit	1	The availability replica is running in the asynchronous-commit mode.
Synchronous Commit	2	The availability replica is running in the synchronous-commit mode.

Backup Priority The priority that is specified to perform backup operations on the replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

Note: The value of this attribute is displayed as 0 if you select the **Exclude Replica** check box while configuring an availability replica.

Create Date The date when the replica was created. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Note: 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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Endpoint URL The database mirroring endpoint that is specified by the user. The value format is an alphanumeric string. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Failover Mode Indicates whether the failover mode of the availability replica is manual or automatic. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Manual	2	The failover mode of the replica is manual.
Automatic	1	The failover mode of the replica is automatic.

Group ID The ID of the availability group to which the replica belongs. The value format is an alphanumeric string. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The name of the availability group to which the replica belongs. The value format is an alphanumeric string. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Hub Timestamp The time when data is inserted at the hub.

Modify Date The date when the replica was last modified. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Note: This attribute shows No Data if the SQL Server query returns the NULL value.

Originnode The key for the table view in the format serverid:hostname:MSS.

Owner SID The security identifier (SID) of the external owner of the replica. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Primary Connection Type Indicates the type of connection that is allowed for a primary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
All	3	All types of connections are allowed.
Read Write	4	Connections with read/write access are allowed.

Replica ID The ID of the replica. The value format is an alphanumeric string. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Replica Server Name The server and instance name of the location that hosts the replica. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Role Indicates whether a replica is currently a primary or secondary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Resolving	1	The replica is recovering from a failover.
Primary	2	The replica is a primary replica.
Secondary	3	The replica is a secondary replica.
Invalid	4	The availability group that is associated with the replica is offline.

Secondary Connection Type Indicates the type of connection that is allowed for a secondary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
No	1	No connections are allowed.
Read Only	2	Connections with read-only access are allowed.
All	3	All types of connections are allowed.

Server The name of the SQL Server.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

MS SQL Availability Replicas Details In Cluster attributes

The MS SQL Availability Replicas Details In Cluster attribute group provides detailed information about each availability replica in the Windows Server Failover Clustering (WSFC) cluster. This attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Cluster Node The name of the cluster node where the availability replica is located. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Group ID The availability group ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The availability group name. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Hub Timestamp The time when data is inserted at the hub.

Join State Indicates whether the availability replica is joined to the availability group in the WSFC cluster. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Not Joined	1	The availability replica is not joined to the availability group.

External value	Internal value	Description
Joined As Stand Alone Instance	2	The availability replica is joined to the availability group as a stand-alone failover cluster instance.
Joined As Failover Cluster Instance	3	The availability replica is joined to the availability group as a cluster instance.

Originnode The key for the table view in the format serverid:hostname:MSS.

Replica Server Name The server and instance name of the location where the availability replica is hosted. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Note: If no data has been collected from the query for this attribute, the value of this attribute is displayed as blank on the Tivoli Enterprise Portal.

Replica ID The availability replica ID. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Server The name of the SQL Server.

MS SQL Availability Replicas Statistics attributes

The MS SQL Availability Replicas Statistics attribute group 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 attribute group is supported for SQL Server 2012 Enterprise Edition, or later.

Bytes Received Per Sec The number of bytes that are currently being received from the availability replica per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Flow Control Per Sec The number of flow control operations that have been currently initiated per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Hub Timestamp The time when data is inserted at the hub.

Instance Name The name of the local SQL Server instance. The following value is valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Originnode The key for the table view in the format serverid:hostname:MSS.

Received Messages Per Sec The number of AlwaysOn messages that are currently being received from the primary replica per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Resent Messages Per Sec The number of AlwaysOn messages that are currently being resent to the availability replica per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Sent Messages Per Sec The number of AlwaysOn messages that are currently being sent to the availability replica per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Server The name of the SQL Server.

MS SQL Availability Replicas Status attributes

The MS SQL Availability Replicas Status attribute group provides details about the status of the replicas in each availability group. This attribute group 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 attribute group 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 attribute group displays the details of only those replicas that are hosted by the local instance.

Connection State The current connection state of the availability replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Disconnected	1	The availability replica is not connected.
Connected	2	The availability replica is connected.

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. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Group Name The name of the availability group to which the replica belongs. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

Hub Timestamp The time when data is inserted at the hub.

Last Connect Error Number The last connection error number. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

Local Indicates whether the replica is hosted by the local instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
No	1	The replica is not hosted by the local instance.
Yes	2	The replica is hosted by the local instance.

Operational State The current operational state of the availability replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Pending Failover	1	A failover is being processed for this availability group.
Pending	2	The primary replica currently cannot process any requests because no database worker threads are available to process the requests.
Online	3	The availability replica is currently online.
Offline	4	The availability replica is currently offline.
Failed	5	The availability replica cannot process any requests.
Failed No Quorum	6	The local server instance node does not have quorum.
Invalid	7	The availability group that is associated with the replica is offline.

Note: This attribute shows the value No Data on the replica that is not local.

Originnode The key for the table view in the format serverid:hostname:MSS.

Recovery Health Indicates whether the databases that are connected to the availability group are online or are being recovered after a failover. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
In Progress	1	At least one database that is connected to the availability group is not online.
Online	2	All databases that are connected to the availability group are online.

Note: This attribute shows the value No Data on the replica that is not local.

Replica ID The ID of the replica. The value format is an alphanumeric string with a maximum of 64 characters. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.
No Data	nodata	No data is collected.

Role Indicates whether the replica is a primary or a secondary replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Resolving	1	The system is gathering information about the role of the availability replica.
Primary	2	The replica is a primary replica.
Secondary	3	The replica is a secondary replica.
Invalid	4	The availability group that is associated with the replica is offline.

Server The name of the SQL Server.

Synchronization Health The synchronization state of all the databases that are connected to the availability group on the availability replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Not Healthy	1	The synchronization health of all the replicas that are connected to the availability group is not healthy.

External value	Internal value	Description
Partially Healthy	2	The synchronization health of at least one replica that is connected to the availability group is not healthy.
Healthy	3	The synchronization health of all the replicas that are connected to the availability group is healthy.

Total Databases The total number of databases that are hosted by the availability replica. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.

MS SQL Availability Replicas Status Summary attributes

The MS SQL Availability Replicas Status Summary attribute group provides a summary of the status of availability replicas in all the availability groups that are hosted on the local server instance. This attribute group 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 attribute group 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 attribute group displays the details of only those replicas that are hosted by the local instance.

Hub Timestamp The time when data is inserted at the hub.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Originnode The key for the table view in the format serverid:hostname:MSS.

Server The name of the SQL Server.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Connected Replicas The total number of connected replicas in all the availability groups that are hosted on the local server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Disconnected Replicas The total number of disconnected replicas in all the availability groups that are hosted on the local server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Local Replicas The total number of local replicas in all the availability groups that are hosted on the local server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Primary Replicas The total number of primary replicas in all the availability groups that are hosted on the local server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Secondary Replicas The total number of secondary replicas in all the availability groups that are hosted on the local server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

External value	Internal value	Description
Not Applicable	-2	The value is not supported.

MS SQL Batch Stats attributes

The MS SQL Batch Stats attribute group contains attributes that you can use to monitor detailed information about batch requests.

Auto-Param Attempts Per Sec The number of auto-parameterization attempts per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Batch Requests Per Sec The number of Transact-SQL command batch requests received per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Failed Auto-Params Per Sec The number of failed auto-parameterization attempts per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Forced Parameterization Per Sec The number of successful forced parameterization attempts per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Originnode The key for the table view of the format serverid:hostname:MSS.

Safe Auto-Params Per Sec The number of safe auto-parameterization attempts per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server The name of the SQL Server.

Server Attention Rate The number of attentions per second. An attention is a request by the client to end the currently running request. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server Compilations Per Sec The number of server compilations per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Server Re-Compilations Per Sec The number of times the server recompiles per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Unsafe Auto-Params Per Sec The number of unsafe auto-parameterization attempts per second. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Blocked Queries Details attributes

The MS SQL Blocked Queries Details attribute group contains attributes that provide information about the queries that are currently blocked on the SQL Server.

Blocked Database Name (Unicode) The name of the database against which the blocked query is currently being run. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocked User Name (Unicode) The name of the user who has run the blocked SQL query. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocked Query Text The SQL text of the query that is currently blocked on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocked Session ID The session ID that is blocking the query that is currently being run on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocking Query Text The SQL text of the query blocks one or more than one queries on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocking Reason The reason for blocking the query. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocking Resource The resource for which the query is currently blocked on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Blocking Session ID The ID of the session that is currently blocking the execution of a query on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Blocking User Name The login name of the user who ran the blocking query. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Hub Timestamp The time when data is inserted at the hub.

Originnode The key for the table view in the format serverid:hostname:MSS.

Server The name of the SQL Server.

Total Wait Time (Sec) The total wait time (in seconds) since the query has been blocked on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Configuration attributes

The MS SQL Configuration attribute group contains attributes that you can use to monitor the configuration of a SQL server.

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.

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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Parm Type The type of configuration parameter. If the configuration parameter is dynamic, a server restart is not required. The following values are also valid:

External value	Internal value	Description
Dynamic	1	Configuration parameter is dynamic.
Static	0	Configuration parameter is static.
Advanced	2	Configuration parameter is advanced.
Dynamic_and_Advanced	3	Configuration parameter is dynamic and advanced.
Unknown	-1	Configuration parameter is unknown.

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.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

MS SQL Database Detail attributes

The MS SQL Database Detail attribute group contains attributes that you can use to monitor the performance and usage of MS SQL database.

Accessible Access validation of the database. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Yes	0	The database is accessible.
No	1	The database is not accessible.

Note: When the database is in offline, restoring, or transition state, the value of this attribute is displayed as No.

Active Transactions The number of active transactions for the database. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Data File Names The names of all physical files for data that make up this database.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Data Freespace The number of megabytes (MB) of free space for the database. The value includes the megabytes of free data only. The value does not include the free space available in the database. The value format is a decimal number with two decimal places allowed; for example, 10.00. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Data Freespace Percent The percentage of maximum available free space for the database. The value format is a percentage with two decimal places allowed; for example, 20.00. The product calculates this percentage from the maximum available free space for the database and the amount of space already allocated for the database. Consider using the alter database command to expand the database size. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Note: Calculating value for this attribute depends on the database file space configuration options that you specify when creating a database. Depending on these options, the SQL Server agent uses a different algorithm to calculate the data freespace percent.

Option one: Database created with the dynamic growth option

This algorithm considers the hard disk drive freespace for calculation.

Formula for calculation:

$$\text{Data Freespace Percent} = (\text{DB freespace} / \text{DB max size}) \times 100$$

Where:

- DB freespace is the total free space on all hard disk drives that contain the database files
- DB max size = total allocated database file size + database freespace

Option two: Database created with a fixed maximum size and the database grows dynamically till the maximum limit is reached.

This algorithm considers the maximum size of the database for calculation.

Formula for calculation:

$$\text{Data Freespace Percent} = (\text{DB freespace} / \text{DB max size}) \times 100$$

Where:

- DB freespace = (DB max size - DB current size)
- DB max size and DB current size read from the sysfiles table

Option three: Database created with the no growth option

This algorithm differs from the first two options. The database maximum size is the initial file size of the database at the time of database creation. The freespace is the freespace in the database file. In the first two options, reports freespace is derived from the physical hard disk drive space.

Formula for calculation:

$$\text{Data Freespace Percent} = (\text{DB freespace} / \text{DB max size}) \times 100$$

Where:

- DB freespace = unallocated database pages in the database file
- DB max size and DB current size read from the sysfiles table

Data Size 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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Database Growth Percent The percentage of growth for the database from the last sample to the current sample. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Database Maximum Growth Size The maximum size to which the database can grow in MB. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Database Name (Unicode) The name of the database. Valid values include text strings with a maximum of 384 bytes; for example, KOQ3. Each database name is unique. The SQL server also assigns each database its own identification number.

Database Space Percent Used The amount of 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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Database State Reports the database state. The possible values for this attribute are:

Value	Description
Online	Database is available for access.
Offline	Database is unavailable.
Restoring	One or more files of the primary filegroup are being restored, or one or more secondary files are being restored offline.
Recovering	Database is being recovered.
Recovery Pending	SQL Server has encountered a resource-related error during recovery.
Suspect	At least the primary filegroup is suspect and might be damaged.
Emergency	User has changed the database and set the status to Emergency.

When the agent fails to collect data, the following value is displayed:

External value	Internal value	Description
No Data	nodata	No data is collected.

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 following values are valid:

External value	Internal value	Description
Available	0	The database is available.
Not Available	1	The database is not available.

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. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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.

DB Owner (Unicode) The SQL server-assigned user ID for the owner of the database. Valid values include text strings with a maximum of 72 bytes. Use the create database command to establish this identifier.

DBO Only Access Indicates whether the database has a status of DBO only. The following value is valid:

External value	Internal value	Description
No	No	Database can be accessed by authorized users.
Yes	Yes	Database has a status of DBO only.

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	Year
MM	Month
DD	Day

For example, 20020501 indicates May 1, 2002. The following value is valid:

External value	Internal value	Description
Not Collected	nodata	The value cannot be collected.

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 following values are valid:

External value	Internal value	Description
No	No	Database does not have an error status.
Yes	Yes	Database has an error status.

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 following values are valid:

External value	Internal value	Description
No	No	Free space accounting is enabled.
Yes	Yes	Free space accounting is disabled.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Last Database Size The database size in MB from the previous sample. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Log Bytes Flushed per Sec The number of bytes of the log file that are flushed for the current interval. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Log File Names The names of all physical files for logs that make up this database.

Log Freespace 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 allowed; for example, 8.00. Various types of transactions, such as mass updates and bulk copying in can involve extensive logging. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 allowed; for example, 20.00. Various types of transactions, such as mass updates and bulk copying in can involve extensive logging. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Log Maximum Growth Size The maximum size to which the log can grow in MB. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Log Size 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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Log Space Percent 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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Long Running Transaction Name The long running transaction name. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Long Running Transaction Time Per Sec The time for which the longest running transaction is active (in seconds). The format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

No CKPT After Recovery Indicates whether a record for the checkpoint is added to the transaction log when the database is recovered. The following values are valid:

External value	Internal value	Description
No	No	No CKPT After Recovery option is disabled.
Yes	Yes	No CKPT After Recovery option is enabled.

Oldest Open Transaction 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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Page Verify The option that identifies and notifies incomplete I/O transactions that have occurred because of disk I/O errors. The following values are valid:

External value	Internal value
NONE	0
TORN PAGE DETECTION	1
CHECKSUM	2
Not collected	-1

Read Only Access Indicates whether the database has a status of read only. The following value is valid:

External value	Internal value	Description
No	No	Database is read and write database.
Yes	Yes	Database has a status of read-only.

Recovery Model The method used to maintain the transaction log. The three types of recovery models include Simple, Full, and Bulk-logged. The following values are valid:

External value	Internal value
FULL	1
BULK LOGGED	2
SIMPLE	3
Not collected	-1

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Replication Latency 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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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	Month
DD	Day
YY	Year
HH	Hour

mm	Minute
SS	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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Select Into Bulkcopy Enabled Indicates whether the select into/bulkcopy option is enabled for the database. The following values are valid:

External value	Internal value	Description
No	No	Select Into/Bulkcopy option is disabled.
Yes	Yes	Select Into/Bulkcopy option is enabled.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters. For example, CFS_SVR5.

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 following values are valid:

External value	Internal value	Description
No	No	Database can be accessed by multiple users.
Yes	Yes	Database has a status of single user.

Suspect Database Reports whether a database is marked as suspect. Use this attribute to warn you that a database cannot be loaded. The following values are valid:

External value	Internal value
True	0
False	1

Table Count The number of tables that exist in the database. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Transactions per Second The number of transactions started for the database per second for the current interval. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Write Transactions per sec The number of transactions written to the database and committed in the last second. The format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Database Mirroring attributes

The MS SQL Database Mirroring attribute group provides information about the database mirroring objects.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express® Edition.

Bytes Received Per Sec The number of bytes received per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Bytes Sent Per Sec The number of bytes sent per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Log Bytes Received Per Sec The number of bytes of log received per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Bytes Sent Per Sec The number of log bytes sent per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Log Compressed Bytes Rcvd Per Sec The number of compressed log bytes received in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Compressed Bytes Sent Per Sec The number of compressed log bytes sent in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Harden Time The number of milliseconds the log blocks waited to be hardened to disk in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Remaining For Undo 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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Scanned For Undo The total number of log bytes (in KB) that are scanned by the new mirror server after failover. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Send Flow Control Time The time (in milliseconds) awaited by the log stream messages for the send flow control in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Send Queue The total number of megabytes of log that are not sent to the mirror server. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Mirror Name The name of the database. The following values are valid:

External value	Internal value	Description
No Data	nodata	The value could not be collected.

Mirrored Write Transactions Per Sec 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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Mirroring Role The current role of the local database in the database mirroring session. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Mirror	0
Principal	1

Mirroring State The state of the mirror database. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Suspended	0
Disconnected	1

External value	Internal value
Synchronizing	2
Pending Failover	3
Synchronized	4

Originnode The key to this table of the format serverid:hostname:MSS.

Pages Sent Per Sec The number of pages sent per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Receives Per Sec The number of mirroring messages received per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Redo Bytes Per Sec The number of bytes of log redo by the mirror database per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Redo Queue The number of megabytes that redo on the mirror database is behind the hardened log. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Sample Timestamp The time when data is collected.

Send Receive Ack Time The time (in milliseconds) awaited by the messages to receive an acknowledgment from the partner. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Sends Per Sec The number of sends initiated per second at the time of monitoring interval. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Server The name of the SQL Server.

Synchronization State The state of synchronization (is failover possible). The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Failover not possible	0
Failover possible	1

Transaction Delay The time (in milliseconds) awaited by the transaction termination to receive an acknowledgment. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

MS SQL Database Summary attributes

The MS SQL Database Summary attribute group contains attributes that you can use to monitor summary information about SQL server databases.

Current Interval 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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day

YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters. For example: CFS_SVR5.

Total Databases The number of databases for the SQL server. The value format is an integer; for example, 10.

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.

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.

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.

Total DBs Read Only The number of databases with a status of read only. The value format is an integer; for example, 5.

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.

MS SQL Device Detail attributes

The MS SQL Device Detail attribute group contains attributes that you can use to monitor the usage and performance of devices.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Disabled	0	The autogrowth feature is disabled.
Restricted	1	The growth is restricted.
Unrestricted	2	The growth is unrestricted.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Not Applicable	-2	The value is not supported.
Percentage	1	The file growth is displayed in percentage.

External value	Internal value	Description
MB	2	The file growth is displayed in MB.

Autogrowth Value Indicates the growth value of a file. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Not Applicable	-2	The value is not supported.

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 following values are also valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Database Name (Unicode) The Database name. Valid values include text strings with a maximum of 384 bytes. The following values are also valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Device Free Space 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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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.

Device Name (Unicode) The Device name. Valid values include text strings with a maximum of 384 bytes.

Device Size 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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Option 1 for calculation :Autogrowth is enabled and has unrestricted growth This algorithm considers the hard disk drive freespace for calculation

Disk size = Total free space on all hard disk drives that contain the database files

Device Size = Disk size + Total database file size

Device Freespace = Device Size - Current allocated space of database file

Device Freespace Percent = (Device Freespace / Device Size) X 100

Option 2 for calculation: Autogrowth is enabled and has restricted growth This algorithm considers the maximum size of the database and hard disk drive free space for calculation.

Device maximum size = The restricted maximum size allocated to the database file

Disk size = Total free space on all hard disk drives that contain the database files

Condition (a): If device maximum size < Disk size

Device Size = Device maximum size

Device Freespace = Device Size - Current allocated space of database file

Device Freespace Percent = (Device Freespace / Device Size) X 100

Condition (b): if Device maximum size > Disk size

Device Size = Disk size

Device Freespace = Device Size - Current allocated space of database file

Device Freespace Percent = (Device Freespace /Device Size) X 100

Option 3 for calculation: Autogrowth disabled The database maximum size is the initial file size of the database at the time of database creation.

Device max size = Initial file size of the database at the time of database creation.

Device Size = Device maximum size

Device Freespace = Device Size - Current allocated space of database file

Device Freespace Percent = (Device Freespace / Device Size) X 100

Device Type Indicates the type of device allocated for the database. The following values are also valid:

External value	Internal value	Description
Log file	1	Indicates that the device stores the transaction log for the database.
Database file	2	Indicates that the device stores data for the database.
For future use	3	Indicates that the device supports future compatibility.

External value	Internal value	Description
Full text	4	Indicates that the device supports a full-text catalog.
FILESTREAM	5	Indicates that the device is a FILESTREAM device.
Not collected	-1	The value could not be collected.

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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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 following values are also valid:

External value	Internal value	Description
Not applicable	N/A	The value is not supported.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

Physical Device Name (Unicode) The Name of the physical device. Valid values include text strings with a maximum of 192 bytes.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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.

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.

Note: For the complete version of the SQL Server, see the Server Version attribute.

MS SQL Enterprise Database Mirroring attributes

The MS SQL Enterprise Database Mirroring attribute group provides information about the database mirroring objects.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

Bytes Received Per Sec The number of bytes received per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Bytes Sent Per Sec The number of bytes sent per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Log Bytes Received Per Sec The number of log bytes received per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Bytes Sent Per Sec The number of log bytes sent per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Log Compressed Bytes Rcvd Per Sec The number of compressed log bytes received in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Compressed Bytes Sent Per Sec The number of compressed log bytes sent in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Harden Time The time (in milliseconds) awaited by the log blocks to be hardened to the disk in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Remaining For Undo 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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Scanned For Undo The total number of log bytes (in KB) that are scanned by the new mirror server after failover. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Send Flow Control Time The time (in milliseconds) awaited by the log stream messenger for send flow control in the last second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Log Send Queue The total number of megabytes of log that are not sent to the mirror server. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Mirror Name The name of the database. The following values are valid:

External value	Internal value	Description
No Data	nodata	The value could not be collected.

Mirrored Write Transactions Per Sec 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. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Mirroring Role The current role of the local database in the database mirroring session. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Mirror	0
Principal	1

Mirroring State The state of the mirror database. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Suspended	0
Disconnected	1
Synchronizing	2
Pending Failover	3
Synchronized	4

Originnode The key to this table of the format serverid:hostname:MSS.

Pages Sent Per Sec The number of pages sent per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Receives Per Sec The number of mirroring messages received per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Redo Bytes Per Sec The number of bytes of log redone by the mirror database per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Redo Queue The number of megabytes that redo on the mirror database is behind the hardened log. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Sample Timestamp The time when data is collected.

Send Receive Ack Time The time (in milliseconds) awaited by the messages to receive an acknowledgment from the partner. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Note: This attribute is not supported on the SQL Server 2005 Enterprise Edition.

Sends Per Sec The number of sends initiated per second at the time of monitoring interval. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Server The name of the SQL Server.

Synchronization State The state of synchronization (is failover possible). The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Failover not possible	0
Failover possible	1

Transaction Delay The time (in milliseconds) awaited by the transaction termination to receive an acknowledgment. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

MS SQL Enterprise Log Shipping DB Details attributes

The MS SQL Enterprise Log Shipping DB Details attribute group provides detailed information about log shipping databases.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

Database ID The ID of the primary or secondary database for the log shipping configuration.

Database Name The name of the database.

Hub Timestamp The time when data is inserted at the hub.

Last File Processed The absolute path of the most recent transaction log backup or restore.

Last File Processed Time The date and time at which the file is processed. The format is YY/MM/DD HH:mm, where:

YY	Year
MM	Month
DD	Day
HH	Hour
mm	Minute

Message The job completion message text.

Originnode The key to this table of the format serverid:hostname:MSS.

Sample Timestamp The time when data is collected.

Server The name of the SQL Server.

Server Type Indicates whether the server is primary or secondary. The format is an integer. The following values are also valid:

External value	Internal value
Primary	0
Secondary	1

Source The source of the error message or event.

Threshold The number of minutes allowed to elapse between backup or restore operations before an alert is generated.

Threshold Exceeded Indicates whether the threshold is exceeded or not. The format is an integer. The following values are valid:

External value	Internal value
No	0
Yes	1

MS SQL Error Event Details attributes

The MS SQL Error Event Details attribute group 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 attribute group is an event-type attribute group.

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.

Error SPID The ID of the session in which the event has occurred. This attribute is available for reports, queries, and workspaces, but not for situations.

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 situations or for historical data collection.

Hub Timestamp The time at which data was inserted in the hub. This attribute is available for reports, queries, and workspaces, but not for situations.

Message Age 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.

Message Issuer The source of the error message. The value format is an alphanumeric string with a maximum of 8 characters, for example, JSMITH.

Message Text (Unicode) The message text. This attribute is available for reports, queries, and workspaces, but not for situations.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Originnode The key for the table view. The format is serverid:hostname:MSS. This attribute is available for reports, queries, and workspaces, but not for situations.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server. The value format is an alphanumeric string with a maximum of 32 characters. For example, CFS_SVR.

Severity Level The severity level of the error.

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.

MS SQL Filegroup Detail attributes

The MS SQL Filegroup Detail attribute group provides details about the filegroups for each database in the SQL Server instance.

Database Name The name of the database.

Filegroup ID The filegroup identification number.

Filegroup Max Growth Size The maximum amount of space (in MB) that the filegroup can grow to.

Filegroup Name The name of the filegroup.

Filegroup Size The size of the filegroup in MB.

Filegroup Space Percent Used The amount of space (in MB) used in the filegroup as a percentage of total space allowed.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day

YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server.

MS SQL FileTable Detail attributes

The MS SQL FileTable Detail attribute group contains attributes that provide information about each FileTable in the SQL Server.

Database ID The database ID. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

FileTable Directory The name of the directory of the FileTable. The following values are valid:

External value	Internal value	Description
No Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

FileTable Enabled Indicates whether the FileTable data is available for non-transactional access. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.
No	1	FileTable data is not enabled for non-transactional access.
Yes	2	FileTable data is enabled for non-transactional access.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Originnode The key for the table view in the format serverid:hostname:MSS.

Server The name of the SQL Server.

Table ID The table ID. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Table Name The table name. The value format is an alphanumeric string with a maximum of 30 characters. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value cannot be collected.
Not Applicable	Not_applicable	The value is not supported.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

MS SQL Individual Queries Details attributes

The MS SQL Individual Queries Details attribute group contains attributes that provide information about each query in the query batch that is currently running on the SQL Server.

Average CPU Time (Sec) The average CPU time (in seconds) that is taken to run an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Average Logical Reads The average number of logical read operations that are completed by an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Average Logical Writes The average number of logical write operations that are completed by an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Average Physical Reads The average number of physical read operations that are completed by an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Database Name (Unicode) The name of the database against which the SQL query is running on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	The value cannot be collected.

Hub Timestamp The time when data is inserted at the hub.

Last Execution Time The time when an SQL query was last run on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	The value cannot be collected.

The format is MM/DD/YY HH:mm:SS, where:

MM	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Max CPU Time (Sec) The maximum CPU time (in seconds) that is used for a single execution of an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Max Logical Reads The maximum number of logical read operations that are completed by a single execution of an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Max Logical Writes The maximum number of logical write operations that are completed by a single execution of an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Max Physical Reads The maximum number of physical read operations that are completed by a single execution of an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Max Waiting Time (Sec) The maximum waiting time (in seconds) that is taken by a single execution of an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Originnode The key for the table view in the format serverid:hostname:MSS.

Query Execution Count The total number of times an SQL query has been run since the query was last compiled. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Query Hash The query hash of an SQL query. The value format is an alphanumeric string with a maximum of 2048 characters. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

This attribute, with the Session ID attribute, is used to create the primary key.

Query Text The SQL text of a query in an SQL query batch. The following value is valid:

External value	Internal value	Description
No Data	nodata	The value cannot be collected.

Server The name of the SQL Server.

Session ID The ID of the session in which the query is run. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

SQL Handle The SQL handle of the query. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Total CPU Time (Sec) The total amount of CPU time (in seconds) that is used by all the executions of an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Percent CPU Used The percentage of the available CPU memory that is used for executing an SQL query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Total Percent Waiting The percentage of total wait operations that are completed when an SQL query is run. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Integration Service Details attributes

The MS SQL Integration Service Details attribute group provides details about the SQL Server integration service pipeline for an instance of the SQL Server.

Buffer In Use The number of buffers that are currently used for the pipeline. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Buffer Memory (MB) The amount of memory in MB that is currently allocated to the buffers that are in the pipeline. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Buffer Spooled The number of buffers that are currently spooled to disk. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Flat Buffers In Use The number of flat memory buffers that are currently used for the pipeline. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Flat Buffer Memory The amount of memory that is currently allocated to flat memory buffers. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, it is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Private Buffer Memory The amount of memory that is currently allocated to private transformation buffers. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Private Buffers In Use The number of private transformation buffers that are currently used for the pipeline. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Server The name of the SQL Server.

Total Rows Read The total number of rows that are currently being read from all data sources. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Rows Written The total number of rows that are currently being written to all data destinations. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Job Detail attributes

The MS SQL Job Detail attribute group contains attributes that you can use to monitor SQL Server jobs.

Current Execution Status The current execution status of the job. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Current Retry Attempt The current number of retry attempts that are done on a step of a running job. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Current Status The current job status. A job is in one of the following states:

External value	Internal value
Executing	1
Waiting_for_thread	2
Between_retries	3
Idle	4
Suspended	5
Performing_completion_actions	7
Unknown	-1

Enabled Whether or not the job is enabled to run. The following values are valid:

External value	Internal value
No	0
Yes	1
Unknown	-1

Job Type The type of the job. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Local Job	1	The job is run on a local server.
Multiserver	2	The job is run on multiple servers.
No Target Server	3	No target server is defined for the job.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Job Category ID Job category ID for the job. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Job Category Name Job category name for the job.

Job Description The description of the job. The value format is an alphanumeric string with a maximum of 512 characters. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Job Duration The amount of time it took for the job to complete (in seconds). The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Job Error Code Error code for the last completion of the job. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Job Execution Duration The time (in seconds) that has elapsed since the job started. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Not Applicable	-2	The value is not supported.

Jobs ID The SQL Server job ID.

Jobs Name The SQL Server job 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 following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Job Status Current status of the job. The following values are valid:

External value	Internal value
Failed	0
Succeeded	1
Retry	2
Canceled	3
In-progress	In-progress
Unknown	5 or -1

Last Run Outcome The last job execution status. The following values are valid:

External value	Internal value
Failed	0
Succeeded	1
Canceled	3
Unknown	5 or -1

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'.

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).

Notify Level Eventlog The value that indicates when a notification event must be sent to the Microsoft Windows application log. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Never	0	The notification must not be sent to the Microsoft application log.
When a job succeeds	1	The notification must be sent to the Microsoft application log when a job succeeds.
When the job fails	2	The notification must be sent to the Microsoft application log when a job fails.
When the job completes	3	The notification must be sent to the Microsoft application log when a job is complete, regardless of the outcome of the job.

Number of Steps The current number of steps in the job. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour

mm	Minute
SS	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.

Server The name of the SQL Server.

Start Step ID The step identifier from where the execution of the job starts. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Number The version number of the job. The version number is automatically updated every time that the job is modified. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Job Summary attributes

The MS SQL Job Summary attribute group contains attributes that you can use to monitor jobs.

Active Jobs The number of active jobs since the agent was started. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Failed Jobs The number of failed jobs since the SQL Server agent was started. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Jobs Failed Current Interval The number of failed jobs for the current sample interval. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Jobs Failed Since Startup The number of failed jobs since the agent was started.

Job Success Rate The rate at which jobs are successful since the agent was started. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Max Jobs Failed Interval Maximum number of failed jobs across all samples since the agent was started.

Max Running Jobs Interval Maximum number of running jobs across all samples since the agent was started.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Queued Jobs The number of jobs in a queue since the agent was started. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Running Jobs Current Interval The number of running jobs for the current sample interval. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server.

Successful Jobs The number of successful jobs since the agent was started. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Lock Conflict Detail attributes

The MS SQL Lock Conflict Detail attribute group contains attributes that you can use to monitor detailed information about a selected lock conflict.

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 situations 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.

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.

Client User ID ID of the user who executed the command. The value format is an alphanumeric string with a maximum of 8 characters.

Database ID ID of the database that is locked. The value format is an integer.

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.

Database Name (Unicode) The name of the database. Valid values include text strings with a maximum of 384 bytes.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The number of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

Server User ID The ID of the server user. The value format is an alphanumeric string with a maximum of 8 characters.

MS SQL Lock Detail attributes

The MS SQL Lock Detail attribute group contains attributes that you can use to monitor detailed information about lock contention by lock type.

Database ID The ID of the database that is locked. The value format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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.

Database Name (Unicode) The name of the locked database.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Lock Request Status Enumeration that identifies the Lock Request Status. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Granted	1
Converting	2
Waiting	3

Lock Resource Type E Enumeration that identifies the Lock Resource Type. The format is an integer. The following values are also valid:

External value	Internal value
NULL_Resource	1
Database	2
File	3
Index	4
Table	5
Page	6
Key	7
Extent	8

External value	Internal value
RID	9
Application	10
Metadata	11
Heap_or_Btree	12
Allocation_Unit	13

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Page Number The page number of the table being locked. The value format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Process Holding Lock The ID of the process holding the lock. The value format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example CFS_SVR5.

Table Name The name of the table being locked. The value format is an alphanumeric string with a maximum of 16 characters.

Table Name (Unicode) Name of Table being locked. Valid values include text strings with a maximum of 384 bytes.

Total Locks The total number of locks for the server. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Lock Conflicts The total number of processes that are involved in lock conflicts. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Type Indicates the type of lock on the resource that is being requested. The following values are also valid:

External value	Internal value	Description
No Data	nodata	No data is collected.
Unknown	Unknown	The value is unknown.

and

Value	Description
NULL	
Sch-S	Schema stability
Sch-M	Schema modification
IS	Intent Shared
SIU	Shared Intent Update
IS-S	Intent Shared-Shared
IX	Intent Exclusive
SIX	Shared Intent Exclusive
S	Shared
U	Update
IIn-Nul	Intent Insert-NULL
IS-X	Intent Shared-Exclusive
IU	Intent Update
IS-U	Intent Shared-Update
X	Exclusive
BU	Used by bulk operations
UIX	Update Intent Exclusive
RangeS-S	Shared Key-Range and Shared Resource
RangeS-U	Shared Key-Range and Update Resource
RangeI-N	Insert Key-Range and Null Resource
RangeI-S	Key-Range conversion lock
RangeI-U	Key-Range conversion lock
RangeI-X	Key-Range conversion lock
RangeX-S	Key-Range conversion lock
RangeX-U	Key-Range conversion lock

Value	Description
RangeX-X	Exclusive Key-Range and Exclusive Resource

MS SQL Lock Resource Type Summary attributes

The MS SQL Lock Resource Type Summary attribute group provides information about locks on individual resource types.

Average Wait Time Locks The average amount of wait time (in milliseconds) for each lock request that resulted in a wait.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

Lock Resource Type An enumeration of the resources that the SQL Server can lock. The following values are also valid:

External value	Internal value
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

Lock Timeouts per Second The number of lock requests per second that timed out, including requests for NOWAIT locks.

Lock Wait Time The total wait time (in milliseconds) for locks in the last second.

Number Deadlocks per Second The number of Deadlocks per second for the current sample interval.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server.

MS SQL Lock Summary attributes

The MS SQL Lock Summary attribute group contains attributes that you can use to monitor the number of database locks that are currently open.

Exclusive Extent Locks The number of granted Exclusive Extent locks.

Exclusive Intent Locks The number of granted Exclusive Intent locks.

Exclusive Locks The number of granted Exclusive locks.

Exclusive Page Locks The number of granted Exclusive Page locks.

Exclusive Table Locks The number of granted Exclusive Table locks.

Extent Locks The number of granted Extent locks.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Intent Locks The number of granted Intent locks.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Page Locks The number of granted Page locks.

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 situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server.

Shared Intent Locks The number of granted Shared Intent locks.

Shared Locks The number of granted Shared locks.

Shared Page Locks The number of granted Shared Page locks.

Shared Table Locks The number of granted Shared Table locks.

Table Locks The number of granted Table locks.

Table Lock Escalations per Sec The number of times that the locks on a table were escalated.

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 situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Update Extent Locks The number of granted Update Extent locks.

Update Page Locks The number of granted Update Page Locks.

MS SQL Log Shipping DB Details attributes

The MS SQL Log Shipping DB Details attribute group provides detailed information about log shipping databases.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

Database ID The ID of the primary or secondary database for the log shipping configuration.

Database Name The name of the database.

Hub Timestamp The time when data is inserted at the hub.

Last File Processed The absolute path of the most recent transaction log backup or restore.

Last File Processed Time The date and time at which the file is processed. The format is YY/MM/DD HH:mm, where:

YY	Year
MM	Month
DD	Day
HH	Hour
mm	Minute

Message The job completion message text.

Originnode The key to this table of the format serverid:hostname:MSS.

Sample Timestamp The time when data is collected.

Server The name of the SQL Server.

Server Type Indicates whether the server is primary or secondary. The format is an integer. The following values are also valid:

External value	Internal value
Primary	0
Secondary	1

Source The source of the error message or event.

Threshold The number of minutes allowed to elapse between backup or restore operations before an alert is generated.

Threshold Exceeded Indicates whether the threshold is exceeded or not. The format is an integer. The following values are valid:

External value	Internal value
No	0
Yes	1

MS SQL Log Shipping Errors attributes

The MS SQL Log Shipping Errors attribute group provides detailed information about log shipping errors.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

Agent Type The type of log shipping job. The following values are also valid:

External value	Internal value
Backup	0
Copy	1
Restore	2

Database ID The ID of the primary or secondary database for the log shipping configuration.

Database Name The name of the database.

Error Source The source of the error message or event.

Hub Timestamp The time when data is inserted at the hub.

Log Time The date and time at which the record is created. The format is YY/MM/DD HH:mm, where:

YY	Year
MM	Month
DD	Day
HH	Hour
mm	Minute

Message The job completion message text.

Originnode The key to this table of the format serverid:hostname:MSS.

Sample Timestamp The time when data is collected.

Sequence Number An incremental number that indicates the correct order of information for errors that span multiple records.

Server The name of the SQL Server.

MS SQL Log Shipping Jobs Detail attributes

The MS SQL Log Shipping Jobs Detail attribute group provides detailed information about log shipping jobs.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

Agent Type The type of log shipping job. The format is an integer. The following values are also valid:

External value	Internal value
Backup	0

External value	Internal value
Copy	1
Restore	2

Database ID The ID of the primary or secondary database for the log shipping configuration.

Note: For jobs that do not have an associated database, the value for the Database ID attribute is displayed as -1.

Database Name The name of the database.

Hub Timestamp The time when data is inserted at the hub.

Job Name The log shipping job name.

Job Status The status of the session. The following values are also valid:

External value	Internal value
Starting	0
Running	1
Success	2
Error	3
Warning	4

Log Time The date and time at which the record is created. The format is YY/MM/DD HH:mm, where:

YY	Year
MM	Month
DD	Day
HH	Hour
mm	Minute

Message The job completion message text.

Originnode The key to this table of the format serverid:hostname:MSS.

Sample Timestamp The time when data is collected.

Server The name of the SQL Server.

Session ID The primary ID for backup or the secondary ID for copy or restore.

MS SQL Memory Manager attributes

The MS SQL Memory Manager attribute group provides information about the memory usage.

Connection Memory The total amount of dynamic memory the server is using for maintaining connections in megabytes. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Granted Workspace Memory 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 following values are also valid:

External value	Internal value
Not Collected	-1

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Lock Blocks The current number of lock blocks that are in use on the server. Refreshed periodically. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Lock Blocks Allocated The current number of allocated lock blocks. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Lock Memory The total amount of dynamic memory the server is using for locks in megabytes. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

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 following values are also valid:

External value	Internal value
Not Collected	-1

Lock Owner Blocks Allocated The current number of allocated lock owner blocks. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Maximum Workspace Memory 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 following values are also valid:

External value	Internal value
Not Collected	-1

Memory Grants Outstanding The current number of processes that have successfully acquired a workspace memory. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Memory Grants Pending The current number of processes waiting for a workspace memory grant. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Optimizer Memory The total amount of dynamic memory the server is using for query optimization in megabytes. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Originnode The key to this table of the format serverid:hostname:MSS.

Sample Timestamp The time when data is collected.

Server The name of the SQL Server.

SQL Cache Memory The total amount of dynamic memory the server is using for the dynamic SQL cache in megabytes. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Target Server Memory The total amount of dynamic memory the server is willing to consume in megabytes. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Total Server Memory The total amount of dynamic memory the server is currently consuming in megabytes. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

MS SQL Problem Detail attributes

The MS SQL Problem Detail attribute group contains attributes that you can use to display detailed information about a selected error.

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.

Error SPID The ID of the session on which the event occurred. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Message Age 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.

Message Issuer The source of the error message. The value format is an alphanumeric string with a maximum of 8 characters; for example, JSMITH.

Message Text The message text. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Message Text (Unicode) The message text. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Message Timestamp The timestamp that indicates the date and time the error occurred. The format is MM/DD/YY HH:mm:SS, where:

MM	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR.

Severity Level Indicates the severity level of the error.

SQL State Code The SQL state value for the error message. The value format is an integer; for example, 37.

MS SQL Problem Summary attributes

The MS SQL Problem Summary attribute group contains attributes that you can use to monitor summary information about SQL server errors.

Age of Last Error The number of minutes that have elapsed since the last error message occurred. The value format is an integer; for example, 2.

Current Interval 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.

Error Log Size The number of bytes in the error log file. The value format is an integer; for example, 50000.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

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	Month
DD	Day
YY	Year

HH	Hour
mm	Minute
SS	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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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.

Total Errors Current Interval The number of error messages that occurred during the current interval. The value format is an integer; for example, 1.

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.

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.

MS SQL Process Detail attributes

The MS SQL Process Detail attribute group contains attributes that you can use to monitor detailed information about a selected SQL server process.

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 situations 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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Command (Unicode) The name of the executing command. The value format is an alphanumeric string with a maximum of 48 bytes. The following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Current CPU Percent 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 following value is valid:

External Value	Internal Value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Database Name (Unicode) The database name. The value format is an alphanumeric string with a maximum of 384 bytes. The following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Duration How long, in seconds, the process has been running. The following value is valid:

External Value	Internal Value	Description
Not Collected	-1	The value could not be collected.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Login Name ID used by the process to log into the SQL Server. The value format is an alphanumeric string with a maximum of 30 characters. The following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

NT User User name associated with the process. The following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

OS Process ID The Microsoft Windows thread ID. The following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Process ID The ID of the process. The value format is an alphanumeric string with a maximum of 10 characters; for example, 42168. The following value is valid:

External Value	Internal Value	Description
Not Collected	-1	The value could not be collected.

Process Status Indicates the status of the process. The following values are valid:

Value	Description
Alarm sleep	Process is waiting for an alarm.
Background	Process is a Microsoft SQL Server process.
Bad status	Process has errors.
Infected	Process is infected.
Lock sleep	Process is waiting for a lock to be released.
Log suspend	Process is suspended by the log transaction.
Recv sleep	Process is waiting for a network read.
Runnable	Process is in the queue.
Running	Process is running.
Sleeping	Process is sleeping.
Stopped	Process is stopped.
Send sleep	Process is waiting on a network send.
Dormant	Process is resetting.
Not Collected	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Program Name (Unicode) The Program Name. The value format is an alphanumeric string with a maximum of 384 bytes. The following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	nodata	The value could not be collected.

Total CPU Time 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 following value is valid:

External Value	Internal Value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External Value	Internal Value	Description
Not Collected	-1	The value could not be collected.

MS SQL Process Summary attributes

The MS SQL Process Summary attribute group contains attributes that you can use to monitor summary information about processes.

Current Interval 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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

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.

Percent Processes Locksleep The percentage of processes with a status of locksleeper. 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.

Percent Processes Othersleep The percentage of processes with a status of alarm sleep, rcv sleep, send sleep, or sleeping. The value format is a percentage with two decimal places allowed; for example, 10.00.

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.

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.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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.

Server CPU Percent 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.

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.

Total Other Sleep The number of processes with a status of othersleep. The value format is an integer; for example, 2.

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.

Total Processes Bad The number of processes with a status of bad. The value format is an integer; for example, 1.

Total Processes Blocked The number of processes that are being blocked. The value format is an integer; for example, 2.

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.

Total Processes Locksleep The number of processes with a status of locksleepp. A process with a status of locksleepp is a process waiting for a lock on a resource to be released. The value format is an integer; for example, 5.

Total Processes Stopped The number of processes with a status of stopped. The value format is an integer; for example, 1.

MS SQL Remote Servers attributes

The MS SQL Remote Servers attribute group contains attributes that you can use to monitor summary information about remote SQL servers.

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.

Current Interval 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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

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 following values are valid:

External value	Internal value	Description
Active	Active	SQL server is active.
Inactive	Inactive	SQL server status is inactive.
Unknown	Unknown	SQL server status is unknown.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

MS SQL Resource Pool Stats attributes

The MS SQL Resource Pool Stats attribute group provides statistics associated with resource pools.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

Active Memory[™] Grant Amount The total amount of granted memory in megabytes in the resource pool. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Active Memory Grants Count The count of memory grants that are assigned to pool workers. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Cache Memory Target The current memory target for cache memory in megabytes. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Compile Memory Target The current memory target for compile memory in megabytes. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

CPU Control Effect Percent 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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

CPU Usage Percent The system CPU Usage by all requests in the specified instance of the performance object. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

CPU Usage Target Percent 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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Max Memory 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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Memory Grants Per Sec The number of query memory grants per second occurring in the resource pool. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key to this table of the format serverid:hostname:MSS.

Pending Memory Grants Count The number of queries waiting for memory grants in the resource pool. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Query Exec Memory Target The current memory target for query execution in megabytes. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Resource Pool Name The resource governor resource pool name. The following values are valid:

External value	Internal value	Description
No Data	nodata	The value could not be collected.

Sample Timestamp The time when data is collected.

Server The name of the SQL Server.

Target Memory 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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Used Memory The used amount of memory in megabytes in the resource pool. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Running Queries Details attributes

The MS SQL Running Queries Details attribute group contains attributes that provide information about the queries that are currently running on the SQL Server.

Blocking Session ID The ID of the session that is currently blocking the execution of a query on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Command Type The type of command that is currently being processed by the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

CPU Time (Sec) The CPU time (in seconds) of the SQL query that is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Database Name (Unicode) The name of the database against which the SQL query is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Granted Query Memory The number of pages that are currently allocated to the execution of a query. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Hub Timestamp The time when data is inserted at the hub.

Login Name The name that is used to log on to the SQL Server on which the query is currently running. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Originnode The key for the table view in the format serverid:hostname:MSS.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Program Name The name of the client program that initiated the request, which is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Query Start Time The time when the query arrived on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

The format is MM/DD/YY HH:mm:ss, where:

MM	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Query Status The status of the SQL query that is currently running on the SQL Server. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Background	1	The query is running at the background.
Running	2	The query is currently running.
Runnable	3	The query can be run.
Sleeping	4	The query is currently inactive.

External value	Internal value	Description
Suspended	5	The query is currently suspended.

Query Reads The number of query read operations completed by the SQL query that is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Query Writes The number of query write operations completed by the SQL query that is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

Row Count The number of rows returned to the client by the SQL query that is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

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 following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Server The name of the SQL Server.

Session ID The session ID of the query that is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

SQL Handle The SQL handle of the query that is currently running on the SQL Server. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

Total Elapsed Time (Sec) The total time (in seconds) that has elapsed since the SQL query arrived on the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

MS SQL Service Broker Activation attributes

The MS SQL Service Broker Activation attribute group monitors the details of the Service Broker activation for each database in an SQL Server instance.

Total Tasks Started The total number of activated tasks that are currently started. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Database Name (Unicode) The name of the database in the Unicode format.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Server The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters.

Stored Procedure Invoked Per Sec The number of stored procedures that are currently called per second.

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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.

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Tasks Aborted Per Sec The number of activated tasks that are currently ended per second. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Tasks Started Per Sec The number of activated tasks that are currently being started per second. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Task Limit Reached The number of times that the task limit that is currently activated for a queue has been reached. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Service Broker Statistics attributes

The MS SQL Service Broker Statistics attribute group displays the summary of the Service Broker statistics.

Broker Transaction Rollbacks The current number of transactions related to the Service Broker that have been rolled back. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Enqueued Local Messages Per Sec The number of messages per second from local endpoints that are currently being delivered into the queues. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Enqueued Transport Messages Per Sec The number of messages from the transport that are being delivered into the queues per second. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Enqueued Messages Per Sec The number of messages from local endpoints and the transport that are currently delivered into the queues per second.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Forwarded Pending Messages The number of forwarded messages that are currently in the queue, and that are not yet sent.

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Server The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters.

SQL Receives Per Sec The number of SQL RECEIVE commands that are processed per second by the Service Broker. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

SQL Sends Per Sec The number of SQL SEND commands that are processed per second by the Service Broker. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Enqueued Local Messages The total number of message fragments from the local endpoints that are currently being delivered into the queues. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Enqueued Transport Messages The total number of messages from the transport that are currently delivered into the queues. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Forwarded Messages The total number of forwarded messages from the local endpoints and the transport that are currently being sent. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total SQL Receives The total number of SQL RECEIVE commands that are currently processed by the Service Broker. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total SQL Sends The total number of SQL SEND commands that are currently processed by the Service Broker. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Service Broker Transport attributes

The MS SQL Service Broker Transport attribute group monitors the details of the Service Broker transport for the SQL Server.

Average Size Message Fragments Received The average byte size of the message fragments that are received by the transport receive I/O operations. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Average Size Message Fragments Sent The average byte size of the message fragments that are sent by the transport send I/O operations. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Message Fragment Received Per Sec The number of message fragments that are currently received per second in the transport receive I/O operations. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Message Fragments Sent Per Sec The number of message fragments that are currently sent per second in the transport send I/O operations. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Server The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters.

Total Open Connection Count The total number of transport connections that are currently open. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Services Detail attributes

The MS SQL Services Detail attribute group provides details about the services that make up the SQL Server.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server.

Service Name The service name for the SQL Server.

Service State Current service state. The following values are also valid:

External value	Internal value
Stopped	0
Start Pending	1
Stop Pending	2
Running	3
Continue Pending	4
Pause Pending	5
Paused	6
Unknown	-1

Service Start Mode Defined start mode for the service. The following values are also valid:

External value	Internal value
Automatic	1
Manual	2
Disabled	3
Unknown	-1

Service Status Current service status. The following values are also valid:

External value	Internal value
OK	0
Error	1
Degraded	2
Unknown	3
Pred Fail	4
Starting	5
Stopping	6
Service	7

Service Type The service type for this service to the Microsoft SQL Server. The following values are also valid:

External value	Internal value
SQLServer	1
SQLAgent	2
DTC	3
Analysis_OLAP	4
Search	5
ADHelper	6
Browser	7
DTS	8
SQLWriter	9

External value	Internal value
Report	10
Distributed Replay Controller	11
Distributed Replay Client	12
Unknown	-1

MS SQL Server Detail attributes

The MS SQL Server Detail attribute group contains attributes that you can use to monitor detailed information about a selected SQL server.

Current Locks The number of current locks for the SQL server. The value format is an integer; for example, 73.

Data Cache Size 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.

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 situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Error Log Name (Unicode) The name of the file that contains the error log for the SQL server (Unicode). Valid values include text strings with a maximum of 392 characters. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Error Log Size 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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Max Locks Allowed The maximum number of allowable locks. The value format is an integer; for example, 2000. This value is a configuration parameter.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

OS Type The operating system for the SQL server. The value format is an alphanumeric string; for example, AIX. The following values are valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

External value	Internal value	Description
Windows Srv 2008	Win2008	The operating system is Windows Server 2008.
Windows Srv 2003	Win2003	The operating system is Windows Server 2003.
Windows Srv 2000	Win2000	The operating system is Windows Server 2000.
Windows NT	WinNT	The operating system is Windows NT.
Windows Srv 2012	Win2012	The operating system is Windows Server 2012.
Windows Srv 2008 R2	Win8R2	The operating system is Windows Server 2008 R2.

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.

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.

Procedure Buffers Percent 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.

Procedure Buffers Percent 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.

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.

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.

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.

Note: This attribute is deprecated from the SQL Server agent V6.2.1.1.

Procedure Cache Percent 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.

Procedure Cache Percent 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.

Procedure Cache Size 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.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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 following values are valid:

External value	Internal value	Description
Active	Active or 1	SQL server is active.
Inactive	Inactive	SQL server status is inactive.
Unknown	Unknown or 0	SQL server status is unknown.

Server Type The type of SQL server. The value format is an alphanumeric string; for example, SQL Server.

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.

Server Version (Supceded) The version of the SQL Server. The value format is the version in the format x.y.z; for example, 10.0.253.

Note: For the complete version of the SQL Server, see the Server Version attribute.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Time Since Startup The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360.

MS SQL Server Properties attributes

The MS SQL Server Properties attribute group contains attributes that you can use to monitor the properties of the SQL Server.

Collation The name of the default collation for the server. The value format is a string. The following value is valid:

External value	Internal value	Description
Not Collected	Not_collected	The value could not be collected.

Collation ID The identification number of this SQL Server collation. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Edition The installed product edition of this instance of the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	Not_collected	The value could not be collected.

Edition ID The identification number that represents the installed product edition of this instance of the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Engine Edition The database engine edition of the instance of the SQL Server that is currently installed. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

FileStream Level The current level of FileStream support that is enabled for the SQL Server instance. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Not Applicable	-2	The value is not supported.
Disabled	1	FileStream support is disabled.
Transactional Access	2	FileStream support is enabled for transactional access.
Local Full Access	3	FileStream support is enabled only for local full access.

External value	Internal value	Description
Local and Remote Full Access	4	FileStream support is enabled for local and remote full access.

FileStream Share Name The Windows share name where the FileStream data is stored. The value format is an alphanumeric string. The following values are valid:

External value	Internal value	Description
Not Collected	Not_collected	The value could not be collected.
Not Applicable	Not_applicable	The value is not supported.

HADR Manager Status Indicates whether the HADR manager for the AlwaysOn availability group has started. The value format is an integer. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Not Applicable	-2	The value is not supported.
No Data	-3	No data is collected.
Communication Pending	1	The communication with the HADR manager is pending, and the HADR manager has not started.
Running	1	The HADR manager is running.
Failed	2	The HADR manager has failed.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
No	1	The SQL Server instance is not configured in a failover cluster.
Yes	2	The SQL Server instance is configured in a failover cluster.

Is HADR Enabled Indicates whether the AlwaysOn Availability Group feature is enabled on the SQL Server instance. The value format is an integer. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
No Applicable	-2	The attribute is not supported.
No Data	-3	No data is collected.
No	1	The AlwaysOn Availability Group feature is not enabled on the SQL Server instance.

External value	Internal value	Description
Yes	2	The AlwaysOn Availability Group feature is enabled on the SQL Server instance.

Is Integrated Security Only The server is in the integrated security mode. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
No	1	The SQL Server is not in the integrated security mode.
Yes	2	The SQL Server is in the integrated security mode.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
No	1	The SQL Server is not in the single-user mode.
Yes	2	The SQL Server is in the single-user mode.

License Type The mode of this instance of the SQL Server. The mode can be per-seat or per-processor. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Physical NetBIOS Name The NetBIOS name of the machine where this instance of the SQL Server is currently running. The following value is valid:

External value	Internal value	Description
Not Collected	Not_collected	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
No Data	nodata	The value could not be collected.
CTP	1	The level of the version of the SQL Server is CTP.
RTM	2	The level of the version of the SQL Server is RTM.
SP1	3	The level of the version of the SQL Server is SP1.
SP2	4	The level of the version of the SQL Server is SP2.
SP3	5	The level of the version of the SQL Server is SP3.
SP4	6	The level of the version of the SQL Server is SP4.
SP5	7	The level of the version of the SQL Server is SP5.
SP6	8	The level of the version of the SQL Server is SP6.
SP7	9	The level of the version of the SQL Server is SP7.
SP8	10	The level of the version of the SQL Server is SP8.
SP9	11	The level of the version of the SQL Server is SP9.
SP10	12	The level of the version of the SQL Server is SP10.

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 following value is valid:

External value	Internal value	Description
Not Collected	Not_collected	The value could not be collected.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters.

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.

MS SQL Server Summary attributes

The MS SQL Server Summary attribute group contains attributes that you can use to monitor summary information for SQL servers.

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.

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 following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

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 95% or greater. The value format is an integer.

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.

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.

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 following values are also valid:

External value	Internal value	Description
Inactive	Inactive	Data collector is inactive.
Unknown	Unknown	Data collector is unknown.
Active	Active	Data collector is active.

CPU Idle Delta The difference in the CPU Idle between the current and the previous samples.

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.

CPU Percent 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.

CPU Usage Delta The difference in the CPU Usage between the current and the previous samples.

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.

Current Interval 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.

Data Cache Size 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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Page Life Expectancy The duration (in seconds) for which an SQL Server block or page is stored in the memory.

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.

Page Writes per Sec The number of physical database page writes issued.

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.

Procedure Cache Size 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.

Readahead Pages per Sec The number of pages read in anticipation of use.

Repl Distribution Delivered Cmds per Sec The number of distribution commands delivered per second to the Subscriber.

Repl Distribution Delivered Trans per Sec The number of distribution transactions delivered per second to the Subscriber.

Repl Distribution Delivery Latency The distribution latency (in MS), the time it takes for transactions to be delivered to the Distributor and applied at the Subscriber.

Repl Logreader Delivered Cmds per Sec The number of Log Reader Agent commands delivered per second to the Distributor.

Repl Logreader Delivered Trans per Sec The number of Log Reader Agent transactions delivered per second to the Distributor.

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.

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.

Repl Merge Downloaded Changes per Sec The number of rows merged (inserted, updated, and deleted) per second from the Publisher to the Subscriber.

Repl Merge Uploaded Changes per Sec The number of rows merged (inserted, updated, and deleted) per second from the Subscriber to the Publisher.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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.

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.

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 following values are also valid:

External value	Internal value	Description
Active	Active or 1	SQL server is active.
Inactive	Inactive	SQL server status is inactive.
Unknown	Unknown or 0	SQL server status is unknown.

Server Type The type of SQL server. The value format is an alphanumeric string; for example, SQL Server.

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.

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.

Note: For the complete version of the SQL Server, see the Server Version attribute.

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.

Stolen Pages The number of pages used for miscellaneous server purposes. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Stolen Pages Growth The growth of the number of stolen pages between the current sample and previous sample. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Time Since Startup The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360.

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.

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.

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.

Total Pages The total number of stolen, free, and database pages in the buffer pool. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.
Not Applicable	-2	The value is not supported.

Total Server Memory The total amount of dynamic memory (in KB) that the server is using currently.

MS SQL Server Transactions Summary attributes

The MS SQL Server Transactions Summary attribute group contains attributes that you can use to monitor the transaction summary information for SQL Servers.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Longest Transactions Running Time The transaction that is currently active for the longest duration of time (in seconds). The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, this attribute is available for reports, queries, and workspaces.

Server The name of the SQL Server. The value is an alphanumeric string with a maximum of 30 characters.

Tempdb Free Space The amount of free space (in KB) that is currently available in the tempdb database. The value format is an integer. The following value is valid.

External value	Internal value
Not Collected	-1

Total Active Transactions The number of transaction enlistments (local, dtc, and bound) that are currently active. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Non-Snapshot Version Transactions The number of currently active non-snapshot transactions that generate version records. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Snapshot Transactions The number of snapshot transactions that are currently active. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Update Snapshot Transactions The number of currently active snapshot transactions that perform update operations. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Update Conflict Ratio The percentage of snapshot transactions that currently have update conflicts. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Cleanup Rate The current rate (in KB per second) at which versions are removed from the version store. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Generation Rate The current rate (in KB per second) at which versions are generated in the version store. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Store Size The amount of space (in KB) that is currently available in the version store. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Store Unit Count The number of allocation units that are currently active in the version store. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Store Unit Creation The number of units that are currently created in the version store. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Version Store Unit Truncation The number of units that are currently truncated in the version store. The value format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Statistics Detail attributes

The MS SQL Statistics Detail attribute group contains attributes that you can use to monitor detailed information about SQL server statistics.

Average Value 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.

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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

Statistic Name The name of the statistic. The value format is an alphanumeric string with a maximum of 32 characters; for example, cpu_busy. The following values are also valid:

External value	Internal value
CPU_busy	cpu_busy
IO_busy	io_busy
I/O_errors	io errors

External value	Internal value
connections	connections
idle	idle
packets_error	packets_error
packets_received	packets_received
packets_sent	packets_sent
physical_reads	physical reads
physical_writes	physical writes

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.

MS SQL Statistics Summary attributes

The MS SQL Statistics Summary attribute group contains attributes that you can use to monitor summary information about SQL server statistics.

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.

Client Count Percent 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.

Current Interval 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.

Current Logons The number of active connections (logons). The value format is an integer; for example, 50.

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.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

IO ErrorsCurrent 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.

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.

Max User Connection 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.

Network Read Rate 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.

Network Write Rate 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.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Percent IO Errors Current Interval 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.

Percent 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.

Percent 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.

Physical Reads per Second 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.

Physical Writes per Second 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.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

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 following value is possible:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Total Logons per Second 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.

Total Logouts per Second Total number of logout operations started per second.

Total OS CPU Percent 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.

Total OS IO Percent 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.

MS SQL Table Detail attributes

The MS SQL Table Detail attribute group contains attributes that you can use to monitor tables within one or more databases.

Database ID ID number of this database. The format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Database Name The Database name. The value format is an alphanumeric string with a maximum of 30 characters. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Index ID The ID number of this table index. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Index Name The table index name. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

FileTable Indicates whether the current table is a FileTable. The format is an integer. The following values are valid:

External value	Internal value	Description
Not Applicable	-2	The value is not supported.
Not Collected	-1	The value could not be collected.
No	1	The current table is not a FileTable.
Yes	2	The current table is a FileTable.

Number of Rows The number of rows that are present in a database table. The format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Optimizer Statistics Age The time (in minutes) since statistics were 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 format is an integer. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL Server. The value format is an alphanumeric string with a maximum of 32 characters

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Table ID ID number of this table. The format is an integer. The following value is also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Table Name The table name. The value format is an alphanumeric string with a maximum of 30 characters. The following value is valid:

External value	Internal value	Description
No Data	nodata	No data is collected.

MS SQL Table Summary attributes

The MS SQL Table Summary attribute group 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 attribute group display the last collected value of data.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute. For example, it is available for reports, queries, and workspaces.

Last Collection Duration The time (in minutes) that is required to collect data of the MS SQL Table Detail attribute group. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Last Collection Start Time The date and time when data collection of the MS SQL Table Detail attribute group last started. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Maximum Fragmentation Percent Since Startup The highest value of fragmentation (in a percent) among all tables since the agent was started. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Maximum Fragmentation Percent The highest fragmentation value (in a percent) among all the tables in the current interval. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Minimum Table Size (MB) The minimum space in MB that is used among all the tables of the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Minimum Fragmentation Percent The lowest value of fragmentation (in a percent) among all the tables in the current interval. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Server The name of the SQLServer. The value format is an alphanumeric string with a maximum of 32 characters

Total FileTables The total number of FileTables that are present in the current instance of the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.
Not Applicable	-2	The value is not supported.

Total Tables The total number of tables that are present in the current instance of the SQL Server. The following value is valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

MS SQL Text attributes

The MS SQL Text attribute group contains attributes that you can use to monitor information about SQL text strings associated with a selected process.

Note: For custom queries, this attribute group displays data only for the specified available process ID.

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.

Client Process ID The process ID assigned by the host client. The value format is an alphanumeric string with a maximum of 16 characters.

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.

Database Name (Unicode) The Database name. Valid values include text strings with a maximum of 384 bytes.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.

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.

Process Status Indicates the status of the process. The following values are also valid:

Value	Description
Alarm sleep	Process is waiting for an alarm.
Background	Process is a Microsoft SQL Server process.
Bad status	Process has errors.
Infected	Process is infected.
Lock sleep	Process is waiting for a lock to be released.
Log suspend	Process is suspended by the log transaction.
Recv sleep	Process is waiting for a network read.
Runnable	Process is in the queue.

Value	Description
Running	Process is running.
Sleeping	Process is sleeping.
Stopped	Process is stopped.
Send sleep	Process is waiting on a network send.
Dormant	Process is resetting.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5.

Text The SQL text for the process. The value format is an alphanumeric string with a maximum of 255 characters.

Text (Unicode) The SQL text for the process.

Total CPU Time 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.

MS SQL Workload Group Stats attributes

The MS SQL Workload Group Stats attribute group provides statistics associated with workload groups.

Note: This attribute group is not supported on the SQL Server 2005 and 2008 Express Edition.

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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Active Requests The number of requests currently running in the workload group. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Blocked Tasks The number of blocked tasks in the workload group. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

CPU Usage Percent The system CPU usage by all requests in the specified instance of the performance object. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 situations or for historical data collection.

Hub Timestamp The time when data is inserted at the hub.

Max Requests CPU Time The maximum requests for CPU time in milliseconds. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Max Requests Memory Grant The maximum value of memory grant in megabytes used by a query in the workload group. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Originnode The key to this table of the format serverid:hostname:MSS.

Query Optimization Per Sec The number of query optimizations per second occurring in the workload group. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

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 following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Requests Completed Per Sec The number of completed requests per second in the workload group. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Sample Timestamp The time when data is collected. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Server The name of the SQL Server.

Suboptimal Plans Per Sec The number of suboptimal query plans generated per second in the workload group. The format is an integer. The following values are also valid:

External value	Internal value	Description
Not Collected	-1	The value could not be collected.

Workload Group Name The name of the workload group. The following values are valid:

External value	Internal value	Description
No Data	nodata	The value could not be collected.

VMS SQL Remote Servers attributes

The VMS SQL Remote Servers attribute group contains attributes that you can use to monitor summary information about remote SQL servers.

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. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

Current Interval 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. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group. The following values are valid:

External value	Internal value	Description
Not Collected	-1	The value cannot be collected.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

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. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

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. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group. To create a situation, use the MS SQL Remote Servers attribute group.

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. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

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 following values are valid:

External value	Internal value	Description
Active	Active	SQL server is active.
Inactive	Inactive	SQL server status is inactive.
Unknown	Unknown	SQL server status is unknown.

This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

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	Month
DD	Day
YY	Year
HH	Hour

mm	Minute
SS	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 attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Remote Servers attribute group.

VMS SQL Server Enterprise View attributes

The VMS SQL Server Enterprise View attribute group contains attributes that you can use to monitor the SQL servers in your enterprise.

Age of Last Error The age in minutes of the last error encountered. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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 following values are also valid:

External value	Internal value	Description
Unknown	Unknown	Data collector is unknown.
Inactive	Inactive	Data collector is inactive.
Active	Active	Data collector is active.

This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Current Logons The number of logons and connections that are currently active. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Database Max Blocks The name of the database blocking the largest number of processes. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Database Max Blocks (Unicode) The name of the database blocking the largest number of processes.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Database Max Locks (Unicode) The name of the database with largest number of locks. Valid values include text strings with a maximum of 32 bytes. This attribute is not available for use in situations or for

historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Error Log Size The size in bytes of the error log. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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 situations or for historical data collection.

Hub Timestamp The time when this data was inserted at the hub. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

IO Errors Current Interval The number of disk errors in the current interval. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Maximum Sev Level The highest severity level of an error message since startup. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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 attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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. This attribute is not available for use in situations or for historical

data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Mirror Name The name of the database. The following values are valid:

External value	Internal value
No Data	nodata

Mirroring Role The current role of the local database in the database mirroring session. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Mirror	0
Principal	1

Mirroring State The state of the mirror database. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Suspended	0
Disconnected	1
Synchronizing	2
Pending Failover	3
Synchronized	4

Originnode The key for the table view in the format serverid:hostname:MSS. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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 situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Percent 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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Percent Process Block The percentage of total processes in conflict. The value format is a decimal number in the 0.00 through 100.00. This attribute is not available for use in situations or for historical data

collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Physical Reads per Second The number of physical reads per second during the current interval. The value format is a decimal number with 2 decimal places. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Physical Writes per Second The Number of physical writes per second during the current interval. The value is a decimal number with 2 decimal places. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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	Month
DD	Day
YY	Year
HH	Hour
mm	Minute
SS	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 attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Server The name of the SQL server. The value format is an alphanumeric string with a maximum of 30 characters; for example, CFS_SVR5. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Server CPU Percent System The percentage of CPU being used by system processes for the server. The value format is a decimal number in the range 0.00 through 100.00. This attribute is not available for use

in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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 following values are also valid:

External value	Internal value	Description
Active	Active or 1	SQL server is active.
Inactive	Inactive	SQL server is not active.
Unknown	Unknown or 0	SQL server status is unknown.

This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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.

Server Version (Superseded) The version of the SQL Server. The value format is the version in the format x.y.z; for example, 10.0.253. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Note: For the complete version of the SQL Server, see the Server Version attribute.

Synchronization State The synchronization state (is failover possible). The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1
Failover not possible	0
Failover possible	1

Table Max Locks The name of the table with largest number of locks. The value format is an alphanumeric string with a maximum of 32 characters.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Time Since Startup The number of minutes that have elapsed since the SQL server was started. The value format is an integer; for example, 360. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Total Databases The number of databases for the SQL server. The value format is an integer; for example, 10. This attribute is not available for use in situations or for historical data collection. Otherwise, this

attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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 situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Total Locks The total number of locks for the server. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

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. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Total Processes The total number of processes. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Total Processes Blocked The total number of processes blocked. The value format is an integer. This attribute is not available for use in situations or for historical data collection. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces. To create a situation, use the MS SQL Server Enterprise View attribute group.

Transaction Delay The time (in milliseconds) awaited by the transaction termination to receive an acknowledgment per second. The format is an integer. The following values are also valid:

External value	Internal value
Not Collected	-1

Disk capacity planning for historical data

Disk capacity planning for a monitoring agent is a prediction of the amount of disk space to be consumed for each attribute group whose historical data is being collected. Required disk storage is an important factor to consider when you are defining data collection rules and your strategy for historical data collection.

Calculate expected disk space consumption by multiplying the number of bytes per instance by the expected number of instances, and then multiplying that product by the number of samples. Table 8 on page 223 provides the following information required to calculate disk space for the Microsoft SQL Server agent:

- *DB table name* is the table name as it would appear in the warehouse database, if the attribute group is configured to be written to the warehouse.

- *Bytes per instance (agent)* is an estimate of the record length for each row or instance written to the agent disk for historical data collection. This estimate can be used for agent disk space planning purposes.
- *Bytes per instance (warehouse)* is an estimate of the record length for detailed records written to the warehouse database, if the attribute group is configured to be written to the warehouse. Detailed records are those that have been uploaded from the agent for long-term historical data collection. This estimate can be used for warehouse disk space planning purposes.
- *Bytes per summarized instance (warehouse)* is an estimate of the record length for aggregate records written to the warehouse database, if the attribute group is configured to be written to the warehouse. Aggregate records are created by the Summarization agent for attribute groups that have been configured for summarization. This estimate can be used for warehouse disk space planning purposes.
- *Expected number of instances* is a guideline that can be different for each attribute group, because it is the number of instances of data that the agent will return for a given attribute group, and depends upon the application environment that is being monitored. For example, if your attribute group is monitoring each processor on your machine and you have a dual processor machine, the number of instances is 2.

The *IBM Tivoli Monitoring Installation and Setup Guide* contains formulas that can be used to estimate the amount of disk space used at the agent and in the warehouse database for historical data collection of an attribute group.

Table 8. Capacity planning for historical data logged by the Microsoft SQL Server agent

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KOQAADBD	MS_SQL_Availability_Database_Details	832	834	1086
KOQADDBD	MS_SQL_Availability_Database_Detail	1448	1451	1488
KOQSASDM	MS_SQL_Analysis_Services_Data_Mining	220	250	658
KOQSASMX	MS_SQL_Analysis_Services_MDX_Processing	364	380	1178
KOQSASSM	MS_SQL_Analysis_Services_Memory_Statistics	676	899	3926
KOQSASPR	MS_SQL_Analysis_Services_Rows_Processing	196	215	486
KOQSASSQ	MS_SQL_Analysis_Services_Storage_Engine_Query	516	815	2950
KOQAVDBD	MS_SQL_Availability_Database_Details	1332	1342	1379
KOQCAVDB	MS_SQL_Availability_Database_Details_In_Cluster	1312	1317	1354
KOQADBST	MS_SQL_Availability_Database_Statistics	612	688	1227
KOQADBSU	MS_SQL_Availability_Database_Summary	208	212	522
KOQAVGLD	MS_SQL_Availability_Group_Listener_Details	1188	1191	1228
KOQGRPDT	MS_SQL_Availability_Groups_Details	1272	1283	1398
KOQCAVGD	MS_SQL_Availability_Groups_Details_In_Cluster	928	928	965

Table 8. Capacity planning for historical data logged by the Microsoft SQL Server agent (continued)

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KOQGRPSM	MS_SQL_Availability_Groups_Summary	248	247	479
KOQAVARD	MS_SQL_Availability_Replicas_Details	1232	1240	1277
KOQCAVRD	MS_SQL_Availability_Replicas_Details_In_Cluster	1560	1567	1604
KOQAVRST	MS_SQL_Availability_Replicas_Statistics	284	355	851
KOQAVARS	MS_SQL_Availability_Replicas_Status	820	824	900
KOQAVRSU	MS_SQL_Availability_Replicas_Status_Summary	208	212	522
KOQBTCHS	MS_SQL_Batch_Stats	248	355	851
KOQBLKQD	MS_SQL_Blocked_Queries_Details	7524	7548	7628
KOQSCFG	MS_SQL_Configuration	313	310	347
KOQDBD	MS_SQL_Database_Detail	2735	3028	4334
KOQDBMIR	MS_SQL_Database_Mirroring	700	980	2088
KOQDBS	MS_SQL_Database_Summary	266	288	559
KOQDEVD	MS_SQL_Device_Detail	1444	1495	1724
KOQERROR	MS_SQL_Error_Event_Detail	684	687	724
KOQFTBLD	MS_SQL_FileTable_Detail	868	872	987
KOQFGRPD	MS_SQL_Filegroup_Detail	1004	1053	1243
KOQINQRD	MS_SQL_Individual_Queries_Details	2776	2857	3551
KOQSSISD	MS_SQL_Integration_Service_Details	192	226	658
KOQJOB	MS_SQL_Job_Detail	1928	1954	2346
KOQJOBS	MS_SQL_Job_Summary	268	281	522
KOQLOCK	MS_SQL_Lock_Conflict_Detail	678	681	718
KOQLOCKS	MS_SQL_Lock_Detail	1078	1087	1124
KOQLRTS	MS_SQL_Lock_Resource_Type_Summary	244	285	529
KOQLOKSU	MS_SQL_Lock_Summary	300	333	1096
KOQLSDBD	MS_SQL_Log_Shipping_DB_Details	1796	1815	1852
KOQLSERR	MS_SQL_Log_Shipping_Errors	1182	1191	1228
KOQLSJBD	MS_SQL_Log_Shipping_Jobs_Detail	1442	1456	1493
KOQMEMGR	MS_SQL_Memory_Manager	276	457	1208
KOQPROBD	MS_SQL_Problem_Detail	804	808	845
KOQPROBS	MS_SQL_Problem_Summary	278	277	437
KOQPRCD	MS_SQL_Process_Detail	934	975	1207
KOQPRCS	MS_SQL_Process_Summary	302	417	1225
KOQSRVR	MS_SQL_Remote_Servers	284	281	318

Table 8. Capacity planning for historical data logged by the Microsoft SQL Server agent (continued)

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KOQRPOOL	MS_SQL_Resource_Pool_Stats	660	846	1597
KOQRUNQD	MS_SQL_Running_Queries_Details	3144	3167	3513
KOQSRVD	MS_SQL_Server_Detail	616	702	1294
KOQSRVRE	MS_SQL_Server_Enterprise_View	1632	1801	2639
KOQSVRPR	MS_SQL_Server_Properties	980	987	1024
KOQSRVS	MS_SQL_Server_Summary	450	712	2387
KOQSTRNS	MS_SQL_Server_Transactions_Summary	276	303	844
KOQSBACT	MS_SQL_Service_Broker_Activation	572	604	875
KOQSBSTA	MS_SQL_Service_Broker_Statistics	268	320	1050
KOQSBTRP	MS_SQL_Service_Broker_Transport	260	311	982
KOQSRVCD	MS_SQL_Services_Detail	620	628	665
KOQSTATD	MS_SQL_Statistics_Detail	280	289	431
KOQSTATS	MS_SQL_Statistics_Summary	310	451	1298
KOQTBLD	MS_SQL_Table_Detail	1412	1450	1706
KOQTBLs	MS_SQL_Table_Summary	204	276	442
KOQSQL	MS_SQL_Text	2997	3017	3105
KOQWLGS	MS_SQL_Workload_Group_Stats	648	795	1393

For more information about historical data collection, see the *IBM Tivoli Monitoring Administrator's Guide*.

Chapter 5. Situations reference

This chapter contains an overview of situations, references for detailed information about situations, and descriptions of the predefined situations included in this monitoring agent.

About situations

A *situation* is a logical expression involving one or more system conditions. Situations are used to monitor the condition of systems in your network. You can manage enterprise situations from the Tivoli Enterprise Portal by using the Situation editor. You can manage enterprise situations from the Tivoli Enterprise Portal by using the Situation Editor or from the command line interface using the `tacmds` for situations.

The IBM Tivoli Monitoring agents that you use to monitor your system environment are shipped with a set of predefined situations that you can use as-is or you can create new situations to meet your requirements. Predefined situations contain attributes that check for system conditions common to many enterprises.

Using predefined situations can improve the speed with which you can begin using the Microsoft SQL Server agent. You can examine and, if necessary, change the conditions or values being monitored by a predefined situation to those best suited to your enterprise.

Note: The predefined situations provided with this monitoring agent are not read-only. Do not edit these situations and save over them. Software updates will write over any of the changes that you make to these situations. Instead, clone the situations that you want to change to suit your enterprise.

You can display predefined situations and create your own situations using the Situation editor. The left frame of the Situation editor initially lists the situations associated with the Navigator item that you selected. When you click a situation name or create a new situation, the right frame opens with the following tabs:

Formula

Condition being tested

Distribution

List of managed systems (operating systems, subsystems, or applications) to which the situation can be distributed.

Expert Advice

Comments and instructions to be read in the event workspace

Action

Command to be sent to the system

Until Duration of the situation

IBM Tivoli Monitoring situations are distributed to managed systems and the Microsoft SQL Server agent only has one type of managed system, the SQL server instance. When a situation is distributed to an SQL server instance, the situation obtains all rows of data for the attribute group that is used by that situation. If all database rows are returned, the situation occurs for all of the databases.

To restrict a situation to evaluate only a subset of data rows, add filters to modify the situation. You can add one or more situation filters so that the situation evaluates only the desired database or database table targets. For example, add the name of the database you want to monitor, such as `Database Name = "Northwind"`, to the situation so that the situation only monitors the Northwind database.

More information about situations

The *IBM Tivoli Monitoring User's Guide* contains more information about predefined and custom situations and how to use them to respond to alerts.

For a list of the predefined situations for this monitoring agent and a description of each situation, see [Predefined situations](#) and the information for each individual situation.

Predefined situations

This monitoring agent contains the following predefined situations:

- MS_SQL_Block_Critical
- MS_SQL_Block_Warning
- MS_SQL_Cache_Ave_FreePage_Crit
- MS_SQL_Cache_Ave_FreePage_Warn
- MS_SQL_Cache_Hit_Ratio_Crit
- MS_SQL_Cache_Hit_Ratio_Warn
- MS_SQL_Cache_Max_FreePage_Crit
- MS_SQL_Cache_Max_FreePage_Warn
- MS_SQL_Client_Cnt_Pct_Used_Crit
- MS_SQL_Client_Cnt_Pct_Used_Warn
- MS_SQL_Collection_Status_Warning
- MS_SQL_CPU_Critical
- MS_SQL_CPU_Warning
- MS_SQL_DB_Error_Status
- MS_SQL_DB_FreeSpace_Critical
- MS_SQL_DB_FreeSpace_Warning
- MS_SQL_DB_Num_Errors_Warning
- MS_SQL_DB_Space_Pct_Used_Crit
- MS_SQL_DB_Space_Pct_Used_Warn
- MS_SQL_DB_Status_Crit
- MS_SQL_DB_Suspect_Crit
- MS_SQL_Device_Free_PCT_Warning
- MS_SQL_Error_Warning
- MS_SQL_ErrorLog_Size_Warning
- MS_SQL_Failed_Jobs_Crit
- MS_SQL_Fragmentation_Warn
- MS_SQL_Fragmentation_Crit
- MS_SQL_IO_Disk_Errors_Crit
- MS_SQL_IOERR_Startup_Warning
- MS_SQL_IOError_Curintvl_Warning
- MS_SQL_Log_Freespace_Critical
- MS_SQL_Log_Freespace_Warning
- MS_SQL_Log_Shipping_Warning
- MS_SQL_Log_Shipping_Status_Err
- MS_SQL_Log_Space_Pct_Used_Crit
- MS_SQL_Log_Space_Pct_Used_Warn

- MS_SQL_Log_Suspend_Warning
- MS_SQL_LogonPct_Crit
- MS_SQL_LogonPct_Warning
- MS_SQL_Long_Running_Transaction_Warning
- MS_SQL_Mirroring_Failover_Critical
- MS_SQL_Mirroring_Lagging_Warning
- MS_SQL_Mirroring_Suspended_Error
- MS_SQL_Network_Read_Rate_Crit
- MS_SQL_Network_Read_Rate_Warn
- MS_SQL_Network_Write_Rate_Crit
- MS_SQL_Network_Write_Rate_Warn
- MS_SQL_Num_Process_Blocked_Crit
- MS_SQL_Num_Process_Blocked_Warn
- MS_SQL_Oldest_Transaction_Crit
- MS_SQL_Oldest_Transaction_Warn
- MS_SQL_Opt_Stats_Age_Crit
- MS_SQL_Opt_Stats_Age_Warn
- MS_SQL_Pct_Block_Warning
- MS_SQL_Pct_CPU_Yields_Warning
- MS_SQL_Pct_IO_Warning
- MS_SQL_PCT_MAX_Locks_Critical
- MS_SQL_PCT_MAX_Locks_Warning
- MS_SQL_ProbAge_GT_17_Warning
- MS_SQL_Proc_Buffs_Active_Crit
- MS_SQL_Proc_Buffs_Active_Warn
- MS_SQL_Proc_Buffs_Used_Crit
- MS_SQL_Proc_Buffs_Used_Warn
- MS_SQL_Proc_Cache_Active_Crit
- MS_SQL_Proc_Cache_Active_Warn
- MS_SQL_Proc_Cache_Used_Crit
- MS_SQL_Proc_Cache_Used_Warn
- MS_SQL_Process_Blocked_Warning
- MS_SQL_Process_Infected_Warning
- MS_SQL_ProcessLockSleep_Warning
- MS_SQL_ProcessOthSleep_Warning
- MS_SQL_Processes_Bad_Warning
- MS_SQL_Processes_Stop_Warning
- MS_SQL_Rem_Serv_Stat_Critical
- MS_SQL_Repl_Latency_Crit
- MS_SQL_Repl_Latency_Warn
- MS_SQL_Status_Critical
- MS_SQL_Status_Inactive
- MS_SQL_Total_Locks_Critical
- MS_SQL_Total_Locks_Warning

Situation descriptions

Each situation description provides information about the situation that you can use to monitor the condition of systems in your network.

MS_SQL_Block_Critical

Declares a critical condition if the number of processes in conflict is greater than or equal to 60.

MS_SQL_Block_Warning

Declares a warning condition if the number of processes in conflict is greater than or equal to 20 and less than 60.

MS_SQL_Cache_Ave_FreePage_Crit

Declares a critical condition if the average number of cache buffers scanned to find a free buffer exceeds the critical threshold. The default threshold is Cache Average Free Scan greater than 15.

MS_SQL_Cache_Ave_FreePage_Warn

Declares a warning condition if the average number of cache buffers scanned to find a free buffer exceeds the warning threshold. The default threshold is Cache Average Free Scan greater than 10 and less than or equal to 15.

MS_SQL_Cache_Hit_Ratio_Crit

Declares a critical condition if the ratio of data cache hits to total data requests exceeds the critical threshold. The default threshold is Cache Hit Ratio less than 70.00.

MS_SQL_Cache_Hit_Ratio_Warn

Declares a warning condition if the ratio of data cache hits to total data requests exceeds the warning threshold. The default threshold is Cache Hit Ratio less than 90.00 and greater than or equal to 70.00.

MS_SQL_Cache_Max_FreePage_Crit

Declares a critical condition if the maximum number of cache buffers scanned to find a free buffer exceeds the critical threshold. The default threshold is Cache Maximum Free Page Scan greater than 15.

MS_SQL_Cache_Max_FreePage_Warn

Declares a warning condition if the maximum number of cache buffers scanned to find a free buffer exceeds the warning threshold. The default threshold is Cache Maximum Free Page Scan greater than 5 and less than or equal to 15.

MS_SQL_Client_Cnt_Pct_Used_Crit

Declares a critical condition if the percentage of client licenses being used exceeds the critical threshold. The default threshold is Client Count Pct Used greater than 90.00.

MS_SQL_Client_Cnt_Pct_Used_Warn

Declares a warning condition if the percentage of client licenses being used exceeds the warning threshold. The default threshold is Client Count Pct Used greater than 70.00 and less than or equal to 90.00.

MS_SQL_Collection_Status_Warning

Declares a warning condition if the collection status is not Active.

MS_SQL_CPU_Critical

Declares a critical condition if the percentage of CPU usage is greater than or equal to 80 and elapsed time since the SQL server was started is at least 10 minutes.

MS_SQL_CPU_Warning

Declares a warning condition if the percentage of CPU usage is greater than or equal to 60 and less than 80 and elapsed time since the SQL server was started is at least 10 minutes.

MS_SQL_DB_Error_Status

Provides an alert if the database has a serious error.

MS_SQL_DB_Freespace_Critical

Declares a critical condition if the percentage of freespace on the database is less than or equal to 10.

MS_SQL_DB_Join_State_Warn

Declares a warning condition if the availability database is in the Not Joined state.

MS_SQL_DB_Space_Pct_Used_Warn

Declares a warning condition if the percentage of space used in the database compared to the total space allowed exceeds the warning threshold. The default threshold is Database Space Pct Used greater than 70.00 and less than or equal to 90.00.

MS_SQL_DB_Space_Pct_Used_Crit

Declares a critical condition if the percentage of space used in the database compared to the total space allowed exceeds the critical threshold. The default threshold is Database Space Pct Used greater than 90.00.

MS_SQL_DB_Num_Errors_Warning

Declares a warning condition if the number of databases with an error status is greater than 0.

MS_SQL_DB_Freespace_Warning

Declares a warning condition if the percentage of freespace in the database is less than or equal to 20 and greater than 10.

MS_SQL_DB_Status_Crit

Declares a critical condition if the database is unavailable because of an error or being taken offline.

MS_SQL_DB_Suspect_Crit

Declares a critical condition if the database is in an inconsistent state because it cannot be restored.

MS_SQL_DB_Sync_Health_Crit

Declares a critical situation when the synchronization health of an availability database is Not Healthy.

MS_SQL_DB_Sync_State_Warn

Declares a warning condition if an availability database is in the Not Synchronizing state.

MS_SQL_Device_Free_PCT_Warning

Declares a warning condition if the percentage of freespace on the device is less than or equal to 10.

MS_SQL_Error_Event_Warning

Declares a warning condition if the SQL Server error is generated with severity greater than 16.

MS_SQL_Error_Warning

Declares a warning condition if the number of SQL Server errors during the current interval is greater than 0.

MS_SQL_ErrorLog_Size_Warning

Declares a warning condition if the error log size is greater than or equal to 10000 bytes.

MS_SQL_Failed_Jobs_Crit

Declares a critical condition if the number of jobs run by the SQLServerAgent service that have failed over the last monitoring interval has exceeded the critical threshold. The default threshold is SQLServerAgent Failed Jobs greater than 0.

MS_SQL_Fragmentation_Crit

Declares a critical condition if the percentage of fragmentation exceeds the critical threshold. The default threshold is Fragmentation greater than 80.00.

MS_SQL_Fragmentation_Warn

Declares a warning condition if the percentage of fragmentation exceeds the warning threshold. The default threshold is Fragmentation greater than 30.00 and less than or equal to 80.00.

MS_SQL_Group_Role_Crit

Declares a critical condition if the role of an SQL Server instance in an availability group changes from Primary to Secondary or Resolving, or if the role changes to Invalid.

MS_SQL_IOERR_Startup_Warning

Declares a warning condition if the number of IO errors since startup is greater than or equal to 10 and the percentage of IO errors in the current interval is greater than or equal to 5.

MS_SQL_IO_Disk_Errors_Crit

Declares a critical condition if the number of SQL Server read/write disk errors exceeds the critical threshold. The default threshold is Statistics Name equals IO errors and the Current Value is greater than 0.

MS_SQL_IOError_Curintvl_Warning

Declares a warning condition if the percentage of IO errors during the current interval is greater than or equal to 80.

MS_SQL_Log_Freespace_Critical

Declares a critical condition if the percentage of freespace in the log is less than or equal to 10.

MS_SQL_Log_Freespace_Warning

Declares a warning condition if the percentage of freespace in the log is less than or equal to 20 and greater than 10.

MS_SQL_Log_Shipping_Status_Err

Declares an error condition if the backup or restore job results in an error. The Job Status is displayed as Error.

MS_SQL_Log_Shipping_Warning

Declares a warning condition if the configured threshold, minutes between backup or restore jobs are exceeded. Test for Threshold Exceeded = Yes.

MS_SQL_Log_Space_Pct_Used_Crit

Declares a critical condition if the percentage of space used by the transaction log compared to the total allocated transaction log size exceeds the critical threshold. The default threshold is Log Space Pct Used greater than 90.00.

MS_SQL_Log_Space_Pct_Used_Warn

Declares a warning condition if the percentage of space used by the transaction log compared to the total allocated transaction log size exceeds the warning threshold. The default threshold is Log Space Pct Used greater than 70.00 and less than or equal to 90.00.

MS_SQL_Log_Suspend_Warning

Declares a warning condition if the number of processes in log suspend is greater than or equal to 1.

MS_SQL_LogonPct_Crit

Declares a critical condition if the number of current user connections as a percentage of the available user connections has exceeded the critical threshold. The default threshold is Pct Max Logons Active greater than 90.00.

MS_SQL_LogonPct_Warning

Declares a warning condition if the percentage of active logins is greater than or equal to 90.

MS_SQL_Long_Running_Transaction_Warning

Declares a warning condition if the same database transaction is identified to be the longest running transaction over two consecutive 15 minutes interval cycles.

MS_SQL_Mirror_Failover_Critical

Declares a critical condition that is found only on the principal server after a failover, but the server is not changed to the mirror role (Mirroring_State is equal to 'Pending Failover').

MS_SQL_Mirror_Lagging_Warning

Declares a warning condition if the contents of the mirror database lag the contents of the principal database (Mirroring State is equal to 'Synchronizing'). The principal server instance sends log records to the mirror server instance, which applies the changes to the mirror database to roll it forward.

MS_SQL_Mirror_Suspended_Warning

Declares a warning condition if the mirror copy of the database is not available (Mirroring_State is equal to 'Suspended'). The principal database runs without sending any logs to the mirror server, a condition known as running exposed.

MS_SQL_Network_Read_Rate_Crit

Declares a critical condition if the rate (packets per second) of Tabular Data Stream (TDS) packets being read from the network exceeds the critical threshold. The default threshold is Network Read Rate (per Sec.) greater than 150.

MS_SQL_Network_Read_Rate_Warn

Declares a warning condition if the rate (packets per second) of Tabular Data Stream (TDS) packets being read from the network exceeds the warning threshold. The default threshold is Network Read Rate (per Sec.) greater than 50 and less than or equal to 150.

MS_SQL_Network_Write_Rate_Crit

Declares a critical condition if the rate (packets per second) of Tabular Data Stream (TDS) packets being written from the network exceeds the critical threshold. The default threshold is Network Write Rate (per Sec.) greater than 150.

MS_SQL_Network_Write_Rate_Warn

Declares a warning condition if the rate (packets per second) of Tabular Data Stream (TDS) packets being written from the network exceeds the warning threshold. The default threshold is Network Write Rate (per Sec.) greater than 50 and less than or equal to 150.

MS_SQL_Num_Process_Blocked_Crit

Declares a critical condition if the number of blocked processes exceeds the critical threshold. The default threshold is Total Processes Blocked greater than 10.

MS_SQL_Num_Process_Blocked_Warn

Declares a warning condition if the number of blocked processes exceeds the warning threshold. The default threshold is Total Processes Blocked greater than 2 and less than or equal to 10.

MS_SQL_Oldest_Transaction_Crit

Declares a critical condition if the age (in minutes) of the oldest transaction in the database's transaction log exceeds the critical threshold. The default threshold is Oldest Open Transaction (min) greater than 15.

MS_SQL_Oldest_Transaction_Warn

Declares a warning condition if the age (in minutes) of the oldest transaction in the database's transaction log exceeds the warning threshold. The default threshold is Oldest Open Transaction (min) greater than 5 and less than or equal to 15.

MS_SQL_Opt_Stats_Age_Crit

Declares a critical condition if the time (in minutes) since the table statistics were last updated exceeds the critical threshold. The default threshold is Optimizer Statistics Age (in Min.) greater than 1440.

MS_SQL_Opt_Stats_Age_Warn

Declares a warning condition if the time (in minutes) since the table statistics were last updated exceeds the warning threshold. The default threshold is Optimizer Statistics Age (in Min.) greater than 180 and less than or equal to 1440.

MS_SQL_Pct_Block_Warning

Declares a warning condition if the percentage of processes in conflict is greater than or equal to 50.

MS_SQL_Pct_CPU_Yields_Warning

Declares a warning condition if the percentage of CPU yields is greater than 75.

MS_SQL_Pct_IO_Warning

Declares a warning condition if the percentage of time spent for IO operations is greater than or equal to 90.

MS_SQL_PCT_MAX_Locks_Critical

Declares a critical condition if the number of locks reaches or exceeds 80 percent of the maximum locks allowed.

MS_SQL_PCT_MAX_Locks_Warning

Declares a warning condition if the number of locks ranges from 60 to 80 percent of the maximum locks allowed.

MS_SQL_ProbAge_GT_17_Warning

Declares a warning condition if the age of the last error is 30 minutes or greater and the highest severity level is greater than or equal to 17.

MS_SQL_Proc_Buffs_Active_Crit

Declares a critical condition if the percentage of active procedure buffers (buffers containing procedure(s) actively being executed) to total available procedure buffers exceeds the critical threshold. The default threshold is Procedure Buffers Pct Active less than 75.00.

MS_SQL_Proc_Buffs_Active_Warn

Declares a warning condition if the percentage of active procedure buffers (buffers containing one or more procedures) actively being executed) to total available procedure buffers exceeds the warning threshold. The default threshold is Procedure Buffers Pct Active less than 95.00 and greater than or equal to 75.00.

MS_SQL_Proc_Buffs_Used_Crit

Declares a critical condition if the percentage of used procedure buffers to total available procedure buffers exceeds the critical threshold. The default threshold is Procedure Buffers Pct Used less than 75.00.

MS_SQL_Proc_Buffs_Used_Warn

Declares a warning condition if the percentage of used procedure buffers to total available procedure buffers exceeds the warning threshold. The default threshold is Procedure Buffers Pct Used less than 95.00 and greater than or equal to 75.00.

MS_SQL_Proc_Cache_Active_Crit

Declares a critical condition if the percentage of the procedure buffer cache with currently executing procedures exceeds the critical threshold. The default threshold is Procedure Cache Pct Active less than 75.00.

MS_SQL_Proc_Cache_Active_Warn

Declares a warning condition if the percentage of the procedure buffer cache with procedures in it exceeds the warning threshold. The default threshold is Procedure Cache Pct Active less than 95.00 and greater than or equal to 75.00.

MS_SQL_Proc_Cache_Used_Crit

Declares a critical condition if the percentage of used procedure buffers to total available procedure buffers exceeds the critical threshold. The default threshold is Procedure Cache Pct Used less than 75.00.

MS_SQL_Proc_Cache_Used_Warn

Declares a warning condition if the percentage of the procedure buffer cache with currently executing procedures exceeds the warning threshold. The default threshold is Procedure Cache Pct Used less than 95.00 and greater than or equal to 75.00.

MS_SQL_Process_Blocked_Warning

Declares a warning condition if the percentage of blocked processes is greater than or equal to 25 and the number of blocked processes is greater than or equal to 5.

MS_SQL_Process_Infected_Warning

Declares a warning condition if the number of infected processes is greater than or equal to 1.

MS_SQL_Processes_Bad_Warning

Declares a warning condition if the number of bad processes is greater than or equal to 1.

MS_SQL_Processes_Stop_Warning

Declares a warning condition if the number of stopped processes is greater than or equal to 5 and the percentage of stopped processes is greater than or equal to 10.

MS_SQL_ProcessLockSleep_Warning

Declares a warning condition if the percentage of processes in locksleeep is greater than or equal to 10 and the number of processes in locksleeep is greater than or equal to 5.

MS_SQL_ProcessOthSleep_Warning

Declares a warning condition if the number of processes in other sleep is greater than or equal to 2 and the percentage of processes in othersleeep is greater than or equal to 50.

MS_SQL_Rem_Serv_Stat_Critical

Declares a critical condition if the remote server has a status of Inactive.

MS_SQL_Repl_Latency_Crit

Declares a critical condition if the time (in milliseconds) between the time a transaction marked for replication is entered into the transaction log and the time the transaction is read out of the log and delivered to the distribution database exceeds the critical threshold. The default threshold is Replication Latency (ms) greater than 900,000.

MS_SQL_Repl_Latency_Warn

Declares a warning condition if the time (in milliseconds) between the time a transaction marked for replication is entered into the transaction log and the time the transaction is read out of the log and delivered to the distribution database exceeds the warning threshold. The default threshold is Replication Latency (ms) greater than 120,000 and less than or equal to 900,000.

MS_SQL_Replica_Join_State_Warn

Declares a warning condition if the availability replica is in the Not Joined state.

MS_SQL_Replica_Role_Health_Crit

Declares a critical condition if the role of the availability replica is not Primary or Secondary. This situation is supported for SQL Server 2012, or later.

MS_SQL_Replica_Sync_Health_Warn

Declares a warning condition if the synchronization health of an availability replica is not healthy.

MS_SQL_Replicas_Conn_State_Crit

Declares a critical condition if the state of an availability replica is Disconnected.

MS_SQL_Status_Critical

Declares a critical condition if the SQL Server status is not active.

MS_SQL_Status_Inactive

Initiates an action if the SQL server is inactive.

MS_SQL_Total_Locks_Critical

Declares a critical condition if the number of locks reaches or exceeds 4000.

MS_SQL_Total_Locks_Warning

Declares a warning condition if the number of locks reaches or exceeds 200.

Chapter 6. Take Action commands reference

This chapter contains an overview of Take Action commands, references for detailed information about Take Action commands, and descriptions of the Take Action commands included in this monitoring agent.

About Take Action commands

Take Action commands can be run from the desktop or included in a situation or a policy.

When included in a situation, the command executes when the situation becomes true. A Take Action command in a situation is also called *reflex automation*. When you enable a Take Action command in a situation, you automate a response to system conditions. For example, you can use a Take Action command to send a command to restart a process on the managed system or to send a text message to a cell phone.

Advanced automation uses policies to perform actions, schedule work, and automate manual tasks. A policy comprises a series of automated steps called activities that are connected to create a workflow. After an activity is completed, Tivoli Enterprise Portal receives return code feedback, and advanced automation logic responds with subsequent activities prescribed by the feedback.

More information about Take Action commands

For more information about working with Take Action commands, see the *IBM Tivoli Monitoring User's Guide*.

For a list of the Take Action commands for this monitoring agent and a description of each command, see "Predefined Take Action commands" section and the information for each individual command.

Predefined Take Action commands

The remaining sections of this chapter contain descriptions of each of these Take Action commands, which are listed alphabetically. The following information is provided following the description of each Take Action command:

Authorization role

The required authorization role, plus any required Microsoft SQL Server permissions

Arguments

List of arguments, if any, for the Take Action with a short description and default value for each one. Each argument is positionally dependent and mandatory. The arguments for the MS SQL Take Actions are organized according to their respective positions on the GUI. If any argument requires an embedded space, the argument must be enclosed in double quotation marks.

When an argument is not required, a blank argument value must be used. A blank argument value is double quotation marks: "".

Destination systems

Where the command is to be executed: on the Managed System (monitoring agent) where the agent resides or on the Managing System (Tivoli Enterprise Monitoring Server) to which it is connected

Usage notes

Additional relevant notes for using the Take Actions

Note: See the Windows Help or the *Microsoft SQL Server Books On-line* system for additional information, formatting considerations, and so on, for the Take Actions.

User ID and password combinations

All Take Actions for this monitoring agent can optionally use an explicit user ID and password when connecting to the MS SQL Server. If you specify a blank user ID and password, the "Log on as" Windows user ID and password credentials of the Microsoft SQL Server agent are used and the connection to the server is established using Windows authentication. Table 9 describes possible combinations for specifying the user ID and password.

Table 9. User ID and password combinations

User ID	Password	Outcome
A given string	A given string	Both strings are used for user ID and password when connecting to the server.
A given string	Blank (use double quotation marks)	The user ID string and blank password are used when connecting to the server.
Blank (use double quotation marks)	A given string	The windows authentication information associated with the agent process (service) is used when connecting to the server.
Blank (use double quotation marks)	Blank (use double quotation marks)	The windows authentication information associated with the agent process (service) is used when connecting to the server.

This monitoring agent contains the following Take Action commands:

- Database Check-up and Tune-up
- Dump Database
- Dump Transaction Log
- Rebuild Fragmented Indexes
- Refresh Query Optimizer Statistics
- Start SQL server
- Stop SQL server
- Update Space Usage Information

Database Check-up and Tune-up action

Performs a series of regular maintenance tasks to correct inconsistencies in the database and tune the performance of the system. The Database Check-up and Tune-up Take Action includes the following tasks:

- Runs DBCC CHECKDB, DBCC CHECKALLOC, DBCC CHECKCATALOG, and DBCC CHECKIDENT.
 - DBCC CHECKDB checks each table to ensure that the following statuses are true:
 - Index and data pages are correctly linked.
 - Indexes are in proper sorted order.
 - All pointers are consistent.
 - Data information on each page is reasonable.
 - Page offsets are reasonable.
 - DBCC CHECKCATALOG identifies inconsistencies between system tables and checks the validity of the last checkpoint.
 - DBCC CHECKIDENT corrects disrupted identity column values that can result from ungraceful shutdowns or unlogged operations.
- Rebuilds fragmented indexes and runs UPDATE STATISTICS to ensure that the query optimizer is using the most recent information.

- Drops and reloads stored procedures.
- Runs DBCC UPDATEUSAGE to ensure more accurate space usage statistics.

Authorization role

Microsoft SQL Server permissions are **sysadmin** or **db_owner**.

Arguments

SQL Server Name (**sql_server_name**)

SQL Server instance name

Database Name (**database_name**)

Name of the database

Fragmentation Limit (**fragmentation_limit**)

Value for the index fragmentation limit. When the actual index fragmentation percentage is greater than the entered number, the index is rebuilt. The default value is 30.

SQL Server ID (**sql_server_id**)

Note: For information about specifying the SQL Server login ID, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

SQL Server Password (**sql_server_password**)

Note: For information about specifying the SQL Server login password, see see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

Destination system

Managed system

Usage notes

Not applicable

Dump Database action

Performs a full database backup. The Dump Database Take Action does the following tasks:

- Provides a way to check the database integrity before backing up to ensure that the backup is viable
- Provides a way to designate an existing backup device, a temporary backup device, or the previous backup device when possible
- Provides a way to specify a list of backup options. For example, you can append to a previous backup or overwrite it. You can specify the type of backup to perform.

Note: The available backup options vary by version of Microsoft SQL Server.

- Checks the readability of the backup after it is complete. Verification consists of reading header information to ensure that the backup set is complete and all volumes are readable. Data structure in the backup volume is not verified.
- Writes some statistics on the backup procedure along with the header information on the backup device to a log file.

Authorization role

Microsoft SQL Server permissions are **sysadmin**, **db_owner**, or **db_backupoperator**.

Arguments

See the Microsoft SQL Server documentation for more detailed information about the arguments, including which options are required and formatting.

SQL Server Name (**sql_server_name**)

SQL Server instance name

Database Name (database_name)

Name of the database

Backup Device (backup_device)

Name of the logical backup device. This value is ignored when the Use Last Backup Device argument is 1.

Note: The Dump Database take action does not work if you specify physical backup device name for this parameter.

Backup Options (backup_options)

A comma-separated list of options. For all options that take an argument, use either a Microsoft SQL Server variable or literal text. You can use the following options:

BLOCKSIZE=number

Sets the block size to use when writing the backup. When writing to disk or tape, the size is set automatically unless overridden. When writing to PIPE, the default size is 65,536 unless overridden.

DESCRIPTION='text'

Specifies a description for the backup set. The maximum length is 255 characters. The default value is a blank.

DIFFERENTIAL

Specifies a differential backup rather than a full backup. The default value is a full backup.

EXPIREDATE=date or RETAINDDAYS=number

Specifies when the backup file can be overwritten. Only one of the following options may be used:

- EXPIREDATE specifies the expiration date.
- RETAINDDAYS specifies how many days to retain the backup.

FORMAT or NOFORMAT

Writes the media header on all volumes used for the backup and rewrites the backup device. An existing media header is overwritten, and all media contents are invalidated. Any existing password is ignored.

Use this option on complete backup sets only. Using it on a single tape that belongs to an existing striped backup set makes the entire backup set useless.

Using FORMAT implicitly sets INIT and SKIP. They do not have to be specified separately.

INIT or NOINIT

INIT means the backup overwrites any previous backup on the backup device. The media header is preserved. NOINIT means the backup is appended to the previous backup, if any. The default value is NOINIT.

The backup does not overwrite if it encounters any of the following conditions:

- The media contains unexpired backup sets. You can specify SKIP to override this condition.
- An explicitly given backup set name does not match the name on the backup media. You can specify SKIP to override this condition.
- The backup media is password-protected.
- The backup media is encrypted.

MEDIADESCRIPTION= 'text'

Specifies a description for the media set. The maximum length is 128 characters. The default value is BLANK.

MEDIANAME='text'

Specifies the media name for the backup. When specified, if the existing MEDIANAME on the backup media does not match, the backup fails.

NAME='text'

Specifies the name of the backup set. The maximum length is 128 characters. The default value is BLANK.

RESTART

Use this option when restarting a backup operation that was interrupted. SQL Server attempts to restart the backup where it was interrupted, which can save time.

SKIP or NOSKIP

SKIP disables the checking for name match and expiration date of all backup sets on the media. SKIP can be used with INIT. NOSKIP can be used with FORMAT to override its default behavior. The default value is NOSKIP except where otherwise noted.

STATS=percent

Specifies how often to write a progress message, which is expressed as a percent complete. The default value is 10.

UNLOAD or NOUNLOAD

Use this option for tape devices only. UNLOAD specifies rewinding and unloading (ejecting) the media when the backup is complete. The default value is NOUNLOAD.

Validate Database Integrity (validate_db)

Runs DBCC CHECKDB, NEWALLOC, CHECKCATALOG, and CHECKIDENT to determine if the database is corrupted. If the database fails the integrity checks, the backup is not performed. Specify 1 to execute the integrity check, otherwise specify 0.

Use Last Backup Device (use_last_backup_device)

You can use the same backup device that you used the last time you backed up the database. Specify 1 for this argument if you want to use the previous backup device. If you specify 0, you must provide the name of the backup device, and the integrity of the database is not validated.

SQL Server ID (sql_server_id)

Note: For information about specifying the SQL Server login ID, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

SQL Server Password (sql_server_password)

Note: For information about specifying the SQL Server login password, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

Destination system

Managed system

Usage notes

When using this Take Action, consider the following information:

- To do backups to remote machines, this monitoring agent must have rights on the remote machine and cannot be running under the local system account.
- Because the Take Action could fail and return a warning, do not make master backups to striped devices.

- Before backing up a database, back up a transaction log using the “Dump Transaction Log action,” which clears the log and decreases the total time that is required to back up the database. However, if the database backup is not done immediately or if it fails, all the logged transactions could be lost. This Take Action does not truncate the log before backing up the database.

Dump Transaction Log action

Backs up, truncates, or backs up and truncates a database’s transaction log. The Dump Transaction Log Take Action does the following tasks:

- Provides a way to designate an existing backup device, a temporary backup device, or the previous backup device when possible.
- Provides a way to specify a list of backup options. For example, you can append a backup to a previous backup or overwrite the previous backup. You can specify the type of backup to perform. The backup options available vary according to the version of Microsoft SQL Server.
- Checks the backup’s readability after the backup is complete. Verification consists of reading header information to ensure that the backup set is complete and all volumes are readable. Data structure in the backup volume is not verified.
- Writes some statistics about the backup procedure and writes the header information on the backup device to a log file.

Authorization role

Microsoft SQL Server permissions are **sysadmin**, **db_owner**, or **db_backupoperator**.

Arguments

See the Microsoft SQL Server documentation for more detailed information about the arguments, including which options are required and formatting.

SQL Server Name (**sql_server_name**)

SQL Server instance name

Database Name (**database_name**)

Name of the database

Backup Device (**backup_device**)

Name of the logical backup device. This value is ignored when the Use Last Backup Device argument is set to 1.

Note: The Dump Transaction Log take action does not work if you specify physical backup device name for this parameter.

Backup Options (**backup_options**)

A comma-separated list of options. For all options that take an argument, use either a Microsoft SQL Server variable or literal text. You can use the following options:

BLOCKSIZE=number

Sets the block size to use when writing the backup. When writing to disk or tape, the size is set automatically unless overridden. When writing to PIPE, the default size is 65,536 unless overridden.

DESCRIPTION='text'

Specifies a description for the backup set. The maximum length is 255 characters. The default value is a blank.

DIFFERENTIAL

Specifies a differential backup rather than a full backup. The default value is a full backup.

EXPIREDATE=*date* or RETAINDAYS=*number*

Specifies when the backup file may be overwritten. Only one of these options may be used. EXPIREDATE specifies the expiration date. RETAINDAYS specifies how many days to retain the backup.

FORMAT or NOFORMAT

Writes the media header on all volumes used for the backup and rewrites the backup device. An existing media header is overwritten and all media contents are invalidated. Any existing password is ignored.

Use this option on complete backup sets only. Using it on a single tape that belongs to an existing striped backup set renders the entire backup set useless.

Using FORMAT implicitly sets INIT and SKIP. They do not have to be specified separately.

INIT or NOINIT

INIT means the backup overwrites any previous backup on the backup device. The media header is preserved. NOINIT means the backup is appended to the previous backup, if any. The default value is NOINIT.

The backup does not overwrite if it encounters any of the following conditions:

- The media contains unexpired backup sets. You can specify SKIP to override this condition.
- An explicitly given backup set name does not match the name on the backup media. You can specify SKIP to override this condition.
- The backup media is password-protected.
- The backup media is encrypted.

MEDIADESCRIPTION= '*text*'

Specifies a description for the media set. The maximum length is 128 characters. The default value is BLANK.

MEDIANAME= '*text*'

Specifies the media name for the backup. When specified, if the existing MEDIANAME on the backup media does not match, the backup fails.

NAME= '*text*'

Specifies the name of the backup set. The maximum length is 128 characters. The default value is BLANK.

RESTART

Use this option when restarting a backup operation that was interrupted. SQL Server attempts to restart the backup where it was interrupted, which can save time.

SKIP or NOSKIP

SKIP disables the checking for name match and expiration date of all backup sets on the media. SKIP can be used with INIT. NOSKIP can be used with FORMAT to override its default behavior. The default value is NOSKIP except where otherwise noted.

STATS=*percent*

Specifies how often to write a progress message, expressed as a percent complete. The default value is 10.

UNLOAD or NOUNLOAD

Use this option for tape devices only. UNLOAD specifies rewinding and unloading (ejecting) the media when the backup is complete. The default value is NOUNLOAD.

Dump Options (dump_option)

Specifies how to backup and/or truncate a transaction log.

Note: The dump option names should be passed literally as strings when issuing this Take Action.

Use one of the following valid choices:

- **backup** backs up the transaction log, but does not truncate it.
- **BackupAndTruncate** backs up the transaction log and truncates it. Truncating the log clears it of all non-active transactions. Use this option only when the transaction log and the data are on separate device fragments.
- **Truncate** truncates the transaction log, but does not back it up. This operation is logged.
- **TruncateNoLog**, which is used when log is full, truncates the transaction log, but does not back it up. It is the only choice when the log is full. This operation is not logged.

Use Last Backup Device (use_last_backup_device)

You can use the same backup device that you used the last time you backed up the database. Specify 1 for this argument if you want to use the previous backup device. If you specify 0, you must provide the name of the backup device, and the integrity of the database is not validated.

SQL Server ID (sql_server_id)

Note: For information about specifying the SQL Server login ID, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

SQL Server Password (sql_server_password)

Note: For information about specifying the SQL Server login password, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

Destination system

Managed system

Usage notes

When using this Take Action, observe the following information:

- If the media fails, specify backup so that the transaction log can be dumped even if the database is inaccessible. The backup option provides up-to-the-minute media recovery when the master database and the log portion of the user database reside on undamaged database devices and when only the data portion of the user database is damaged. Specify the backup option to back up the master database log on its own.
- The "Dump Database action" on page 241 backs up both the data and the log. The Dump Transaction Take Action backs up only the transaction log and not the data. When you dump a database, you are doing a database backup. When you dump a transaction log, you could be backing it up, truncating it, or both.

Rebuild Fragmented Indexes action

Rebuilds clustered and non-clustered indexes on tables. You can specify the level of fragmentation that triggers reindexing.

Authorization role

Microsoft SQL Server permissions are **sysadmin** or **db_owner**.

Arguments

SQL Server Name (sql_server_name)

SQL Server instance name

Database Name (database_name)

Name of the database

Table Name (table_name)

Updates the indexes for a specified table or for all tables. May have an empty string value. If blank, this argument updates indexes for all tables.

Index Name (index_name)

Updates indexes of a particular name that you specify in this field, or all indexes. May have an empty string value. If this argument is blank, it updates all indexes.

Rebuild if fragmentation % greater than (fragmentation_limit)

Rebuilds the index when the index reaches this percentage of fragmentation. If the actual index fragmentation percentage is greater than the number you enter here, the Take Action rebuilds the index. To rebuild the indexes under all circumstances, use zero. The default value is 30 percent.

Fill Factor (fill_factor)

Specifies how full you want the new page to be. Enter a value that is the percentage of the page to be filled. For example, entering 20 means that you want to fill one-fifth of the page, leaving 80 percent free. To rebuild with the same FILLFACTOR used in the previous rebuild, specify zero. If the actual fill factor is greater than the number you enter, this Take Action rebuilds the index using the specified fill factor. Leaving some space on a page can improve the performance on clustered indexes. If you specify a low fill factor, remember that the database takes up more disk space.

SQL Server ID (sql_server_id)**Note:** For information about specifying the SQL Server login ID, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.**SQL Server Password (sql_server_password)****Note:** For information about specifying the SQL Server login password, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.Table 10 shows how entries in the **Table Name** and **Index Name** arguments determine which indexes and tables are updated.*Table 10. Index and table updates*

Table Name	Table Exists	Index Name	Index Exists	Outcome
Entered	Y	Entered	Y	Updates the specified index on the specified table.
Entered	Y	Entered	N	Error. Specified index does not exist.
Entered	Y	Blank (use double quotation marks "")	-	Updates all indexes on the specified table.
Entered	N	Entered	Y	Error. Specified index does not exist.
Entered	N	Entered	N	Error. Specified index does not exist.

Table 10. Index and table updates (continued)

Table Name	Table Exists	Index Name	Index Exists	Outcome
Entered	N	Blank (use double quotation marks "")	-	Error. Specified index does not exist.
Blank (use double quotation marks "")	-	Entered	Y	Updates all indexes on the specified table
Blank (use double quotation marks "")	-	Entered	N	Error. Specified index does not exist.
Blank (use double quotation marks "")	-	Blank (use double quotation marks "")	-	Updates all indexes.

Destination system

Managed system

Usage notes

You can improve table fragmentation only if the table has a clustered index that can be rebuilt. Table data is directly related to the clustered index because it is the leaf level of the index. The Take Action first checks the clustered index (INDID = 1).

Rebuilding the clustered index reorders the data pages and causes all other indexes to be rebuilt. If a table has no clustered index or the clustered index is not fragmented (non-clustered index fragmentation), you might still want to run this Take Action. It checks the non-clustered indexes for fragmentation, which the Take Action automatically rebuilds on an individual basis if necessary. Non-clustered indexes have an INDID > 1.

Refresh Query Optimizer Statistics action

Refreshes the optimizer statistics for all indexes on a table or for a specific index and forces every affected stored procedure to recompile before it is run again.

Authorization role

Microsoft SQL Server permissions are **sysadmin** or the table owner.

Arguments

SQL Server Name (sql_server_name)

SQL Server instance name

Database Name (database_name)

Name of the database

Table Name (table_name)

Updates the indexes for a specified table. May have an empty string value. If blank, this argument updates indexes for all tables.

Index Name (index_name)

Updates the specified indexes. May have an empty string value. If this argument is blank, it updates all indexes.

SQL Server ID (sql_server_id)

Note: For information about specifying the SQL Server login ID, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

SQL Server Password (sql_server_password)

Note: For information about specifying the SQL Server login password, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

Table 11 shows how entries in the **Table Name** and **Index Name** arguments determine which indexes and tables are updated.

Table 11. Index and table updates

Table Name	Table Exists	Index Name	Index Exists	Outcome
Entered	Y	Entered	Y	Updates the specified index on the specified table.
Entered	Y	Entered	N	Error. Specified index does not exist.
Entered	Y	Blank (use double quotation marks "")	-	Updates all indexes on the specified table.
Entered	N	Entered	Y	Error. Specified index does not exist.
Entered	N	Entered	N	Error. Specified index does not exist.
Entered	N	Blank (use double quotation marks "")	-	Error. Specified index does not exist.
Blank (use double quotation marks "")	-	Entered	Y	Updates all indexes on the specified table.
Blank (use double quotation marks "")	-	Entered	N	Error. Specified index does not exist.
Blank (use double quotation marks "")	-	Blank (use double quotation marks "")	-	Updates all indexes.

Destination system

Managed system

Usage notes

This Take Action is limited by the sp_recompile stored procedure. For example, you cannot use sp_recompile with system tables. If you must have stored procedures based solely on system tables that were dropped and recreated, use the "Database Check-up and Tune-up action" on page 240. If you update a single index on a table, all stored procedures based on that table are recompiled.

Start SQL Server action

Starts any or all of the SQL Server services. The command starts an instance of SQL Server. By default, the instance name is MSSQLServer. If you have a non-default instance name, replace net start MSSQLSERVER with net start MSSQL\$*instancename* in the command area.

Authorization role

None

Arguments**MS SQL Server Windows Service Name**

The Windows service name of the MS SQL Server instance you want to start. The default service name that the default SQL Server instance uses is `MSSQLSERVER`. This Take Action uses that service name as the default. To start a named MS SQL Server instance, set the value of this argument to the service name that is associated with the named SQL Server instance.

Destination system

Managed system

Usage notes

Not applicable

Stop SQL Server action

Stops any or all of the SQL Server services. The command stops the instance of the SQL Server. By default, the instance name is `MSSQLSERVER`. If you are using a non-default instance name, replace `MSSQLSERVER` with `MSSQL$instancename`.

Authorization role

None

Arguments**MS SQL Server Windows Service Name**

The Windows service name of the MS SQL Server instance you want to stop. The default service name that the default SQL Server instance uses is `MSSQLSERVER`. This Take Action uses that service name as the default. To stop a named MS SQL Server instance, set the value of this argument to the service name that is associated with the named SQL Server instance.

Destination system

Managed system

Usage notes

In the command, you can add `/Y` to stop dependent services too if they are running (for example, `net stop MSSQSERVER /Y`). If the dependent service `SQLServerAgent` is running and `/Y` is not added in the command, none of the SQL services is stopped.

Update Space Usage Information action

Corrects inaccuracies in the `sysindexes` table that can cause incorrect space usage reports by the `sp_spaceused` system stored procedure.

Authorization role

Microsoft SQL Server permissions are `sysadmin` or `db_owner`.

Arguments**SQL Server Name (sql_server_name)**

SQL Server instance name

Database Name (database_name)

Name of the database

Table Name (table_name)

Updates the space usage information for a specific table or for all tables. If blank, this argument updates indexes for all tables.

Index Name (index_name)

Updates the space usage information for indexes with a particular name, which you specify in this field, or all indexes. If this argument is blank, it updates all indexes. The Table 12 table shows how entries in the Table Name and Index Name arguments determine which indexes and tables are updated.

Table 12. Index and table updates

Index name	Table Name	Outcome
Exists	Exists	Updates only the specified index on the specified table.
Exists	Does not exist	Error. Specified index does not exist.
Exists	Blank (use double quotation marks "")	Updates all indexes on the specified table.
Does not exist	Exists	Error. Specified table does not exist.
Does not exist	Does not exist	Error. Specified table does not exist.
Does not exist	Blank (use double quotation marks "")	Error. Specified table does not exist.
Blank (use double quotation marks "")	Exists	Updates all indexes that share the specified name.
Blank (use double quotation marks "")	Does not exist	Error. Specified table does not exist.
Blank (use double quotation marks "")	Blank (use double quotation marks "")	Updates all indexes.

SQL Server ID (sql_server_id)

For information about specifying the SQL Server login ID, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

SQL Server Password (sql_server_password)

For information about specifying the SQL Server login password, see the "User ID and password combinations" table in the "Take Action commands reference" chapter of the *IBM Tivoli Monitoring for Microsoft Applications: Microsoft SQL Server User's Guide*.

Destination system

Managed system

Usage notes

In the command, you can add /Y to stop dependent services too if they are running (for example, net stop MSSQSERVER /Y). If the dependent service SQLServerAgent is running and /Y is not added in the command, none of the SQL services is stopped.

Chapter 7. Policies reference

Policies are an advanced automation technique for implementing more complex workflow strategies than you can create through simple automation.

A *policy* is a set of automated system processes that can perform actions, schedule work for users, or automate manual tasks. You use the Workflow Editor to design policies. You control the order in which the policy executes a series of automated steps, which are also called activities. Policies are connected to create a workflow. After an activity is completed, Tivoli Enterprise Portal receives return code feedback and advanced automation logic responds with subsequent activities prescribed by the feedback.

Note: The predefined policies provided with this monitoring agent are not read-only. Do not edit these policies and save over them. Software updates will write over any of the changes that you make to these policies. Instead, clone the policies that you want to change to suit your enterprise.

More information about policies

For more information about working with policies, see the *IBM Tivoli Monitoring User's Guide*.

For information about using the Workflow Editor, see the *IBM Tivoli Monitoring Administrator's Guide* or the Tivoli Enterprise Portal online help.

For a list of the policies for this monitoring agent and a description of each policy, see Predefined policies section and the information for each individual policy.

Predefined policies

You can create policies for any agent. The Microsoft SQL Server agent provides one predefined policy.

Policy descriptions

The policy descriptions provide information about the actions that you can take when situations are triggered in your network.

MS_SQL_DB_free_space_critical

This policy provides actions to take on low free space on the MS SQL server. This policy is triggered by the MS_SQL_DB_Freespace_Critical situation. When triggered, the workflow requires that you select one of the following actions:

- Run the **koqignor** command.
- Run the **koqmail** command as follows:

Where:

database user id

Connection user ID

password

Password as a sysadmin authority

email id

Email address of the database administrator

Chapter 8. Tivoli Common Reporting for the monitoring agent

You can use the Tivoli Common Reporting application to generate reports for the Microsoft SQL Server agent. For more information about Tivoli Common Reporting, such as prerequisites, importing reports, and running reports, see the *IBM Tivoli Monitoring Administrator's Guide V6.2.3*.

IBM Tivoli Monitoring for Microsoft Applications V6.3 reports are historical reports that include summarized data that is collected in the Tivoli Data Warehouse. These reports are built to run only against the IBM Tivoli Monitoring for Microsoft SQL Server agent V6.3.

The reports can be administered and run on Tivoli Common Reporting V1.3, V2.1, and V2.1.1, Fix Pack 6.

For more information about Tivoli Common Reporting, see the developerWorks® website (<http://www.ibm.com/developerworks/spaces/tcr>).

This version of Tivoli Common Reporting includes Cognos Business Intelligence and Reporting V8.4.

Cognos-based report packages

The Cognos-based Tivoli Common Reporting tool is used to create, view, and manage reports for the Tivoli group of products.

You can use the Cognos reports to analyze resource information such as availability, utilization, performance, and so on. With Cognos reports, you can evaluate the key metrics of the computers that are on the managed environment of your organization.

You can use the Tivoli Common Reporting tool to:

- Create custom reports by using the drag-and-drop feature integrated with the web-based editor.
- Schedule, share, secure, and administer reports in a single interface.
- Save the report in HTML, PDF, Excel, XML, or CSV file formats.
- Share reports by email or save the reports in a file system for later use.

The following databases are supported for all reports:

- DB2® V9.5 and V9.7, Fix Pack 2
- Oracle 10g and 11g
- SQL Server 2005 and 2008

You can use Tivoli Common Reporting V1.3, V2.1, and V2.1.1, Fix Pack 6 software that is shipped with *IBM Tivoli Monitoring V6.2.3*, or later to administer, run, and edit Cognos reports. For more information about Tivoli Common Reporting, see the Tivoli Common Reporting Information Center (http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ic-home.html).

Prerequisites

Before you run the Cognos-based reports, ensure that the prerequisites that are required for installing and running Tivoli Common Reporting packages are met.

Procedure

1. Install Tivoli Common Reporting. For more information, see the documentation in the IBM Tivoli Common Reporting Information Center (http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ic-home.html).

To ensure that Tivoli Common Reporting is running, go to:
https://machine_name:port_number/ibm/console/.

2. Obtain the reports package from the product media and extract the package.
The Cognos reports are available in the following package: Product Media root/REPORTS/KOQ/ITCAMMA_SQLServer_V630_Cognos_Reports.
3. Copy this package in a directory on any drive of the same computer where the Tivoli Common Reporting Server is installed, and extract the package in the same directory.
The directory contents include:
 - Database scripts required to prepare the Tivoli Data Warehouse for Cognos reports.
 - A report installer that imports the reports into Tivoli Common Reporting and sets up the database connection.
4. Open the setup.batch file, and verify that the specified drive is the same drive where the Tivoli Common Reporting Server is installed. If the specified drive is not correct, edit the setup.batch file to specify the correct drive.
For example, change the TCR_DEFAULT_DIR=C:\IBM\tivoli\tip to TCR_DEFAULT_DIR=F:\IBM\tivoli\tip if the Tivoli Common Reporting Server is installed on the F drive.
5. Configure historical collection for Microsoft SQL Server agent V6.3 and the IBM Tivoli Warehouse Summarization and Pruning Agent. After IBM Tivoli Monitoring V6.2.3, Fix Pack 1, or later is installed and the Microsoft SQL Server agent V6.3 is installed and configured, configure historical collection. Also, configure the Warehouse Summarization and Pruning agent with or without shifts enabled. For more information about how to enable historical collection and configure the Warehouse Summarization and Pruning agent in IBM Tivoli Monitoring, see the following documentation:
 - http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/adminuse/history_manage_intro.htm
 - http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/adminuse/history_manage_collection.htm

Note: Historical collection and daily and hourly summarization must be enabled for all attribute groups of Microsoft SQL Server agent V6.3.

To ensure that the required views are present, run the following query against the Tivoli Data Warehouse:

- DB2: Select distinct "VIEWNAME" from SYSCAT.VIEWS where "VIEWNAME" like '%V'
 - Oracle: Select distinct "OBJECT_NAME" from OBJ where "OBJECT_TYPE" like '%V'
 - MS SQL Server: Select distinct "NAME" from SYS.OBJECTS where "TYPE_DESC" like '%V'
6. Prepare the Tivoli Data Warehouse to support Cognos dimensions.
Preparing the Tivoli Data Warehouse for Tivoli Common Reporting includes creating the IBM_TRAM dimensions, which are required for running the Cognos reports and using the data models.
See "Creating shared dimension tables and populating the time dimensions table" in the *IBM Tivoli Monitoring Administrator's Guide* at the IBM Tivoli Monitoring V6.2.3, Fix Pack 1 Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/adminuse/tcr_reports_dimensionsshared.htm).
 7. Connect to the Tivoli Data Warehouse by using the database client over Open Database Connectivity (ODBC).
Cognos uses ODBC to connect to the database. Install a database client on the Tivoli Common Reporting Server and connect it to the Tivoli Data Warehouse.
See "Connecting to the Tivoli Data Warehouse using the database client over ODBC" in the *IBM Tivoli Monitoring Administrator's Guide* at the IBM Tivoli Monitoring V6.2.3, Fix Pack 1 Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/adminuse/tcr_tdwconnect.htm).

Important: All prerequisites described here must be met or the reports cannot run.

Importing Cognos report packages

You must import the Cognos report package that contains the Microsoft SQL Server agent data model and reports into Tivoli Common Reporting. Before you import the Cognos report package, ensure that all the prerequisites have been met. Otherwise, reports cannot run.

For Tivoli Common Reporting V1.3, use one of the following options to import the Cognos report package:

- Importing by running the reports installer – For information about running the reports installer to import the Cognos report package, see "Installing and running IBM Cognos reports" in the *IBM Tivoli Monitoring Administrator's Guide* at the IBM Tivoli Monitoring V6.2.2, Fix Pack 2 Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/tcr_install_itm.htm).
- Importing manually – For information about manually importing the reports package, see "Importing report packages" at the IBM Tivoli Systems Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr_cog.doc/tcr_cog_import.html).

For Tivoli Common Reporting V2.1, or later, see "Importing report packages" at the Tivoli Common Reporting Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/tcr_import.html).

If you are using a schema other than ITMUSER, you must complete the following steps for the data model and reports to run:

1. Verify that you have completed the following steps:
 - a. Altered the schema name in the database scripts when you prepared the Tivoli Data Warehouse for Cognos dimensions in Step 6 on page 256 of "Prerequisites" on page 255.
 - b. Installed the reports package by using the Reports installer as explained in "Importing reports by using the report installer" topic in the *IBM Tivoli Monitoring Administrator's Guide*. To verify that you have installed the reports package, go to Tivoli Common Reporting and make sure you can see "ITCAMMA SQLServer V630 Cognos Reports" in the Public Folders of IBM Cognos Connection.
2. Install and configure the Cognos Framework Manager, which is the data modeling tool. See the instructions at the Tivoli Common Reporting Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/tcr_import.html).
3. Open the Framework Manager. Select **File > Open**. Browse to the extracted ITCAMMA SQLServer V630 Cognos Reports reports package. Browse to the model folder and select the ITCAMMA SQLServer V630 Cognos Reports.cpf file.
4. If you are prompted to enter login credentials, enter your tipadmin user ID and password.
5. After the IBM Tivoli Monitoring for SQL Servers data model in the Framework Manager opens, expand **Data Sources** under **Tivoli Reporting and Analytics Model** in the Project Viewer.
6. Select **TDW** under **Data Sources**.
7. When you select **TDW**, the Properties view is updated with information about the Tivoli Data Warehouse data source. By default, the Properties view is located at the bottom center of the screen. If the Properties view is not visible, select **View > Properties**.
8. In the Properties, edit the Schema field. Change it from ITMUSER to your schema name.
9. Save the project.
10. In the Project Viewer, expand **Packages**.
11. Right-click **ITCAMMA SQLServer V630 Cognos Reports**.
12. Select **Publish Packages**. The Publish Wizard opens.
13. Keep the default selection and click **Next**.
14. Click **Next** on the next screen.

15. Clear the **Verify the package before publishing** check box.
16. Click **Publish**. A window is displayed that alerts you that A package with that name already exists and asks Do you want to publish this package?
17. Click **Yes**.
18. Go back to Tivoli Common Reporting and check if the Modified field of "ITCAMMA SQLServer V630 Cognos Reports" in the Public Folders of IBM Cognos Connection shows the time of publishing.

After completing these steps, you can run any report from the IBM Tivoli Monitoring for ITCAMMA SQLServer V630 Cognos Reports package.

Cognos data models and reports

When all the prerequisites are met, you can use the Tivoli Integrated Portal interface to create, modify, and manage Cognos reports.

In Tivoli Common Reporting, the historical data that is collected by the agent is used to build ad hoc reports and queries. The package that you import into Tivoli Common Reporting contains a Cognos data model. All the reports that you create in Tivoli Common Reporting are based on the data model.

Cognos data models are virtual star schema models that contain facts and dimensions. Facts are measurable quantities that can be aggregated, such as CPU utilization and number of processors. Dimensions are the main identifiers by which facts can be grouped, aggregated, and organized. For example, time and server are dimensions by which the fact CPU utilization can be grouped.

The facts in the data model are organized into the folders by their summarization type, such as Daily and Hourly. The Daily and the Hourly folders contain attribute groups that correspond to tables or views in the data warehouse. Each attribute group contains a group of facts and measures.

The data model is built on top of the data warehouse to organize data. The data model contains the Tivoli Reporting and Analytics Model (TRAM) Shared Dimensions, which are shared across Tivoli by products such as Time.

To create reports in Tivoli Common Reporting tool, you can use one of the following report authoring tools:

- **Query Studio:** A web-based product that is used mostly for ad hoc reporting. Users can create simple queries and reports, and apply basic formatting to the reports.
- **Report Studio:** A web-based tool that is used by technical users and professional report writers for creating advanced reports. Users can retrieve data from multiple databases and create sophisticated reports that have multiple pages and multiple queries.

Reports for the monitoring agent

You can verify whether the reporting functionality is installed and configured correctly by running the reports for the Microsoft SQL Server agent. The Microsoft SQL Server agent reporting package that you imported into Tivoli Common Reporting includes 10 reports. By using these reports, you can monitor the reporting activity and see what a typical Cognos report includes. These reports are available in the Common Reporting panel in Tivoli Common Reporting. The following tables describe these reports.

Table 13. Top Processes by CPU report

Name	Top Processes by CPU report
Description	This report displays the relative percentage of the CPU that is used by the selected server processes compared to other server processes and the percentage of processes that are being blocked.

Table 13. Top Processes by CPU report (continued)

Purpose	You can use this report to provide information about the process details of the selected SQL Server.
Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Summarization Selection</p> <p>Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly, Quarterly, Yearly, or Default from the list. If you select the Default option, the summarization type is calculated according to the number of days in the date range.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Number of Processes to Display Select the processes from the list</p>
Tables/Views Used	<ul style="list-style-type: none"> • MS_SQL_Process_Detail_%V • MS_SQL_Process_Summary_%V
Output	<p>The Top Processes by Current CPU % Used chart displays the current percentage of CPU that is used for different processes on a selected server.</p> <p>Note: You can change the number of processes to display to any positive integer. However, the default value is 5.</p> <p>The Processes Blocked % chart displays the percentage of processes that are blocked on a server for a selected period.</p>
Usage	The administrators can use this report to observe the relative percentage of the CPU that is used and the percentage of processes that are blocked.

Table 14. Server Transactions Details report

Name	Server Transactions Details report
Description	This report displays the replication latency between transactions, the number of transactions that occurred in current interval, and the number of active transactions for the database.

Table 14. Server Transactions Details report (continued)

Purpose	You can use this report to observe the performance details of a database.
Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Summarization Selection</p> <p>Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly, Quarterly, Yearly, or Default from the list. If you select the Default option, the summarization type is calculated according to the number of days in the date range.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Database Name Select the name of the database from the list.</p>
Tables or views used	MS_SQL_Database_Detail_%V
Output	<p>The AVG Replication Latency between Transactions chart displays average replication latency for the selected databases on a server.</p> <p>The AVG Transactions Per Sec chart displays the number of transactions per second for selected databases on a server.</p> <p>The Number of Active Transactions chart displays the maximum number of active transactions for the selected databases on a server.</p> <p>The Database Details table that follows the preceding charts, displays the same information</p>
Usage	The IT administrator can use this report to observe hindrance in the replication Reports 19 process, transactions started per second in the current interval, and the number of active transactions for the database.

Table 15. Wait Time and Requests of Locks report

Name	Wait Time and Requests of Locks report
-------------	--

Table 15. Wait Time and Requests of Locks report (continued)

Description	This report displays the total wait time for locks in the last second and the total number of locks that block the other processes.
Purpose	You can use this report to provide information about the lock requests and wait time.
Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Summarization Selection</p> <p>Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly, Quarterly, Yearly, or Default from the list. If you select the Default option, the summarization type is calculated according to the number of days in the date range.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p>
Tables or views used	<ul style="list-style-type: none"> • MS_SQL_Lock_Resource_Type_Summary_%V • MS_SQL_Server_Summary_%V
Output	<p>The Average Lock Wait Time (in MS) Over a Period of Time line chart displays the average wait time (in milliseconds) for locks in the last second over a period of time.</p> <p>The Total Number of Blocking Lock Requests line chart displays the number of lock requests that block the other processes over a period of time.</p>
Usage	The IT administrators can use this report to observe the total wait time for locks in the last second and the number of locks that block the other processes over a period of time.

Table 16. Space Utilization by Database report

Name	Space Utilization by Database report
Description	This report displays the free space (in MB), the percentage of free space, the amount of space used as percentage of the total space, the percentage of free space in the transaction log, and the percentage of transaction log that contains the database.

Table 16. Space Utilization by Database report (continued)

Purpose	You can use this report to provide the space details of a database.
Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select an end date from the calendar and an end time from the time widget. You must select both date and time.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Database Name Select the name of the database from the list.</p> <p>Attribute Selection Select the space used or the space available.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Thresholds You can enter the threshold values that are represented by the status. The values represent the default threshold values for the Space Used attribute.</p>
Tables or views used	MS_SQL_Database_Detail_HV
Output	<p>The Average Lock Wait Time (in MS) Over a Period of Time chart displays the average wait time (in MS) for locks of the selected database over a period of time.</p> <p>The Average Database Free Space % Per Hour Over a Period of Time chart displays the percentage of the used space (in megabytes) of the selected database over a period of time.</p> <p>The Average Log Free Space % Per Hour Over a Period of Time chart displays the percentage of the used space of the selected database log over a period of time.</p>
Usage	The IT administrators can use this report to observe the used space for data and used space for log, for a selected database.

Table 17. Free Space Availability report

Name	Free Space Availability report
Description	This report displays the lowest percentage of free space in the data segment on a device that is allocated to a database and the lowest percentage of free space in the log segment on a device that is allocated to a transaction log.

Table 17. Free Space Availability report (continued)

Purpose	You can use this report to provide information about the free space in a data segment and a log segment.
Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select an end date from the calendar and an end time from the time widget. You must select both date and time.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Thresholds You can enter the threshold values that are represented by the status. The values represent the default threshold values.</p>
Tables or views used	MS_SQL_Database_Summary_HV
Output	<p>The Minimum Data Free Space % per hour over a period of time chart displays the minimum percentage of data free space on the selected server.</p> <p>The Minimum Log Free Space % per hour over a period of time chart displays the minimum percentage of log free space that is allocated to a database transaction log for a selected server.</p>
Usage	The IT administrator can use this report to observe the free space in a database and in a log.

Table 18. Database Status Availability report

Name	Database Status Availability report
Description	This report displays the availability of the database for usage.
Purpose	You can use this report to provide information about the status of a database.

Table 18. Database Status Availability report (continued)

<p>Parameters</p>	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Database Name Select the name of the database from the list.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Thresholds You can enter the threshold values that are represented by the status. The values represent the default threshold values.</p>
<p>Tables or views used</p>	<p>MS_SQL_Database_Detail_HV</p>
<p>Output</p>	<p>The Average Available % of Database per hour over a period of time chart displays the percentage of the availability of the database.</p>
<p>Usage</p>	<p>The IT administrators can use this report to observe the status of a database on a selected SQL Server.</p>

Table 19. CPU Utilization report

<p>Name</p>	<p>CPU Utilization report</p>
<p>Description</p>	<p>This report displays the percentage of the CPU memory that is used by all the processes including the SQL Server processes .</p>
<p>Purpose</p>	<p>You can use this report to provide information about the CPU utilization of all the processes including the SQL Server processes.</p>

Table 19. CPU Utilization report (continued)

Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Thresholds You can enter the threshold values that are represented by the status. The following values represent the default threshold values:</p> <ul style="list-style-type: none"> • Good: Less than or equal to 20 • Fair: Between 20 and 30 • Warning: Between 30 and 50 • Bad: Between 50 and 80 • Critical: Above 80
Tables or views used	MS_SQL_Server_Summary_HV
Output	<p>The Average Total CPU % Used Per Hour by SQL Processes over a period of time chart displays the percentage of the CPU time that is utilized by the SQL processes.</p> <p>The Average Total CPU % Used Per Hour by All Processes over a period of time chart displays the total percentage of the CPU time that is utilized by the operating system.</p>
Usage	The IT administrator can use this report to compare the relative CPU percentage used by the processes.

Table 20. Processes Details report

Name	Processes Details report
Description	This report displays the total number of processes that are started and are blocked since the server was started.
Purpose	You can use the report to provide information about the processes on a selected SQL Server.

Table 20. Processes Details report (continued)

<p>Parameters</p>	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Thresholds You can enter the threshold values that are represented by the status. The following values represent the default threshold values:</p> <ul style="list-style-type: none"> • Good: Less than or equal to 20 • Fair: Between 20 and 30 • Warning: Between 30 and 50 • Bad: Between 50 and 80 • Critical: Above 80
<p>Tables or views used</p>	<p>MS_SQL_Process_Summary_HV</p>
<p>Output</p>	<p>The Average Total Processes Blocked per hour over a period of time chart displays the total processes that are blocked for a selected period.</p> <p>The Average Processes Blocked (%) per hour over a period of time chart displays the percentage of processes that are blocked for a selected period.</p>
<p>Usage</p>	<p>The IT administrator can use this report to compare the number of blocked processes with the number of total processes that are present.</p>

Table 21. Server Lock Details report

<p>Name</p>	<p>Server Lock Details report</p>
<p>Description</p>	<p>This report displays the number of lock conversions per second, the number of lock requests per second that timed out, total wait time for locks in the last second, and the total number of current locks that are blocking other processes.</p>
<p>Purpose</p>	<p>You can use this report to provide information about locks on the different resource types.</p>

Table 21. Server Lock Details report (continued)

<p>Parameters</p>	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Summarization Selection</p> <p>Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly, Quarterly, Yearly, or Default from the list. If you select the Default option, the summarization type is calculated according to the number of days in the date range.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p>
<p>Tables or views used</p>	<ul style="list-style-type: none"> • MS_SQL_Lock_Resource_Type_Summary_%V • MS_SQL_Server_Summary_%V
<p>Output</p>	<p>The Number of Lock Requests/Timeouts in a Sec line chart displays the lock requests per second and lock timeouts per second.</p> <p>The Average Total Wait Time in Last Sec line chart displays total wait time (in milliseconds) for locks in the last second for the selected server.</p> <p>The table following the chart displays bar charts for average blocking requests and maximum lock wait time. Click the server name in the table to drill down to the Wait Time and Requests of Locks report.</p>
<p>Usage</p>	<p>The IT administrator can use this report to observe the number of new locks and lock conversions, and the number of timed out lock requests. The total wait time (in milliseconds) for locks to release in the last second and the number of blocking locks indicate that the processes are held up because of lock contention. Click the server name in the table to drill down to the Wait Time and Requests of Locks report.</p>

Table 22. Servers with Most Failed Jobs report

<p>Name</p>	<p>Servers with Most Failed Jobs report</p>
<p>Description</p>	<p>This report provides the job summary of the top servers that display the highest number of failed jobs since the agent was started.</p>

Table 22. Servers with Most Failed Jobs report (continued)

Purpose	You can use this report to provide information about locks on the different resource types.
Parameters	<p>Date Range</p> <p>Report period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days, and so on. You can also enter a start date, an end date, and the time for the reporting period.</p> <p>Start Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>End Date Select a start date from the calendar and the start time from the time widget. You must select both date and time.</p> <p>Summarization Selection</p> <p>Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly, Quarterly, Yearly, or Default from the list. If you select the Default option, the summarization type is calculated according to the number of days in the date range.</p> <p>Shift Period If shift period is enabled, the hourly tables display the Shift Period as 1 (off-peak hours) or 2 (peak hours), depending on the peak and off-peak hours that are configured in the data warehouse. In the daily tables, the value 1 corresponds to peak hours, the value 2 corresponds to off-peak hours, and the value -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.</p> <p>Vacation Period If Vacation Period is not enabled, the default value is -1 (All Days). You can enter the value 0 (Work Days) or 1 (Vacation Days), if Vacation Period is enabled.</p> <p>Display Options</p> <p>Server Name Select the name of the server from the Server Name list.</p> <p>Number of Servers to Display Select the servers from the list.</p>
Tables or views used	MS_SQL_Job_Summary_%V
Output	<p>The Number of Lock Requests/Timeouts in a Sec line chart displays the lock requests per second and lock timeouts per second.</p> <p>The table following chart displays the same data.</p>
Usage	The IT administrator can use this report to observe and compare the performance of different SQL servers by using the number of failed jobs since the agent started.

Chapter 9. Troubleshooting

This chapter explains how to troubleshoot the IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent. Troubleshooting is the process of determining why a certain product is malfunctioning.

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in the Prerequisites topic for the agent in the information center.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information. Also see “Support information” on page 307 for other problem-solving options.

Gathering product information for IBM Software Support

Before contacting IBM Software Support about a problem you are experiencing with this product, gather the following information that relates to the problem:

Table 23. Information to gather before contacting IBM Software Support

Information type	Description
Log files	Collect trace log files from failing systems. Most logs are located in a logs subdirectory on the host computer. See “Trace logging” on page 270 for lists of all trace log files and their locations. See the <i>IBM Tivoli Monitoring User’s Guide</i> for general information about the IBM Tivoli Monitoring environment.
Microsoft SQL Server information	<ul style="list-style-type: none">• Version number and patch level (For example: start Microsoft SQL Server Management Studio. Select the About... item in the Help menu to see this information.)• Sample application data file (if monitoring a file)
Operating system	Operating system version number and patch level (For example, you can right-click the My Computer icon and select Properties to obtain this information.)
Messages	Messages and other information displayed on the screen
Version numbers for IBM Tivoli Monitoring	Version number and patch level of the following members of the monitoring environment: <ul style="list-style-type: none">• IBM Tivoli Monitoring• IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent
Screen captures	Screen captures of incorrect output, if any.

Built-in troubleshooting features

The primary troubleshooting feature in the IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent is logging. *Logging* refers to the text messages and trace data generated by the IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent. Messages and trace data are sent to a log file.

Trace data captures transient information about the current operating environment when a component or application fails to operate as designed. IBM Software Support personnel use the captured trace information to determine the source of an error or unexpected condition. See “Trace logging” on page 270 for more information.

Problem classification

The following types of problems might occur with the IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent:

- Installation and configuration
- General usage and operation
- Display of monitoring data
- Take Action commands

This chapter provides symptom descriptions and detailed workarounds for these problems, as well as describing the logging capabilities of the monitoring agent. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Trace logging

Trace logs capture information about the operating environment when component software fails to operate as intended. The principal log type is the RAS (Reliability, Availability, and Serviceability) trace log. These logs are in the English language only. The RAS trace log mechanism is available for all components of IBM Tivoli Monitoring. Most logs are located in a logs subdirectory on the host computer. See the following sections to learn how to configure and use trace logging:

- “Principal trace log files” on page 271
- “Viewing trace logs” on page 274
- “Enabling detailed tracing in the collector trace log” on page 274
- “RAS trace parameters” on page 275

Note: The documentation refers to the RAS facility in IBM Tivoli Monitoring as "RAS1".

IBM Software Support uses the information captured by trace logging to trace a problem to its source or to determine why an error occurred. The default configuration for trace logging, such as whether trace logging is enabled or disabled and trace level, depends on the source of the trace logging. Trace logging is always enabled.

Overview of log file management

Table 24 on page 272 provides the names, locations, and descriptions of RAS1 log files. The log file names adhere to the following naming convention:

```
hostname_product_instance_program_timestamp-nn.log
```

where:

- *hostname* is the host name of the computer on which the monitoring component is running.
- *product* is the two-character product code. For Microsoft SQL Server agent, the product code is oq.
- *instance* is the name of a database instance that is being monitored.
- *program* is the name of the program being run.
- *timestamp* is an 8-character hexadecimal timestamp representing the time at which the program started.
- *nn* is a rolling log suffix. See “Examples of trace logging” for details of log rolling.

Examples of trace logging

For example, if a Microsoft SQL Server database dbinst02 is running on computer server01, the RAS log file for the Microsoft SQL Server agent might be named as follows:

```
server01_oq_dbinst02_koqagent_437fc59-01.log
```

For long-running programs, the *nm* suffix is used to maintain a short history of log files for that startup of the program. For example, the koqagent program might have a series of log files as follows:

```
server01_oq_dbinst02_koqagent_437fc59-01.log
server01_oq_dbinst02_koqagent_437fc59-02.log
server01_oq_dbinst02_koqagent_437fc59-03.log
```

As the program runs, the first log (*nm=01*) is preserved because it contains program startup information. The remaining logs *roll*. In other words, when the set of numbered logs reach a maximum size, the remaining logs are overwritten in sequence. Each time a program is started, a new timestamp is assigned to maintain a short program history. For example, if the Microsoft SQL Server agent is started twice, it might have log files as follows:

```
server01_oq_dbinst02_koqagent_437fc59-01.log
server01_oq_dbinst02_koqagent_437fc59-02.log
server01_oq_dbinst02_koqagent_437fc59-03.log
```

```
server01_oq_dbinst02_koqagent_537fc59-01.log
server01_oq_dbinst02_koqagent_537fc59-02.log
server01_oq_dbinst02_koqagent_537fc59-03.log
```

Each program that is started has its own log file. For example, the Microsoft SQL Server agent would have agent logs in this format:

```
server01_oq_dbinst02_koqagent_437fc59-01.log
```

Other logs, such as logs for collector processes and Take Action commands, have a similar syntax, as in the following example:

```
server01_oq_dbinst02_koqsql_447fc59-01.log
```

where **koqsql** is the name of a program.

Note: When you communicate with IBM Software Support, you must capture and send the RAS1 log that matches any problem occurrence that you report.

Principal trace log files

Table 24 on page 272 contains locations, file names, and descriptions of trace logs that can help determine the source of problems with agents.

Table 24. Trace log files for troubleshooting agents

System where log is located	File name and path	Description
<p>On the computer that hosts the monitoring agent</p> <p>See Definitions of variables for descriptions of the variables in the file names in column two.</p>	<p>The RAS1 log files are named <i>hostname_oq_instance_program_timestamp-nn.log</i> and are located in the <i>install_dir\tmaitm6\logs</i> path:</p> <p>Note: File names for RAS1 logs include a hexadecimal time stamp.</p>	<p>Traces activity of the monitoring agent.</p> <p>Note: Other logs, such as logs for collector processes and Take Action commands (if available), have a similar syntax and are located in this directory path.</p>
	<p>The collector trace log file is named <i>hostname_oq_instance_col.out</i>. The file is located in the <i>install_dir\tmaitm6\logs</i> path.</p> <p>Note: A <i>hostname_oq_instance_col.out</i> log file also exists and contains stderr and stdout messages.</p>	<p>A collector trace log file is produced by Microsoft SQL Server agent. New log files are created each time the agent is started. One backup copy, renamed from *.out to *.ou1, is kept for historical purposes. The log file is not pruned. See “Enabling detailed tracing in the collector trace log” on page 274 to learn how to change the default behavior of this type of logging.</p>
	<p>The agent trace log file is named <i>hostname_oq_instance_agt.out</i>. The file is located in the <i>install_dir\tmaitm6\logs</i> path.</p>	<p>An agent trace log file is produced by Microsoft SQL Server agent. New log files are created each time the agent is started. One backup copy, renamed from *.out to *.ou1, is kept for historical purposes.</p>
	<p>The *.LG0 file is located in the <i>install_dir\tmaitm6\logs</i> path.</p>	<p>A new version of this file is generated every time the agent is restarted. IBM Tivoli Monitoring generates one backup copy of the *.LG0 file with the tag .LG1. View .LG0 to learn the following details regarding the current monitoring session:</p> <ul style="list-style-type: none"> • Status of connectivity with the monitoring server. • Situations that were running. • The success or failure status of Take Action commands.

Table 24. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
<p>On the Tivoli Enterprise Portal Server</p> <p>See Definitions of variables for descriptions of the variables in the file names in column two.</p>	<p>On UNIX: The <code>install_dir/logs</code> file in the <code>install_dir/logs</code> path.</p> <p>On Windows: The IBM Tivoli Monitoring <code>timestamp.log</code> file in the <code>install_dir\InstallITM</code> path. Unlike RAS1 log files, the name of the file displays a <i>decimal</i> time stamp. *</p> <p>The RAS1 log file is named <code>hostname_ms_timestamp-nn.log</code> and is located in the following path:</p> <ul style="list-style-type: none"> • On Windows: <code>install_dir\logs</code> • On UNIX: <code>install_dir/logs</code> <p>Note: File names for RAS1 logs include a hexadecimal time stamp</p> <p>Also on UNIX, a log with a decimal time stamp is provided: <code>hostname_ms_timestamp.log</code> and <code>hostname_ms_timestamp.pidnnnnn</code> in the <code>install_dir/logs</code> path, where <code>nnnnn</code> is the process ID.</p>	<p>Provides details about products that are installed.</p> <p>Note: Trace logging is enabled by default. A configuration step is not required to enable this tracing.</p> <p>Traces activity on the monitoring server.</p>
<p>On the computer that hosts the Warehouse Proxy agent</p>	<p>The <code>Warehouse_Configuration.log</code> file is located in the following path on Windows: <code>install_dir\InstallITM\Warehouse_Configuration.log</code>. The <code>Warehouse_Configuration.log</code> file is located in the following path on Unix: <code>install_dir/logs/KHDCFG.log</code>.</p>	<p>Provides details about the configuration of data warehousing for historical reporting.</p>
<p>On the Tivoli Enterprise Portal Server</p> <p>See Definitions of variables for descriptions of the variables in the file names in column two.</p>	<p>The RAS1 log file is named <code>hostname_cq_timestamp-nn.log</code> and is located in the following path:</p> <ul style="list-style-type: none"> • On Windows: <code>install_dir\logs</code> • On UNIX: <code>install_dir/logs</code> <p>Note: File names for RAS1 logs include a hexadecimal time stamp</p> <p>Also on UNIX, a log with a decimal time stamp is provided: <code>hostname_cq_timestamp.log</code> and <code>hostname_cq_timestamp.pidnnnnn</code> in the <code>install_dir/logs</code> path, where <code>nnnnn</code> is the process ID.</p>	<p>Traces activity on the portal server.</p>
	<p>The <code>TEPS_ODBC.log</code> file is located in the following path on Windows: <code>install_dir\InstallITM</code>.</p>	<p>When you enable historical reporting, this log file traces the status of the warehouse proxy agent.</p>
<p>Definitions of variables for RAS1 logs:</p> <ul style="list-style-type: none"> • <code>hostname</code> is the host name of the computer on which the agent is running. • <code>install_dir</code> represents the directory path where you installed the IBM Tivoli Monitoring component. <code>install_dir</code> can represent a path on the computer that hosts the monitoring server, the monitoring agent, or the portal server. • <code>product</code> is the two character product code. For Microsoft SQL Server agent, the product code is <code>oq</code>. • <code>instance</code> refers to the name of the database instance that you are monitoring. • <code>program</code> is the name of the program being run. • <code>timestamp</code> is an eight-character hexadecimal time stamp representing the time at which the program started. • <code>nn</code> is a rolling log suffix. See “Examples of trace logging” on page 270 for details of log rolling. 		

See the *IBM Tivoli Monitoring Installation and Setup Guide* for more information on the complete set of trace logs that are maintained on the monitoring server.

Viewing trace logs

About this task

Typically IBM Software Support applies specialized knowledge to analyze trace logs to determine the source of problems. However, you can open trace logs in a text editor to learn some basic facts about your IBM Tivoli Monitoring environment. Use a text editor that is enabled for UTF-8, such as Notepad.

On Windows, you can use the following alternate method to view trace logs:

Procedure

1. In the Windows **Start** menu, choose **Program Files > IBM Tivoli Monitoring > Manage Tivoli Monitoring Service**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
2. Right-click a component and select **Advanced > View Trace Log** in the pop-up menu. The program displays the Select Log File window that lists the RAS1 logs for the monitoring agent.
3. Select a log file from the list and click **OK**. You can also use this viewer to access remote logs.

Note: The viewer converts time stamps in the logs to a readable format.

Enabling detailed tracing in the collector trace log

Collector trace logs are generated by a database monitoring agent. You can use the following options to modify collector trace logging.

Option one: modify the settings file

This section provides the procedure to modify the settings file.

Procedure

Complete the following steps to enable detailed logging for the collector trace log:

1. Open the `koqtrac.ct1` file. The file is located in the following path: `install_dir\tmaitm6`, where `install_dir` is the location of the IBM Tivoli Monitoring component.
2. Remove the two semicolons (`::`) from the beginning of the line that contains `trace_all`;
3. Save the file.
4. Restart the monitoring agent for the database instance that you want to trace. Logging goes to a collector trace log file that is named in Table 24 on page 272. The log file is stored in the following path: `install_dir\tmaitm6\logs` where `install_dir` is the path where you installed IBM Tivoli Monitoring.
5. When you complete an analysis of detailed trace logs, turn off detailed logging as follows, because trace logs can consume a large amount hard disk space.
 - a. Open the `koqtrac.ct1` file mentioned in Step 1.
 - b. Type two semicolons (`::`) at the beginning of the line that contains `trace_all`;
 - c. Save the file.
 - d. Restart the monitoring agent for the database instance that you have been tracing.

Option two: modify environment variables

This section proves the procedure to modify environment variables.

About this task

You can modify trace logging for the collector log by setting the following environment variables:

COLL_WRAPLINES

By default, new collector log files are created each time you start the agent or when the log file reaches the limit specified by the **COLL_WRAPLINES** environment variable. **COLL_WRAPLINES** defines the number of lines in the col.out file and is 90000 lines by default (about 2 MB).

COLL_NUMOUTBAK

By default, the product creates one backup copy of the collector log file. The file tag is changed from *.out to *.ou1. You can configure the product to create up to nine backup files by defining **COLL_NUMOUTBAK** to be an integer from one to nine. By default, the number of backup files is 5.

Set the environment variables as follows:

Procedure

1. In the Windows **Start** menu, choose **Program Files > IBM Tivoli Monitoring > Manage Tivoli Monitoring Service**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
2. Right-click the row that contains the name of the monitoring agent whose environment variables you want to set.
3. Select **Advanced > Edit Variables** in the pop-up menu.
4. If the agent is running, accept the prompt to stop the agent. A restart is necessary so that the environment variable that you create takes effect.
5. The list dialog is displayed. When only the default settings are in effect, there are no variables displayed. The variables are listed after you override them.
6. Override the variable settings as follows:
 - a. Click **Add**.
 - b. Select the variable you want to modify in the **Variable** pull-down menu. For example, select **COLL_NUMOUTBAK**.
 - c. Type a value in the **Value** field. For example, type the number 9, the maximum number of backup files that can be generated.
 - d. Select **OK** two times.
 - e. Restart the agent.

RAS trace parameters

Pinpoint a problem by setting detailed tracing of individual components of the monitoring agent and modules

See “Overview of log file management” on page 270 to ensure that you understand log rolling and can reference the correct log files when you manage log file generation.

Setting RAS trace parameters by using the GUI

On Windows systems, you can use the graphical user interface to set trace options.

About this task

The IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent uses RAS1 tracing and generates the logs described in “Principal trace log files” on page 271. The default RAS1 trace level is ERROR.

Procedure

1. Open the Manage Tivoli Enterprise Monitoring Services window.
2. Select **Advanced > Edit Trace Params**. The Tivoli Enterprise Monitoring Server Trace Parameters window is displayed.
3. Select a new trace setting in the pull-down menu in the **Enter RAS1 Filters** field or type a valid string.
 - General error tracing. KBB_RAS1=ERROR
 - Intensive error tracing. KBB_RAS1=ERROR
 - Maximum error tracing. KBB_RAS1=ERROR

Note: As this example shows, you can set multiple RAS tracing options in a single statement.

4. Modify the value for Maximum Log Size Per File (MB) to change the log file size (changes LIMIT value).
5. Modify the value for Maximum Number of Log Files Per Session to change the number of log files per startup of a program (changes COUNT value).
6. Modify the value for Maximum Number of Log Files Total to change the number of log files for all startups of a program (changes MAXFILES value).
7. Optional: Click Y (Yes) in the **KDC_DEBUG Setting** menu to log information that can help you diagnose communications and connectivity problems between the monitoring agent and the monitoring server. The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily, while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.
8. Click **OK**. You see a message reporting a restart of the monitoring agent so that your changes take effect.

What to do next

Monitor the size of the logs directory. Default behavior can generate a total of 45 - 60 MB for each agent that is running on a computer. For example, each database instance that you monitor can generate 45 - 60 MB of log data. See the "Procedure" section to learn how to adjust file size and numbers of log files to prevent logging activity from occupying too much disk space.

Regularly prune log files other than the RAS1 log files in the logs directory. Unlike the RAS1 log files that are pruned automatically, other log types can grow indefinitely.

Use collector trace logs as an additional source of troubleshooting information.

Note: The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.

Manually setting RAS trace parameters

You can manually edit the RAS1 trace logging parameters.

About this task

The Microsoft SQL Server agent uses RAS1 tracing and generates the logs described in "Principal trace log files" on page 271. The default RAS1 trace level is ERROR.

Procedure

1. Open the trace options file:
 - **Windows systems:**

`install_dir\tmaitm6\KOQENV_instance name`

- **UNIX systems:**

`install_dir /config/OQ_instance name.config`

2. Edit the line that begins with **KBB_RAS1=** to set trace logging preferences. For example, if you want detailed trace logging, set the **Maximum Tracing** option: `KBB_RAS1=ERROR (UNIT:kOQ ALL) (UNIT:kra ALL)`
3. Edit the line that begins with **KBB_RAS1_LOG=** to manage the generation of log files:
 - **MAXFILES:** The total number of files that are to be kept for all startups of a specific program. When this value is exceeded, the oldest log files are discarded. The default value is 9.
 - **LIMIT:** The maximum size, in megabytes (MB) of a RAS1 log file. The default value is 5.
4. Restart the monitoring agent so that your changes take effect.

What to do next

Monitor the size of the logs directory. Default behavior can generate a total of 45 - 60 MB for each agent that is running on a computer. For example, each database instance that you monitor can generate 45 - 60 MB of log data. See the "Procedure" section to learn how to adjust file size and numbers of log files to prevent logging activity from occupying too much disk space.

Regularly prune log files other than the RAS1 log files in the logs directory. Unlike the RAS1 log files that are pruned automatically, other log types can grow indefinitely.

Use collector trace logs as an additional source of troubleshooting information.

Note: The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.

Problems and workarounds

The following sections provide symptoms and workarounds for problems that might occur with Microsoft SQL Server agent:

- "Installation and configuration troubleshooting"
- "Agent troubleshooting" on page 287
- "Workspace troubleshooting" on page 297
- "Troubleshooting for remote deployment" on page 296
- "Situation troubleshooting" on page 301
- "Take Action command troubleshooting" on page 304
- "Troubleshooting for SQL Server" on page 286

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in the Prerequisites topic for the agent in the information center for IBM Tivoli Composite Application Manager for Microsoft Applications.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Installation and configuration troubleshooting

This section provides tables that show solutions for the following types of installation, configuration, and uninstallation problems:

- Operating system problems
- Problems with database applications

Table 25. Problems and solutions for installation and configuration on UNIX and Windows systems

Problem	Solution
<p>When you upgrade to IBM Tivoli Monitoring, you might need to apply fix packs to Candle, Version 350, agents.</p>	<p>Fix packs for Candle, Version 350, are delivered as each monitoring agent is upgraded to IBM Tivoli Monitoring. Note: The IBM Tivoli Monitoring download image or CD provides application fix packs for the monitoring agents that are installed from that CD (for example, the agents for operating systems such as Windows, Linux, UNIX, and i5/OS™). The upgrade software for other agents is located on the download image or CDs for that specific monitoring agent, such as the agents for database applications.</p> <p>If you do not upgrade the monitoring agent to IBM Tivoli Monitoring, the agent continues to work. However, you must upgrade to have all the functionality that IBM Tivoli Monitoring offers.</p>
<p>Presentation files and customized Omegamon DE screens for Candle monitoring agents need to be upgraded to a new Linux on z/Series system.</p>	<p>The upgrade from version 350 to IBM Tivoli Monitoring handles export of the presentation files and the customized Omegamon DE screens.</p>
<p>The following message is displayed in the installation log for some Windows agents when upgrading from Tivoli OMEGAMON® V350: <REPLACELINE> Pair missing 1=[KBB_RAS1=ERROR] no 2, skipped.</p>	<p>There is no workaround. The previous value of KBB_RAS1 from the OMEGAMON V350 agent is used, preserving prior customer settings for this variable. The problem has no adverse effect on the installation or subsequent operation of the monitoring agent .</p>
<p>Non-ASCII characters entered into the configuration window for the monitoring agent do not show up or are not the correct characters.</p>	<p>Enter only ASCII characters into these fields.</p>
<p>(UNIX only) During a command-line installation, you choose to install a component that is already installed, and you see the following warning: WARNING - you are about to install the SAME version of "component" where <i>component</i> is the name of the component that you are attempting to install. Note: This problem affects UNIX command-line installations. If you monitor only Windows environments, you would see this problem if you choose to install a product component (for example, a monitoring server) on UNIX.</p>	<p>You must exit and restart the installation process. You cannot return to the list where you selected components to install. When you run the installer again, do not attempt to install any component that is already installed.</p>

Table 25. Problems and solutions for installation and configuration on UNIX and Windows systems (continued)

Problem	Solution
<p>A problem can arise when you install and configure a new monitoring agent to a computer where other agents are running as described in this example:</p> <ul style="list-style-type: none"> • Agents are running on computer and communicating with a Tivoli Enterprise Monitoring Server, called TEMS1. • You install a new agent on the same computer and you want this agent to communicate with a different monitoring server, called TEMS2. • When you configure the new agent to communicate with TEMS2, all the existing agents are re-configured to communicate with TEMS2. 	<p>You must reconfigure the previously existing agents to restore their communication connection with TEMS1. For example, you can right-click the row for a specific agent in the Manage Tivoli Enterprise Monitoring Services, and select Reconfigure. See the <i>IBM Tivoli Monitoring Installation and Setup Guide</i> for more information on reconfiguration.</p>
<p>Diagnosing problems with product browse settings.</p>	<p>When you have problems with browse settings, perform the following steps:</p> <ol style="list-style-type: none"> 1. Click Start > Programs > IBM Tivoli Monitoring > Manage Tivoli Enterprise Monitoring Services. The Manage Tivoli Enterprise Monitoring Services is displayed. 2. Right-click the Windows agent and select Browse Settings. A text window is displayed. 3. Click Save As and save the information in the text file. If requested, you can forward this file to IBM Software Support for analysis.
<p>A message similar to Unable to find running CMS on CT_CMSLIST in the log file is displayed.</p>	<p>If a message similar to Unable to find running CMS on CT_CMSLIST is displayed in the Log file, the agent is not able to connect to the monitoring server. Confirm the following points:</p> <ul style="list-style-type: none"> • Do multiple network interface cards (NICs) exist on the system? • If multiple NICs exist on the system, find out which one is configured for the monitoring server. Ensure that you specify the correct host name and port settings for communication in the IBM Tivoli Monitoring environment.
<p>You successfully upgraded from an OMEGAMON monitoring agent to IBM Tivoli Monitoring, Version 6.x. However, when you configure historical data collection, you see an error message that includes, Attribute name may be invalid, or attribute file not installed for warehouse agent.</p>	<p>Copy the attribute files (koq.atr) for the upgraded monitoring agent to <i>install_dir\taitm6\attrib</i> on the computer where you have installed the Warehouse Proxy. The Warehouse Proxy must be able to access the short attribute names for tables and columns. That way, if the longer versions of these names exceed the limits of the Warehouse database, the shorter names can be substituted.</p>
<p>Text for configuration functions is displayed in English instead of native languages when installing and configuring the monitoring agent. For example, when using the Manage Tivoli Enterprise Monitoring Services GUI on a Windows system.</p>	<p>None. You must complete configuration of the monitoring agent using English.</p>

Table 25. Problems and solutions for installation and configuration on UNIX and Windows systems (continued)

Problem	Solution
Installation errors occur when installing the previous 6.1.0 version of the monitoring agent after installing the 6.1.2 version of the agent.	Verify that the remote agent is not at a lower version than V6.1.2 before attempting to run the V6.1.0 remote installation with the -v option. If V6.1.2 is already present, do not force the V6.1.0 agent installation. After an inadvertent back level to V6.1.0, run the V6.1.2 installation (or higher) to upgrade the agent.
After you upgrade ITCAM for Microsoft Applications from an earlier version to V6.2.3 or V6.3.0, the following entries are created in the Add or Remove Programs window: <ul style="list-style-type: none"> • IBM Tivoli Monitoring for Microsoft Applications • IBM Tivoli Monitoring 	Ignore IBM Tivoli Monitoring for Microsoft Applications in the Add or Remove Programs window. If you want to modify or remove the installed agents in IBM Tivoli Monitoring for Microsoft Applications, use IBM Tivoli Monitoring from the Add or Remove Programs window.

Table 26. Problems and solutions for installation and configuration of the Microsoft SQL Server agent

Problem	Solution
The procedure for launching the Microsoft SQL Server agent Trace Parameters window in “RAS trace parameters” on page 275 fails because no trace settings are available in the Enter RAS1 Filters field.	This problem happens when the KOQENV file is missing. If the KOQENV file is missing from the <i>install_dir\tmaitm6</i> directory, the Microsoft SQL Server agent must be reinstalled.
Collector trace logs are not receiving error information.	Review the information in “Trace logging” on page 270 to ensure that you are consulting the correct log file. The file is located in the following path: <i>install_dir\tmaitm6</i> , where <i>install_dir</i> is the location of IBM Tivoli Monitoring.
After running the agent successfully, you reinstall the agent software, and collection of monitoring data stops.	Consult the list of supported versions in the configuration chapter of the agent user's guide for SQL Server. Confirm that you are running a valid version of the target application. If you are monitoring a supported version of the database application, gather log files and other information and contact IBM Software Support, as described in “Gathering product information for IBM Software Support” on page 269.
Alert summary report of Tivoli Enterprise Portal shows no information.	Ensure that the configuration setting of agent's error log file defines a valid path. Confirm that the error log file is correct. If the error log file is correct, the problem can occur because the logon account that you are using does not have sufficient permissions in the IBM Tivoli Monitoring environment. For details see Chapter 2, “Agent installation and configuration,” on page 15.
Error counts are displayed in the Alert summary report in the Tivoli Enterprise Portal, however, error messages are not displayed in the Alert detail report.	Check the time stamp for the reports. If you have set up historical data collection for Alert summary report, it is recommended that you set up historical data collection for the Alert detail report, too.

Table 26. Problems and solutions for installation and configuration of the Microsoft SQL Server agent (continued)

Problem	Solution
<p>When installing the SQL Server agent support file for the Tivoli Enterprise Monitoring Server, you see the following error message: RC=350</p> <p>This problem might occur if the Tivoli Enterprise Portal client fails to display the Enterprise Database Mirroring workspace.</p>	<p>You can continue using the SQL Server agent despite this error. However, the Enterprise Database Mirroring data collection is affected.</p>
<p>During the silent uninstallation of the SQL Server agent, if the Manage Tivoli Enterprise Monitoring Services window is open, the agent fails to uninstall.</p>	<p>Close the Manage Tivoli Enterprise Monitoring Services window before uninstalling the agent using the silent uninstallation method.</p>
<p>Historical Data Collection takes some seconds more than the configured data collection interval time to finish data collection. This delay gives an impression that some data rows are missing when the collection interval is 1 minute.</p>	<p>Keep the collection interval to a minimum of 5 minutes and use a smaller interval for an attribute group if it is critical.</p>
<p>When you upgrade the SQL Server from V2005 to V2008 R2 or V2012 and configure the SQL Server agent locally, the database version in the Database Server Properties window does not display the upgraded version of the SQL Server.</p>	<p>Remove the earlier instance of the SQL Server agent and create another instance.</p>

Table 26. Problems and solutions for installation and configuration of the Microsoft SQL Server agent (continued)

Problem	Solution
<p>Unique names for monitoring components: ORIGINNODE is truncated and doesn't show the product code.</p>	<p>IBM Tivoli Monitoring might not be able to generate a unique name for monitoring components due to the truncation of names that the product automatically generates. IBM Tivoli Monitoring automatically creates a name for each monitoring component by concatenating the subsystem name, host name, and product code separated by colons (<i>subsystem_name:hostname:MSS</i>).</p> <p>Note: When you monitor a multinode systems, like databases, IBM Tivoli Monitoring uses a database instance name as the subsystem name.</p> <p>The length of the name that IBM Tivoli Monitoring generates is limited to 32 characters. Truncation can result in multiple components having the same 32-character name. If this problem happens, shorten the <i>subsystem_name</i> portion of the name as described in the steps as below:</p> <ul style="list-style-type: none"> • On Windows: <p>Note: This procedure involves editing the Windows Registry. An error in editing the Registry may cause other system errors. It is best practice to make a backup copy of the Registry prior to modifying the Registry. If you do not feel comfortable editing the Registry, contact IBM Software Support.</p> <ol style="list-style-type: none"> 1. Run the Registry editor, regedit. Select Start > Run. Type regedit in the field and click OK. 2. Locate the following key: <code>HKEY_LOCAL_MACHINE\SOFTWARE\Candle\K0Q\610\instance_name</code> <p>where <i>instance_name</i> is the Microsoft SQL Server instance name that is associated with the target agent.</p> 3. Open the Environment key. 4. Select the CTIRA_SUBSYSTEM_ID string attribute. If you do not find CTIRA_SUBSYSTEM_ID, add it. step 6. Type a new name for an instance name (subsystem name). 5. Right-click and select Modify. 6. Type a new instance name (subsystem name). Keep in mind that the final concatenated name, including the subsystem name, host name, and OQ, cannot be longer than 32 characters. Note: You must ensure that the resulting name is unique with respect to any existing monitoring component that was previously registered with the Tivoli Enterprise Monitoring Server. 7. Exit from the Registry editor. 8. Restart the agent. 9. A second instance in the Tivoli Enterprise Portal Navigation tree might be automatically added. If this happens, clear the old, offline agent instance navigation item using the Managed System Status workspace of the Enterprise Navigation tree item.

Table 26. Problems and solutions for installation and configuration of the Microsoft SQL Server agent (continued)

Problem	Solution
<p>In Windows 2003 system, a non-administrator user cannot configure the SQL Server agent services through the Change Startup option in the Manage Tivoli Enterprise Monitoring Services window.</p>	<p>To configure the SQL Server agent services on a Windows 2003 system, perform the following steps:</p> <ol style="list-style-type: none"> 1. On the Windows Start menu, click Run. 2. Type <code>services.msc</code>, and click OK. 3. In the Services window, perform these steps for the agent service and the agent collector service: <ol style="list-style-type: none"> a. Right-click the service, and click Properties. b. Click the Log On tab. c. Click This account. d. Type the user name. e. In the Password and Confirm Password field, enter password, and click OK. 4. In the Manage Tivoli Enterprise Monitoring Services window, right-click the SQL Server agent service, and click Recycle.
<p>If you have installed a 64-bit agent for ITCAM for Microsoft Applications, information about the agent processes is not displayed in the Kinviewer window.</p>	<p>There is no solution to this problem.</p>

Table 26. Problems and solutions for installation and configuration of the Microsoft SQL Server agent (continued)

Problem	Solution
<p>When you install ITCAM for Microsoft Applications from the product DVD, the kininfo command does not display the installation date when you install 32-bit or 64-bit agents of ITCAM for Microsoft Applications. For the 64-bit agent, the kininfo command displays two entries.</p>	<p>Complete the following steps to resolve this problem for a 32-bit agent:</p> <ol style="list-style-type: none"> 1. Edit the KXXWICMA.ver file in the <i>Candle_home</i>\Install\ITM\Ver directory, and then add the following properties and values at the end of the file: GSKit=LOCAL[INSTALL INFO] Install Date=<i>Date_Time</i>Install Dir=<i>Candle_home</i>Install GUID={23B43DE3-AE00-E230-ED11-BA05A422701D} 2. Replace <i>Candle_home</i> with the folder path where you have installed the agent, for example, C:\IBM\ITM. 3. Replace <i>Date_Time</i> with the name of the latest folder that is created in the <i>Candle_home</i>\Install\ITM\VerPrevInstall directory, for example, Install Date=20100721 0700. <p>Complete the following steps to resolve this problem for a 64-bit agent:</p> <ol style="list-style-type: none"> 1. Edit the KXX64CMA.ver file in the <i>Candle_home</i>\Install\ITM\Ver directory, and then add the following properties and values at the end of the file: GSKit=LOCAL[INSTALL INFO] Install Date=<i>Date_Time</i>Install Dir=<i>Candle_home</i> Install GUID={A75F7F91-1FF2-49C6-ABFF-1FDAF54E7E67} 2. Replace <i>Candle_home</i> with the folder path where you have installed the agent, for example, C:\IBM\ITM. 3. Replace <i>Date_Time</i> with the name of the latest folder that is created at <i>Candle_home</i>\Install\ITM\VerPrevInstall, for example, Install Date=20100721 0700. 4. Edit the KXXWICMA.ver file in the <i>Candle_home</i>\Install\ITM\Ver directory, and then add the following properties and values at the end of the file: GSKit=LOCAL Dummy32BitVerfile=YES[INSTALL INFO] Install Date=<i>Date_Time</i>Install Dir=<i>Candle_home</i> Install GUID={A75F7F91-1FF2-49C6-ABFF-1FDAF54E7E67} 5. Replace <i>Candle_home</i> and <i>Date_Time</i> values with the values that are used in steps 2 and 3. 6. Edit the KXXWICMA.ver file in the <i>Candle_home</i>\Install\ITM\Ver directory, and then modify the value of the PlatformArch variable from WIX64 to WINNT. For example, if the KXXWICMA.ver file contains PlatformArch=WIX64, change it to PlatformArch=WINNT.

Table 27. General problems and solutions for uninstallation

Problem	Solution
<p>On Windows, uninstallation of IBM Tivoli Monitoring fails to uninstall the entire environment.</p>	<p>Confirm that the following problems do not exist:</p> <ul style="list-style-type: none"> • Ensure that you are the only user who is logging into the computer where you are performing an uninstallation operation. If another user is performing operations during an uninstall process, the uninstall process fails. • Be sure that you follow the general uninstallation process described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>: <ol style="list-style-type: none"> 1. Uninstall monitoring agents first, as in the following examples: <ul style="list-style-type: none"> – Uninstall a single monitoring agent for a specific database. —OR— – Uninstall all instances of a monitoring product, such as IBM Tivoli Monitoring for Microsoft Applications. 2. Uninstall IBM Tivoli Monitoring.
<p>The way to remove inactive managed systems (systems whose status is OFFLINE) from the Navigator tree in the portal is not obvious.</p>	<p>Use the following steps to remove, but not uninstall, an offline managed system from the Navigator tree:</p> <ol style="list-style-type: none"> 1. Click the Enterprise icon in the Navigator tree. 2. Right-click, then click Workspace > Managed System Status. 3. Right-click the offline managed system, and select Clear offline entry. <p>If you also want to uninstall the monitoring agent, use the procedure described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.</p>
<p>Attempting a remote uninstallation of the Microsoft SQL Server agent through the Tivoli Enterprise Portal client triggers a timeout error that does not correctly identify the managed system where the timeout is said to have occurred. The error text does not display another managed system; instead it displays incorrect text in place of the managed system name. However, the agent does appear to be uninstalled. The error message returned is KFWITM290E.</p>	<p>The agent is uninstalled and no other known errors are associated with this problem at this time.</p>
<p>You have installed Tivoli Enterprise Monitoring Agent Framework V6.2.2, Fix Pack 2 on a managed system. If you uninstall the SQL Server agent by using the Modify option in the Add or Remove Programs window, the entry for the SQL Server agent is not removed from the Manage Tivoli Enterprise Monitoring Service window.</p>	<p>Remove the SQL Server agent instances from the Manage Tivoli Enterprise Monitoring Service window manually.</p>

Table 27. General problems and solutions for uninstallation (continued)

Problem	Solution
You have installed an ITCAM for Microsoft Applications agent, V6.2.3 on a computer. Later, you installed earlier version (V6.2.2) of other ITCAM for Microsoft Applications agents on the same computer. The entry for IBM Tivoli Monitoring is removed from the Add or Remove Programs window, and you cannot uninstall any of the ITCAM for Microsoft Applications agents from the computer.	To restore the entry for IBM Tivoli Monitoring in the Add or Remove Programs window, run the setup.exe file of ITCAM Microsoft Applications agent, V6.2.3.
You have installed a 64-bit agent of ITCAM for Microsoft Applications V6.2.3 on a computer that has preinstalled ITCAM for Microsoft Applications agents on the same computer. When you try to uninstall and reinstall the 64-bit agent of ITCAM for Microsoft Applications V6.2.3 on the computer, the 64-bit agent does not start automatically.	Perform the following steps to start the 64-bit agent: <ol style="list-style-type: none"> 1. Unconfigure the agent. 2. Reconfigure and start the agent.
You have the Windows OS agent and one or more agents of ITCAM for Microsoft Applications preinstalled on the same computer. When you run the ITCAM for Microsoft Applications installation wizard to upgrade the agents from V6.2.2 to a later version, the installation wizard does not display the version of the preinstalled ITCAM for Microsoft Applications agents.	There is no solution to this problem.
When you use the Installation Launch Pad to install a 32-bit agent on a 32-bit computer, the following message is displayed: The following components cannot be installed because the installation action failed. IBM Tivoli Composite Application Manager for Microsoft Applications V6.2.3 for Windows on 64-bit AMD and Intel systems (x64)	No action is required. You can ignore this message and continue with the installation.

Troubleshooting for SQL Server

Table 28 lists problems that might occur on the system or application that you are monitoring. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 28. SQL Server problems and solutions

Problem	Solution
The following Microsoft SQL Server metrics might exceed 100 percent values: <ul style="list-style-type: none"> • Procedure Buffers Percent Active • Procedure Cache Percent Active 	Be aware of this problem as you monitor Microsoft SQL Server. There is no known solution. Monitor the product support website for Microsoft SQL Server for updated information on this problem.

Agent troubleshooting

This section lists problems that might occur with agents.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 29. General agent problems and solutions

Problem	Solution
The system experiences high CPU usage after you install or configure Microsoft SQL Server agent.	View the memory usage of the KOQCMA process. If CPU usage seems to be excessive, recycle the monitoring agent.
A monitoring agent seems to be offline.	<ol style="list-style-type: none"> 1. Allow several minutes to pass whenever you restart the Tivoli Enterprise Portal Server. A monitoring agent will seem to be down while the portal server is restarting. 2. Select Physical View and highlight the Enterprise Level of the navigator tree. 3. Select View > Workspace > Managed System Status to see a list of managed systems and their status. 4. If a system is offline, check network connectivity and status of the specific system or application.
Some TEC events from the Microsoft SQL Server agent have PARSING_FAILED status in wtdump output of TEC server.	<p>Solution This problem can be caused by a wrong data type, as shown in the following wtdump output:</p> <pre>1~14560~1~1158832838(??? 21 14:00:38 2006) ### EVENT ### ITM_MS_SQL_Process_Summary;source='ITM';sub_source='TW06210:TW06210:MSS';cms_hostname='tps06072.tw.twnet';cms_port='3661';integration_type='U';master_reset_flag='';appl_label='';situation_name='MS_SQL_Process0thSleep_Warning';situation_origin='TW06210:TW06210:MSS';situation_time='09/21/2006 \ 14:00:37.000';situation_status='Y';hostname='TW06210';origin='192.168.192.120';adapter_host='HUB_TPS06072';severity='WARNING';date='09/21/2006';msg='MS_SQL_Process0thSleep_Warning[(Total_Other_Sleep>=2 AND Percent_Processes_Othersleep>=50.00) ON TW06210:TW06210:MSS (Total_Other_Sleep=46 Percent_Processes_Othersleep=74.19)]';situation_displayitem='';current_interval='193';host_name='TW06210';hub_timestamp='1060921135139000';originnode='TW06210:TW06210:MSS';percent_processes_bad='0.00';percent_processes_blocked='0.00';percent_processes_infected='0.00';percent_processes_locksleep='0.00';percent_processes_othersleep='74.19';percent_processes_sleeping='74.19';percent_processes_stopped='0.00';sample_timestamp='1060921134858000';server='TW06210';server_cpu_percent_application='0.00';server_cpu_percent_system='0.00';total_log_suspend='0';total_other_sleep='46';total_processes='62';total_processes_bad='0';total_processes_blocked='0';total_processes_infected='0';total_processes_locksleep='0';total_processes_stopped='0';situation_eventdata='~';END ### END EVENT ### PARSING_FAILED~'Line 1: Value does not match type REAL!'</pre> <p>Use the odadmin environ command as follows to set en_US in the oserv environment:</p> <ol style="list-style-type: none"> 1. odadmin environ get > /tmp/env.out 2. Edit /tmp/env.out and set LANG=en_US 3. odadmin reexec all

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>No performance data is displayed in workspace views, no data is available for situations, and no data is available for historical logging.</p>	<p>When the Windows operating system detects a problem in one of its extensible performance monitoring DLL files, it marks the DLL as disabled. Any DLL that is disabled cannot provide performance data through the Windows Performance Monitor interfaces (Perfmon or Performance Monitor APIs). This prevents IBM Tivoli Monitoring agents from gathering data supplied by the disabled DLL. For more information, see Microsoft Support Knowledge Base article 248993 at the following Web address: http://support.microsoft.com/default.aspx?scid=kb;EN-US;248993</p> <p>Follow the Resolution instructions provided in this article (248993) to re-enable any performance monitoring extension DLL files disabled by Windows. Then, restart the monitoring agent.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>A configured and running instance of the monitoring agent is not displayed in the Tivoli Enterprise Portal, but other instances of the monitoring agent on the same system do appear in the portal.</p>	<p>Tivoli Monitoring products use Remote Procedure Call (RPC) to define and control product behavior. RPC is the mechanism that allows a client process to make a subroutine call (such as GetTimeOfDay or ShutdownServer) to a server process somewhere in the network. Tivoli processes can be configured to use TCP/UDP, TCP/IP, SNA, and SSL as the desired protocol (or delivery mechanism) for RPCs.</p> <p>"IP.PIPE" is the name given to Tivoli TCP/IP protocol for RPCs. The RPCs are socket-based operations that use TCP/IP ports to form socket addresses. IP.PIPE implements virtual sockets and multiplexes all virtual socket traffic across a single physical TCP/IP port (visible from the netstat command).</p> <p>A Tivoli process derives the physical port for IP.PIPE communications based on the configured, well-known port for the HUB Tivoli Enterprise Monitoring Server. (This well-known port or BASE_PORT is configured using the PORT: keyword on the KDC_FAMILIES / KDE_TRANSPORT environment variable and defaults to 1918.)</p> <p>The physical port allocation method is defined as (BASE_PORT + 4096*N) where N=0 for a Tivoli Enterprise Monitoring Server process and N={1, 2, ..., 15} for a non-Tivoli Enterprise Monitoring Server. Two architectural limits result as a consequence of the physical port allocation method:</p> <ul style="list-style-type: none"> • No more than one Tivoli Enterprise Monitoring Server reporting to a specific Tivoli Enterprise Monitoring Server HUB can be active on a system image. • No more than 15 IP.PIPE processes can be active on a single system image. <p>A single system image can support any number of Tivoli Enterprise Monitoring Server processes (address spaces) provided that each Tivoli Enterprise Monitoring Server on that image reports to a different HUB. By definition, there is one Tivoli Enterprise Monitoring Server HUB per monitoring Enterprise, so this architecture limit has been simplified to one Tivoli Enterprise Monitoring Server per system image.</p> <p>No more than 15 IP.PIPE processes or address spaces can be active on a single system image. With the first limit expressed above, this second limitation refers specifically to Tivoli Enterprise Monitoring Agent processes: no more than 15 agents per system image.</p> <p>This limitation can be circumvented (at current maintenance levels, IBM Tivoli Monitoring V6.1 Fix Pack 4 and later) if the Tivoli Enterprise Monitoring Agent process is configured to use EPHEMERAL IP.PIPE. (This is IP.PIPE configured with the EPHEMERAL:Y keyword in the KDC_FAMILIES / KDE_TRANSPORT environment variable). There is no limitation to the number of ephemeral IP.PIPE connections per system image. If ephemeral endpoints are used, the Warehouse Proxy Agent is accessible from the Tivoli Enterprise Monitoring Server associated with the agents using ephemeral connections either by running the Warehouse Proxy Agent on the same computer or by using the Firewall Gateway feature. (The Firewall Gateway feature relays the Warehouse Proxy Agent connection from the Tivoli Enterprise Monitoring Server computer to the Warehouse Proxy Agent computer if the Warehouse Proxy Agent cannot coexist on the same computer.)</p>
<p>Attributes do not allow non-ASCII input in the Situation Editor.</p>	<p>None. Any attribute that does not include (Unicode) might support only ASCII characters. For example Attribute (Unicode) will support Unicode but Attribute without (Unicode) might only support ASCII characters.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>When remotely stopping an instance of the Microsoft SQL Server agent shortly after the instance has started, a KDY1024E error message is returned and the monitoring agent instance is stopped.</p>	<p>After starting the instance, wait 10 minutes for the instance to complete initialization before stopping the agent.</p>
<p>Negative value is returned for Transactions per Second in the Database Detail attribute group when the user database has a similar name to that of a system database.</p>	<p>No known workaround at this time.</p>
<p>The Total OS CPU percent attribute of the Tivoli Enterprise Portal Server shows erroneous value in the Tivoli Enterprise Portal.</p> <p>ITM 6.X database agents do not monitor attributes that are related to operating systems. The database agent might not reflect dynamic values collected at every instance because the database agent collects the value at a specific interval of time.</p>	<p>Use the Operating System agents to monitor the Total OS CPU percent attribute.</p>
<p>Queries results that return No_Data</p> <p>If no data is available for an attribute group, the Tivoli Enterprise Portal shows a row with zero values. This triggers false situations. Also, if a row is returned and selected attributes have no data, these attributes are returned as a zero value.</p>	<p>These changes are implemented in the Microsoft SQL Server agent V6.2.1:</p> <ul style="list-style-type: none"> • If no data is available for an attribute group, then no rows are returned to the Tivoli Enterprise Monitoring Server or the Tivoli Enterprise Portal Server. • If a row is returned and the selected attributes have no data, then the selected attributes are set to a value of -1 for integer type attributes and No_Data for string type attributes. -1 values are translated to Not_Collected and displayed in the Tivoli Enterprise Portal Server.
<p>The Server Status attribute of the Server Summary or Server Detail attribute groups returns a value of Inactive or Unknown.</p> <p>A value of Inactive is returned to the Tivoli Enterprise Portal Server when the agent collector service fails to communicate with the SQL Server. A value of Unknown is returned to the Tivoli Enterprise Portal Server when the agent collector service is stopped, being restarted, or not functioning properly.</p>	<p>Ensure that the SQL Server is running. It might take a few minutes for the agent to successfully connect to the SQL Server and return a Server Status value of Active. If the SQL Server is running and the Collection Status value is still Inactive, then restart the SQL Server agent.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>The Collection Status attribute of the Server Summary attribute group returns a value of Inactive.</p> <p>A value of Inactive indicates that the SQL Server agent data collector service has stopped. The collector service might enter this state if:</p> <ul style="list-style-type: none"> • The collector service is manually stopped. • An internal error has occurred. • The SQL Server is stopped. 	<p>Ensure that the SQL Server is running. If the SQL Server is running and the Collection Status value is still Inactive, then restart the SQL Server agent.</p>
<p>If you rename a database in the SQL management studio, the SQL Server agent does not collect data for the following attributes in the MS SQL Database Detail attribute group:</p> <ul style="list-style-type: none"> • Database Growth Percent • Data Size • Data Freespace Percent • Log Freespace • Log Size • Last Database Size 	<p>Restart the SQL Server database service.</p>
<p>When you create two SQL Server databases with the same name, but precede one database name with a space and the other database name without a space, the portal does not display data properly.</p>	<p>The space at the beginning of the database name is truncated by the portal, making the two databases look identical.</p> <p>Avoid spaces at the beginning of a database name.</p>
<p>If the table level locking hint (Holdlock, Tablock, and so on) is specified with the SELECT, INSERT, UPDATE and DELETE statements on multiple tables in an SQL Server database, the portal does not display data till the lock is released.</p>	<p>Release the locks.</p>
<p>The memory usage of data collector service is high (approximately 150 megabytes) in the case of the SQL Server agent that runs on the Windows Server 2008 operating system, whereas the memory usage is about 50 megabytes in a Windows 2003 environment.</p>	<p>This is a limitation with the Windows 2008 environment.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>The remote agent functionality (for example, Start SQL Server, Stop SQL Server actions) and some other functionality does not work in some Windows locales, such as English (India), English (Malaysia), English (Singapore), and so on. Also, a few attributes that are of the String data type show garbage strings instead of displaying correct values.</p>	<p>Configure the Windows server to run on other Windows locale that is supported by the agent. Note that the change of locale affects the functioning of other applications on the Windows server.</p>
<p>If there is a shared or granted locks on the tempdb database, then the Lock Detail workspace displays lock information for the tempdb database on the portal. However, these locks might not be counted under the Shared Locks attribute of the Lock Summary attribute group.</p>	<p>None.</p>
<p>If the SQL Server agent fails to collect data, the Server Version attribute in the Server Summary view of the Enterprise Server Summary workspace displays empty value.</p>	<p>None.</p>
<p>Multiple instances of some agent processes, such as koqllog.exe, koqos.exe, and koqserv.exe are displayed in the Task Manager of the workstation where the SQL Server agent is installed. The memory usage for these processes is constant and the CPU usage is zero.</p>	<p>Generally every instance of the agent creates an entry in the Task Manager. If there are more entries than the total number of configured instances, which do not show any CPU usage for long hours but continuously persist in the Task Manager, end these processes manually.</p>
<p>If you restart the workstation where the SQL Server agent is installed, the agent service koqcoll.exe stops and the event viewer displays the following error message: Faulting application koqcoll.exe, version 0.0.0.0, faulting module ntdll.dll, version 5.2.3790.4455, fault address 0x00000000.</p>	<p>This problem occurs when the SQL Server service is not started. If the agent is in the stopped state even though the agent workstation is restarted, recycle the agent from the Manage Tivoli Enterprise Monitoring Services window.</p>
<p>If a linked server is configured on the SQL Server, the portal displays incorrect data for the Network Name attribute of the Remote Server attribute group.</p>	<p>None.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>The value of the Collection Level attribute in the Remote Servers attribute group is always displayed as 1 on the Tivoli Enterprise Portal. However, the possible values of the Collection level attribute are 1 and 0.</p>	<p>There is no solution to this problem.</p>
<p>Although the incorrect version of the SQL Server is configured with the SQL Server agent, the agent data is displayed on the Tivoli Enterprise Portal.</p>	<p>There is no solution to this problem.</p>
<p>The SQL Server agent takes time to display data when a large number of databases (typically, more than 6000) are created on the SQL Server. The agent data collector service stops running if the memory used by the data collector exceeds the threshold value because of a large number of databases on the SQL Server.</p>	<p>There is no solution to this problem.</p>
<p>You have installed and configured a 64-bit SQL Server agent locally on a computer. When you start the SQL Server agent, the collector service (koqcoll.exe) stops abruptly.</p>	<p>Install the SQL Server agent remotely.</p>
<p>In the "xxxagent.exe Properties" window, the following agent information is not updated:</p> <ul style="list-style-type: none"> • Copyright • File Version • Product Version <p>This problem occurs on both Windows 32-bit and 64-bit operating systems. Note: On Windows 2003 systems, the agent information is shown in the Version tab; however, on Windows 2008 systems, the agent information is shown in the Details tab.</p>	<p>There is no solution to this problem.</p>
<p>On the z/OS® operating system, while configuring all attributes in the History Collection Configuration window, if you select TEMS in the Collection Location field, the attribute group tables are not created in the Warehouse database.</p>	<p>In the History Collection Configuration window, select TEMA in the Collection Location field.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>When you install the 32-bit SQL Server agent on a 64-bit operating system and a 64-bit MS SQL Server, you cannot view data for the following workspaces:</p> <ul style="list-style-type: none"> • Analysis services Datamining Statistics • Analysis Services Memory Statistics • Analysis services Storage Engine Query processing • Analysis Services Row Processing 	<p>Install the 64-bit SQL Server agent.</p>
<p>The 64-bit SQL Server agent displays a high use of memory when the agent is installed and configured on the 64-bit Windows 2008 R2 operating system.</p>	<p>Install the service pack 1 of the 64-bit Windows 2008 R2 operating system.</p>
<p>On the Tivoli Enterprise Portal, no value is displayed for the primary replica in the Availability Group Details attribute group. This problem occurs when you pause or stop the cluster service on the node that hosts the SQL Server agent.</p>	<p>Complete the following steps to resolve the problem:</p> <ol style="list-style-type: none"> 1. In the Failover Cluster Manager, right-click the cluster node that hosts the SQL Server agent. 2. If you have paused the cluster service on the node, click Resume. If you have stopped the cluster service on the node, click Start.
<p>On the Tivoli Enterprise Portal, incorrect values are displayed when a custom query is applied for the following attribute groups:</p> <ul style="list-style-type: none"> • MS SQL Availability Groups Summary • MS SQL Availability Database Summary • MS SQL Availability Database Details • MS SQL Availability Replicas Status Summary • MS SQL Availability Replicas Details In Cluster • MS SQL Service Broker Activation 	<p>Refresh the Tivoli Enterprise Portal after applying the custom query for the following attribute groups:</p> <ul style="list-style-type: none"> • MS SQL Availability Groups Summary • MS SQL Availability Database Summary • MS SQL Availability Database Details • MS SQL Availability Replicas Status Summary • MS SQL Availability Replicas Details In Cluster • MS SQL Service Broker Activation
<p>On the Tivoli Enterprise Portal, the Stolen Pages Growth attribute displays negative values. The value for the Stolen Pages Growth attribute is calculated on the basis of the previous and the current sample data.</p>	<p>There is no solution to this problem.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
<p>You delete the Microsoft SQL Server agent by completing the following steps:</p> <ul style="list-style-type: none"> • Click Start > Programs > IBM Tivoli Monitoring > Manage Tivoli Enterprise Monitoring Services. • In the Manage Tivoli Enterprise Monitoring Services window, right-click Microsoft SQL Server agent, and click Advanced > Remove Instance. <p>After the Microsoft SQL Server agent is deleted, the status of the agent and collector services changes to Disabled. If you configure the SQL Server agent again by using the agent configuration panel without closing the Manage Tivoli Enterprise Monitoring Services window, the agent fails to start.</p>	<p>After deleting the Microsoft SQL Server agent, close the Manage Tivoli Enterprise Monitoring Services window and reopen it before you configure the monitoring agent again.</p>
<p>The Not collected enumeration value is incorrectly displayed as -0.01 in the warehouse database. Similarly, the Not applicable enumeration value is incorrectly displayed as -0.02 in the warehouse database. While designing Tivoli Common Reporting reports, the value -0.01 must be displayed as Not collected, and the value -0.02 must be displayed as Not applicable in the reports. The correct enumeration value for Not collected is -1 and the correct enumeration value for Not applicable attribute is -2.</p>	<p>There is no solution to this problem.</p>
<p>While upgrading the SQL Server agent to version 6.3, the following message appears:</p> <pre>Creating instance at Path c:\IBM\ITM\TMAITM6\koqma.ini Product Code KOQ Instance <i>instance name</i> failed</pre> <p>However, the SQL Server agent is successfully upgraded to version 6.3.</p>	<p>There is no solution to this problem.</p>

Table 29. General agent problems and solutions (continued)

Problem	Solution
The configuration file of the MS SQL Problem Detail attribute group cannot handle non-English ERRORLOG files by default.	<p>Update the koqErrConfig.ini file to resolve this issue. For example, for the French language, update the koqErrConfig.ini file to reflect the following format:</p> <pre>[French_CI_AS] Error=Erreur : Severity=Gravité : State=État:</pre> <p>For the English language, update koqErrConfig.ini the file to reflect the following format:</p> <pre>[SQL_Latin1_General_CP1_CI_AS] Error=Error: Severity=Severity: State=State:</pre> <p>Note: If you are using the 32-bit Microsoft SQL Server agent, the koqErrConfig.ini file is saved in the candle_home\TMAITM6 directory. If you are using the 64-bit SQL Server agent, the koqErrConfig.ini file is saved in the candle_home\TMAITM6_x64 directory.</p>
If the device type is FILESTREAM for any database, the attributes in the MS SQL Device Detail attribute group display negative or incorrect values for the database.	There is no solution to this problem.

Troubleshooting for remote deployment

Table 30 lists problems that might occur with remote deployment. This section provides information about troubleshooting remote deployment of the monitoring agent. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

This section describes problems and solutions for remote deployment and removal of agent software Agent Remote Deploy:

Table 30. Remote deployment problems and solutions

Problem	Solution
Remote deployment fails because the process exceeds the communication timeout.	Increase the amount of time allowed for remote deployment, as described in the "Troubleshooting Tivoli Enterprise Monitoring Agent problems" chapter of the <i>IBM Tivoli Monitoring Troubleshooting Guide</i> .
<p>You have one of following problems after you use the tacmd addSystem command to perform a remote deployment:</p> <ul style="list-style-type: none"> The Tivoli Enterprise Portal fails to display workspaces for the monitoring agent. —OR— You see the KDY1008E error message and a return code of 4. 	<p>Ensure that you have typed the correct values for the parameters in the tacmd addSystem command. For example, to run the command successfully, you must provide correct values for the following parameters:</p> <ul style="list-style-type: none"> the home directory of the database server the name of the Microsoft SQL Server that you want to monitor the name of the Microsoft SQL Server instance that you want to monitor <p>See Chapter 2, "Agent installation and configuration," on page 15 for more information on these parameters.</p>

Table 30. Remote deployment problems and solutions (continued)

Problem	Solution
While you are using the remote deployment feature to install Microsoft SQL Server agent, an empty command window is displayed on the target computer. This problem occurs when the target of remote deployment is a Windows computer. (See the <i>IBM Tivoli Monitoring Installation and Setup Guide</i> for more information on the remote deployment feature.)	Do not close or modify this window. It is part of the installation process and will be dismissed automatically.
The removal of a monitoring agent fails when you use the remote removal process in the Tivoli Enterprise Portal desktop or browser.	This problem might happen when you attempt the remote removal process immediately after you have restarted the Tivoli Enterprise Monitoring Server. You must wait at least 10 minutes for the monitoring agent to refresh its connection with the Tivoli Enterprise Monitoring Server before you begin the remote removal process.
When you run the tacmd command with the DBSETTINGS.db_ver property, the following error message is displayed: The property DBSETTINGS.db_ver is invalid for version 062100000 of the OQ agent.	The property DBSETTINGS.db_ver is not valid for V6.2.1 of the SQL Server agent. Run the tacmd describeSystemType command to view the properties that are permitted for the agent type. If you run the tacmd command by using a script, you must modify the script to remove the DBSETTINGS.db_ver option.
If you have not specified the parameter values in the Table Detail Collection Settings pane in the Database Server Properties window and the size of the tables in the databases is large, agent data is not displayed on the Tivoli Enterprise Portal.	Enter the parameter values in the Table Detail Collection Settings pane in the Database Server Properties window for data collection of Table Detail attribute group.
You have installed an agent of ITCAM for Microsoft Applications remotely from the Tivoli Enterprise Portal Client. On the Agent tab of the Managed System Configuration window, you selected the Use this account option, and typed the account information in the "user@domain.com"; format (for example, administrator@itmagents.com). The Deployment Status Detail view shows the agent deployment status as failed; however, the agent is configured and installed with the LocalSystem account instead of the specified user account.	Specify the account information in the domain\user format.

Workspace troubleshooting

Table 31 shows problems that might occur with workspaces. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 31. Workspace problems and solutions

Problem	Solution
No data is displayed for the monitoring agent in the workspaces of the Tivoli Enterprise Portal.	<ol style="list-style-type: none"> 1. Confirm that the computer that hosts the monitoring agent is communicating with the Tivoli Enterprise Monitoring Server. 2. Verify the Microsoft SQL Server configuration credentials that you used during configuration of the monitoring agent. The credentials must match the options listed in Chapter 2, "Agent installation and configuration," on page 15.

Table 31. Workspace problems and solutions (continued)

Problem	Solution
<p>You start collection of historical data but the data cannot be seen.</p>	<p>Managing options for historical data collection:</p> <ul style="list-style-type: none"> • Basic historical data collection populates the Warehouse with raw data. This type of data collection is turned off by default. See Chapter 2, “Agent installation and configuration,” on page 15 for information on managing this feature including how to set the interval at which data is collected. By setting a more frequent interval for data collection you reduce the load on the system incurred every time data is uploaded. • You use the Summarization and Pruning monitoring agent to collect specific amounts and types of historical data. Be aware that historical data is not displayed until the Summarization and Pruning monitoring agent begins collecting the data. By default, this agent begins collection at 2 AM daily. At that point, data is visible in the workspace view. See the IBM Tivoli Monitoring Administrator's Guide to learn how to modify the default collection settings.
<p>The following workspace views display column headers, but no data: Process Holding Lock SQL Text, Blocking Process SQL Text, Waiting Process SQL Text, and SQL Text</p>	<p>The workspaces containing these views only display valid data when the workspace is accessed from a workspace link. Workspace Administrators may access these workspaces directly (without following the correct workspace link) and this will result in no data being displayed for these views. When you want to view monitored data in these views, use the workspace links. See the workspaces chapter of this user's guide for additional information on these workspace views.</p> <p>The SQL Server ID configured for use by the MS SQL agent requires System Administrators SQL Server Role authorization to access the SQL text data. If you want to display SQL test information, give System Administrators authorization to the MS SQL agent's SQL Server user ID. Refer to the Chapter 2, “Agent installation and configuration,” on page 15 section in this book.</p>
<p>The Microsoft SQL Server agent workspace views have empty data, and there is no <i>hostname_00_sql server instance name__regtitls.txt</i> file in the TMAITM6\logs subdirectory.</p>	<p>The absence of the target MS SQL Server performance counters causes this problem. If these prerequisite counters are not installed and enabled, the monitoring agent is not able to properly collect monitoring data for the target MS SQL Server. When this happens, the RAS1 trace log contains the following error trace message: SQL Server Instance performance data is not enabled!</p> <p>This might be caused by installing the MS SQL Server instance using a Minimum installation setup type, or by not selecting Performance Counters from a Custom installation setup type. Install and configure the performance counters for the target MS SQL Server instance. This might require reinstallation of the MS SQL Server instance. Use either the Typical or Custom installation setup type.</p>

Table 31. Workspace problems and solutions (continued)

Problem	Solution
<p>When configuring the time span for the historical data in the Errorlog Alerts workspace, the time span icon will gray out (become inactive) and no historical data will be displayed.</p>	<p>When configuring the time span, click the Apply to all views associated with this view's query checkbox in the Select the Time Span property box and click OK.</p> <p>–OR–</p> <p>Regenerate the query for each view by:</p> <ol style="list-style-type: none"> 1. Right-click the view with the inactive time span icon. 2. Select Properties. 3. Click Click here to assign a query. 4. Reselect the query assigned for that view. 5. Click Apply on the Query Editor page.
<p>When creating a job in the SQL Server, if you specify the job name in double quotation marks, then the Job Detail and Job Status views of the Job Detail workspace do not display data rows on the portal.</p>	<p>Do not specify the job name in double quotation marks.</p>
<p>When creating a database in the SQL Server, if you specify the database name in single quotation marks, then the last character of the database name is truncated on the portal. If you specify a database name in double quotation marks, then the value of the Database Name attribute is displayed as blank on the portal.</p>	<p>Do not specify the database name in single or double quotation marks.</p>
<p>In the Table Detail workspace, the agent does not show table details for empty tables (tables that do not contain rows) in the SQL Server database.</p>	<p>There is no workaround for this problem.</p>
<p>The portal displays incorrect value for the Database State attribute of the MS SQL Database Detail attribute group if the database name contains single or double quotation marks.</p>	<p>There is no workaround for this problem.</p>
<p>The following views in the Tivoli Enterprise Portal do not display data rows for the databases that are not accessible.</p> <ul style="list-style-type: none"> • Device Detail view of the Device Detail workspace • File group Detail view of the File group Detail workspace <p>Because the database is not accessible, no data is returned by the database that can be displayed in the portal.</p>	<p>The database is not accessible when the database is in one of the following states:</p> <ul style="list-style-type: none"> • Offline • Restoring • Transition <p>Make sure that the state of the database is online.</p>

Table 31. Workspace problems and solutions (continued)

Problem	Solution
<p>The Tivoli Enterprise Portal client displays negative values for the Transactions per second and the Database Growth Percent attributes in the Database Detail view of the Database Information workspace.</p> <p>If you have multiple databases having four or more characters in the database name common, the Tivoli Enterprise Portal client displays negative values for the Transactions per second and the Database Growth Percent attributes in the Database Detail view. The SQL Server agent calculates the values of the Transactions per second and the Database Growth Percent attributes based on the previous and current sample data.</p>	<p>There is no workaround for this problem.</p>
<p>The portal does not display rows for the Statistic Name attribute values, such as connections, physical reads, physical write, I/O Errors in the MS SQL Statistics Detail attribute group.</p>	<p>None.</p>
<p>On the Tivoli Enterprise Portal, the Alert Detail view in the Errorlog Alerts workspace does not display data for more than 200 error messages with a severity level 17 or more.</p>	<p>There is no solution to this problem.</p>
<p>On the Tivoli Enterprise Portal, the Fragmentation (greater than 30%) by Index Name view does not display the index name in the bar chart when the index value is zero. The Fragmentation (greater than 30%) by Index Name view is available in the Table Detail workspace.</p>	<p>There is no solution to this problem.</p>
<p>On the Tivoli Enterprise Portal, the Fragmentation (greater than 30%) by Index Name view displays the fragmentation data in the bar chart although the fragmentation is below 30%.</p>	<p>There is no solution to this problem.</p>
<p>There is an inconsistency in the representation of the value Not_Collected in a bar chart for different workspaces.</p>	<p>There is no solution to this problem.</p>

Table 31. Workspace problems and solutions (continued)

Problem	Solution
You have installed the Tivoli Enterprise Monitoring Server on z/OS. In the Configuration panel of the History Collection Configuration window, if you have selected TEMS from the Collection Location list, then the Tivoli Enterprise Portal does not display historical data for all the attribute groups.	Reconfigure history collection, and set the collection location as TEMA for these attribute groups.

Situation troubleshooting

This section provides information about both general situation problems and problems with the configuration of situations. See the *IBM Tivoli Monitoring Troubleshooting Guide* for more information about troubleshooting for situations.

General situation problems

Table 32 lists problems that might occur with specific situations.

Table 32. Specific situation problems and solutions

Problem	Solution
You want to change the appearance of situations when they are displayed in a Workspace view.	<ol style="list-style-type: none"> 1. Right-click an item in the Navigation tree. 2. Select Situations in the pop-up menu. The Situation Editor window is displayed. 3. Select the situation that you want to modify. 4. Use the Status pull-down menu in the lower right of the window to set the status and appearance of the Situation when it triggers. Note: This status setting is not related to severity settings in IBM Tivoli Enterprise Console.
Monitoring activity requires too much disk space.	Check the RAS trace logging settings that are described in “RAS trace parameters” on page 275. For example, trace logs grow rapidly when you apply the ALL logging option.
Monitoring activity requires too many system resources.	<p>If possible, reduce the number of attributes that you monitor to a smaller set.</p> <p>Be aware that the Table Detail attribute group can consume significant system resources because it gathers data from each table and index in each database that you monitor.</p>
A formula that uses mathematical operators appears to be incorrect. For example, if you were monitoring Linux, a formula that calculates when Free Memory falls under 10 percent of Total Memory does not work: LT <code>#'Linux_VM_Stats.Total_Memory' / 10</code>	<p>This formula is incorrect because situation predicates support only logical operators. Your formulas cannot have mathematical operators.</p> <p>Note: The Situation Editor provides alternatives to math operators. Regarding the example, you can select % Memory Free attribute and avoid the need for math operators.</p>
Situations that you create display the severity UNKNOWN in IBM Tivoli Enterprise Console.	<p>For a situation to have the correct severity in TEC for those situations which are not mapped, you need to ensure that an entry exists in the tecserver.txt file for the situation and that SEVERITY is specified.</p> <p>See the “Configuring Tivoli Enterprise Console integration” chapter in the <i>IBM Tivoli Monitoring Administrator’s Guide</i> for more information.</p>

Table 32. Specific situation problems and solutions (continued)

Problem	Solution
You see the Unable to get attribute name error in the Tivoli Enterprise Monitoring Server log after creating a situation.	<p>Ensure that the agent attribute files are installed on the Tivoli Enterprise Monitoring Server.</p> <p>The following example shows a typical log entry when you have this problem:</p> <pre>(4320916A.0049-F60:kfaottev.c,1572,"Translate_ResultBuffer") \ Unable to get attribute name for tablename/column \ <UAG524400.UA4>. Ignored.</pre>
Situations are triggered in the Tivoli Enterprise Monitoring Server, but events for the situation are not sent to the Tivoli Enterprise Console server. The Tivoli Enterprise Monitoring Server is properly configured for event forwarding, and events for many other situations are sent to the event server.	<p>None. This is a limitation of the Tivoli Enterprise Monitoring Server event forwarding function. Situations that only monitor other situations do not send events to the event server.</p> <p>This condition can occur when a situation is only monitoring the status of other situations. The event forwarding function requires an attribute group reference in the situation in order to determine the correct event class to use in the event. When the situation only monitors other situations, no attribute groups are defined and the event class cannot be determined. Because the event class cannot be determined, no event is sent.</p>
The situation MS_SQL_Rem_Serv_Stat_Critical triggers when the SQL Server is down or the Microsoft SQL Server agent collector service is stopped (that means when the collector service fails to collect data.)	None.
The MS_SQL_Status_Critical situation is triggered when the SQL Server, the SQL Server agent, and the collector service of the SQL Server agent is running.	Enter the parameters values in the Table Detail Collection Settings pane in the Database Server Properties window for the data collection of the Table Detail attribute group.

Problems with configuration of situations

Table 33 lists problems that might occur with situations.

This section provides information for troubleshooting for agents. Be sure to consult the *IBM Tivoli Monitoring Troubleshooting Guide* for more general troubleshooting information.

Table 33. Problems with configuring situations that you solve in the Situation Editor

Problem	Solution
<p>Note: To get started with the solutions in this section, perform these steps:</p> <ol style="list-style-type: none"> 1. Launch the Tivoli Enterprise Portal. 2. Click Edit > Situation Editor. 3. In the tree view, choose the agent whose situation you want to modify. 4. Choose the situation in the list. The Situation Editor view is displayed. 	
The situation for a specific agent is not visible in the Tivoli Enterprise Portal.	<ul style="list-style-type: none"> • Open the Situation Editor. Access the All managed servers view. If the situation is absent, confirm that application support for Microsoft SQL Server agent has been added to the monitoring server. . If not, add application support to the server, as described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>. • Click the Action tab and check whether the situation has an automated corrective action. This action can occur directly or through a policy. The situation might be resolving so quickly that you do not see the event or the update in the graphical user interface.

Table 33. Problems with configuring situations that you solve in the Situation Editor (continued)

Problem	Solution
The monitoring interval is too long.	Access the Situation Editor view for the situation that you want to modify. Check the Sampling interval area in the Formula tab. Adjust the time interval as needed.
The situation did not activate at startup.	Manually recycle the situation as follows: <ol style="list-style-type: none"> 1. Right-click the situation and choose Stop Situation. 2. Right-click the situation and choose Start Situation. <p>Note: You can permanently avoid this problem by placing a check mark in the Run at Startup option of the Situation Editor view for a specific situation.</p>
An Alert event has not occurred even though the predicate has been properly specified.	Check the logs, reports, and workspaces.
A situation fires on an unexpected managed object.	Confirm that you have distributed and started the situation on the correct managed system.
The product did not distribute the situation to a managed system.	Click the Distribution tab and check the distribution settings for the situation.
The situation does not fire. Incorrect predicates are present in the formula that defines the situation. For example, the managed object shows a state that normally triggers a monitoring event, but the situation is not true because the wrong attribute is specified in the formula.	Click the Distribution tab and check the distribution settings for the situation. In the Formula tab, analyze predicates as follows: <ol style="list-style-type: none"> 1. Click the <i>fx</i> icon in the upper-right corner of the Formula area. The Show formula window is displayed. <ol style="list-style-type: none"> a. Confirm the following details in the Formula area at the top of the window: <ul style="list-style-type: none"> • The attributes that you intend to monitor are specified in the formula. • The situations that you intend to monitor are specified in the formula. • The logical operators in the formula match your monitoring goal. • The numerical values in the formula match your monitoring goal. b. (Optional) Click the Show detailed formula check box in the lower left of the window to see the original names of attributes in the application or operating system that you are monitoring. c. Click OK to dismiss the Show formula window. 2. (Optional) In the Formula area of the Formula tab, temporarily assign numerical values that will immediately trigger a monitoring event. The triggering of the event confirms that other predicates in the formula are valid. <p>Note: After you complete this test, you must restore the numerical values to valid levels so that you do not generate excessive monitoring data based on your temporary settings.</p>

Table 34. Problems with configuration of situations that you solve in the Workspace area

Problem	Solution
Situation events are not displayed in the Events Console view of the workspace.	Associate the situation with a navigator item. Note: The situation does not need to be displayed in the workspace. It is sufficient that the situation be associated with any navigator item.
You do not have access to a situation.	Note: You must have administrator privileges to perform these steps. <ol style="list-style-type: none"> 1. Select Edit > Administer Users to access the Administer Users window. 2. In the Users area, select the user whose privileges you want to modify. 3. In the Permissions tab, Applications tab, and Navigator Views tab, select the permissions or privileges that correspond to the user's role. 4. Click OK.

Table 34. Problems with configuration of situations that you solve in the Workspace area (continued)

Problem	Solution
A managed system seems to be offline.	<ol style="list-style-type: none"> 1. Allow several minutes to pass whenever you restart the Tivoli Enterprise Portal Server. A managed system will seem to be down while the portal server is restarting. 2. Select Physical View and highlight the Enterprise Level of the navigator tree. 3. Select View > Workspace > Managed System Status to see a list of managed systems and their status. 4. If a system is offline, check network connectivity and status of the specific system or application.

Table 35. Problems with configuration of situations that you solve in the Manage Tivoli Enterprise Monitoring Services window

Problem	Solution
After an attempt to restart the agents in the Tivoli Enterprise Portal, the agents are still not running.	Check the system status and check the appropriate IBM Tivoli Monitoring logs.
The Tivoli Enterprise Monitoring Server is not running.	Check the system status and check the appropriate IBM Tivoli Monitoring logs.

Take Action command troubleshooting

Table 36 lists general problems that might occur with Take Action commands. When each Take Action command runs it generates the log file listed in Table 24 on page 272. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 36. Take Action commands problems and solutions

Problem	Solution
Take Action commands might require several minutes to complete.	Allow several minutes. If you do not see a pop-up message advising you of completion, try to run the command manually. If you are unable to perform the Take Action command manually, see the <i>IBM Tivoli Monitoring Troubleshooting Guide</i> for general information on troubleshooting the Take Action command.
A Take Action command fails with a pop-up message containing a non-zero return value.	Check the trace log of the Take Action command. If you see a log entry containing the "insufficient authorization" string, verify that you are using one of the authorizations for Microsoft SQL Server that are listed in Chapter 6, "Take Action commands reference," on page 239.
Dump Transaction Log Take Action command fails.	The documentation might list a "backup and truncate" option. The task fails if you use this option as documented. You must use the following syntax for this option: BackupAndTruncate

Table 36. Take Action commands problems and solutions (continued)

Problem	Solution
<p>When running koqDumpDatabase take action command against a SQL Server 2005, the following (or similar) trace information might be returned in the take action's trace log:</p> <pre>(48244D4D.0000- 1AF8:koqdumpdatabase.cpp,137,"main") Calling DumpDatabase DLL API(48244D51.0000- 1AF8:koqdumpdatabase.cpp,145,"main") koqdumpdatabase has found an error:(48244D51.0001- 1AF8:koqdumpdatabase.cpp,146,"main") Error code: 1(48244D51.0002- 1AF8:koqdumpdatabase.cpp,147,"main") Error source: dmoDumpDatabase(48244D51.0003- 1AF8:koqdumpdatabase.cpp,148,"main") Error description: +48244D51.0003{ERROR:8011:101102tivoliTIV162\ TIV1620THmsdb6112005-10-14 01:54:05.00049259526700000</pre>	<p>The backup actually occurred. The error can be ignored.</p>

Tivoli Common Reporting troubleshooting

Table 37 contains a list of problems that might occur with the Tivoli Common Reporting predefined reports for Microsoft SQL Server agent.

For information about troubleshooting for the Tivoli Common Reporting tool, see http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc/tcr_welcome.html.

Table 37. Tivoli Common Reporting for Microsoft SQL Server agent problems and solutions

Problem	Solution
<p>When you simultaneously query two tables in the Query Studio interface, no data is displayed. However, when you query the two tables separately, data is displayed.</p>	<p>This problem occurs when a relationship between the tables is not defined. To resolve this problem, ensure that all the ad hoc queries have at least one identifier.</p>
<p>When you create an ad hoc query by dragging some columns in the Query Studio interface, the following error message is displayed:</p> <pre>RQP-DEF-0177 An error occurred while performing operation 'sqlPrepareWithOptions' status='-232'.</pre>	<p>This is an SQL error related to arithmetic flow. This error is generated because the average or the sum for certain columns is more than the threshold size that is defined in the database. To resolve this error, use the limited columns and add a standard timestamp while creating an ad hoc query.</p>
<p>If a view or a table for the selected summarization type does not exist in the database for a report, the report does not open and the following error message is displayed:</p> <pre>RQP-DEF-0177 An error occurred while performing operation 'sqlPrepareWithOptions' status='-232'.</pre>	<p>To resolve this problem, complete the following tasks:</p> <ul style="list-style-type: none"> • Verify that the summarization and pruning agent is working correctly. • Generate data for all the summarization types. • Verify that the warehouse is collecting historical data.

Table 37. Tivoli Common Reporting for Microsoft SQL Server agent problems and solutions (continued)

Problem	Solution
<p>When you run a report, the report is not displayed in the correct format and the following error message is displayed:</p> <p>RQP-DEF-0177 An error occurred while performing operation 'sqlPrepareWithOptions' status='-232'.</p>	<p>This problem occurs due to incorrect data source. To resolve this problem, complete the following tasks:</p> <ul style="list-style-type: none"> • Verify that the datasource configuration parameters are configured correctly. • Verify that the specified values for the parameters of the summarization and pruning agent such as database URL, driver user, and password match with the values of these parameters on the database.
<p>If data is not available in the database for the selected parameters, the following error message is displayed after querying these parameters:</p> <p>Empty data set No data returned by query. Try another set of parameters.</p>	<p>To resolve this error, complete the following tasks:</p> <ul style="list-style-type: none"> • Configure the summarization and pruning agent and verify that it is working correctly. • Generate data for all the summarization types in the database.
<p>Reports are not generated correctly in the Microsoft Excel format.</p>	<p>There are some limitations to generate reports in the Microsoft Excel format. To view these limitations, see the IBM Cognos Business Intelligence Information Center (http://publib.boulder.ibm.com/infocenter/cbi/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.ug_cr_rptstd.10.1.0.doc%2Fug_cr_rptstd_id32474excel_limitations.html).</p>
<p>When you view a report spanning multiple pages in the PDF format, the report parameters section is displayed at the top of each page.</p>	<p>No solution is available for this problem at this time.</p>
<p>Charts are not displayed correctly in Microsoft Excel 2007.</p>	<p>No solution is available for this problem at this time.</p>
<p>Labels for some charts are displayed in the HTML output, but are not displayed in the PDF output.</p>	<p>The font size is rendered differently in the HTML and the PDF output. In the PDF output, some fonts are not displayed because of the large font size. To resolve this issue, reduce the font size by completing the following steps:</p> <ol style="list-style-type: none"> 1. Open the report in Report Studio. 2. Click the chart. 3. In the chart properties, select Font. 4. Modify the font properties, such as family, size, weight, and style. 5. Save the settings, and run the report in the PDF format.
<p>The value of the Not collected attribute is incorrectly displayed as -0.01 in the warehouse database. Similarly, the value of the Not applicable attribute is incorrectly displayed as -0.02 in the warehouse database. While designing the Tivoli Common Reporting reports, the value -0.01 must be displayed as Not collected, and the value -0.02 must be displayed as Not applicable in the reports.</p> <p>The correct value for the Not collected attribute is -1 and the correct value for the Not applicable attribute is -2.</p>	<p>There is no solution to this problem.</p>

Support information

If you have a problem with your IBM software, you want to resolve it quickly.

IBM provides the following ways for you to obtain the support you need:

Online

The following websites contain troubleshooting information:

- Go to the IBM Software Support website (<http://www.ibm.com/support/entry/portal/software>) and follow the instructions.
- Go to the Application Performance Management Wiki (<http://www.ibm.com/developerworks/servicemanagement/apm/index.html>). Feel free to contribute to this wiki.

IBM Support Assistant

The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to the IBM Support Assistant website (<http://www.ibm.com/software/support/isa>).

Appendix A. Upgrading for warehouse summarization

Refer to appendix A only when you are upgrading the agent from V6.1 to V6.2.1. This appendix does not apply when you are upgrading the agent from V6.2 to V6.2.1.

The Microsoft SQL Server agent made changes to the warehouse collection and summarization characteristics for some agent attribute groups. These changes correct and improve the way warehouse data is summarized, producing more meaningful historical reports. This appendix explains those changes and the implications to your warehouse collection and reporting.

Warehouse summarization is controlled on a per-table basis. How the rows in each table are summarized is determined by a set of attributes in each table that are designated as primary keys. There is always one primary key representing the monitored resource, and data is minimally summarized based on this value. For all agents, this primary key is represented internally by the column name, ORIGINNODE; however, the external attribute name varies with each monitoring agent.

One or more additional primary keys are provided for each attribute group to further refine the level of summarization for that attribute group. For example, in an OS agent disk attribute group, a primary key might be specified for the logical disk name that allows historical information to be reported for each logical disk in a computer.

Tables in the warehouse

For a monitoring agent, there are two main types of warehouse tables:

- Raw tables:

These tables contain the raw information reported by a monitoring agent and written to the warehouse by the Warehouse Proxy agent. Raw tables are named for the attribute group that they represent, for example, MS_SQL_DATABASE_DETAIL.

- Summary tables:

These tables contain summarized information based on the raw tables and written to the warehouse by the Summarization and Pruning agent. Summarization provides aggregation results over various reporting intervals, for example, hours, days, and so on. Summary table names are based on the raw table name with an appended suffix, for example, MS_SQL_DATABASE_DETAIL_H, MS_SQL_DATABASE_DETAIL_D, and so on.

Effects on summarized attributes

When tables are summarized in the warehouse, the summary tables and summary views are created to include additional columns to report summarization information. Table 38 contains a list of the time periods and the suffixes for the summary tables and views.

Table 38. Time periods and suffixes for summary tables and views

Data collection time period	Summary table suffixes	Summary view suffixes
Hourly	_H	_HV
Daily	_D	_DV
Weekly	_W	_WV
Monthly	_M	_MV
Quarterly	_Q	_QV
Yearly	_Y	_YV

Table 39 shows the expansion to summary columns of some of the most commonly used attribute types.

Table 39. Additional columns to report summarization information

Attribute name	Aggregation type	Additional summarization columns
MyGauge	GAUGE	MIN_MyGauge MAX_MyGauge SUM_MyGauge AVG_MyGauge
MyCounter	COUNTER	TOT_MyCounter HI_MyCounter LO_MyCounter LAT_MyCounter
MyProperty	PROPERTY	LAT_Property

These additional columns are provided only for attributes that are not primary keys. In the cases when an existing attribute is changed to be a primary key, the Summarization and Pruning agent no longer creates summarization values for the attributes, but the previously created column names remain in the table with any values already provided for those columns. These columns cannot be deleted from the warehouse database, but as new data is collected, these columns will not contain values. Similarly, when the primary key for an existing attribute has its designation removed, that attribute has new summarization columns automatically added. As new data is collected, it is used to populate these new column values, but any existing summarization records do not have values for these new columns.

The overall effect of these primary key changes is that summarization information is changing. If these changes result in the old summarization records no longer making sense, you can delete them. As a part of warehouse upgrade, summary views are dropped. The views will be recreated by the Summarization and Pruning agent the next time it runs. Dropping and recreating the views ensure that they reflect the current table structure.

Upgrading your warehouse with limited user permissions

The IBM Tivoli Monitoring warehouse agents (Warehouse Proxy and Summarization and Pruning agents) can dynamically adjust warehouse table definitions based on attribute group and attribute information being loaded into the warehouse. These types of table changes must be done for this monitoring agent for one or both of the following conditions:

- The monitoring agent has added new attributes to an existing attribute group and that attribute group is included in the warehouse.
- The monitoring agent has added a new attribute group and that attribute group is included in the warehouse.

For the warehouse agents to automatically modify the warehouse table definitions, they must have permission to alter warehouse tables. You might not have granted these agents these permissions, choosing instead to manually define the raw tables and summary tables needed for the monitoring agents. Or, you might have granted these permissions initially, and then revoked them after the tables were created.

You have two options to effect the required warehouse table changes during the upgrade process:

- Grant the warehouse agents temporary permission to alter tables
If using this option, grant the permissions, start historical collection for all the desired tables, allow the Warehouse Proxy agent to add the new data to the raw tables, and allow the Summarization and Pruning agent to summarize data for all affected tables. Then, remove the permission to alter tables

- Make the warehouse table updates manually

If using this option, you must determine the table structures for the raw and summary tables. If you manually created the tables in the earlier warehouse definition, you already have a methodology and tools to assist you in this effort. You can use a similar technique to update and add new tables for this warehouse migration.

For a method of obtaining raw table schema, refer to the IBM Redbook, *Tivoli Management Services Warehouse and Reporting*, January 2007, SG24-7290. The chapter that explains warehouse tuning includes a section on creating data tables manually.

Appendix B. Event mapping

The Tivoli Event Integration Facility (EIF) interface is used to forward situation events to Tivoli Netcool/OMNIBus or Tivoli Enterprise Console.

EIF events specify an event class and the event data is specified as name value pairs that identify the name of an event slot and the value for the slot. An event class can have subclasses. IBM Tivoli Monitoring provides the base event class definitions and a set of base slots that are included in all monitoring events. Agents extend the base event classes to define subclasses that include agent-specific slots. For Microsoft SQL Server agent events, the event classes correspond to the agent attribute groups, and the agent-specific slots correspond to the attributes in the attribute group.

A description of the event slots for each event class is provided in this topic. The situation editor in the Tivoli Enterprise Portal can be used to perform custom mapping of data to EIF slots instead of using the default mapping described in this topic. For more information about EIF slot customization, see the *Tivoli Enterprise Portal User's Guide*.

Tivoli Enterprise Console requires that event classes and their slots are defined in BAROC (Basic Recorder of Objects in C) files. Each agent provides a BAROC file that contains event class definitions for the agent and is installed on the Tivoli Enterprise Monitoring Server in the TECLIB directory (`install_dir/cms/TECLIB` for Windows systems and `install_dir/tables/TEMS_hostname/TECLIB` for UNIX systems) when application support for the agent is installed. The BAROC file for the agent and the base BAROC files provided with Tivoli Monitoring must also be installed onto the Tivoli Enterprise Console. For details, see "Setting up event forwarding to Tivoli Enterprise Console" in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Each of the event classes is a child of KOQ_Base and is defined in the `oq.baroc` (version 6.3) file. The KOQ_Base event class can be used for generic rules processing for any event from the IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent..

Table 40. Overview of attribute groups to event classes and slots

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Summary	<p>MS_SQL_Server_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • server_type: STRING • server_status: STRING • server_status_enum: STRING • collection_status: STRING • collection_status_enum: STRING • data_cache_size: INTEGER • data_cache_size_enum: SRING • procedure_cache_size: INTEGER • procedure_cache_size_enum: STRING • percent_max_locks: REAL • percent_max_locks_enum: STRING • server_cpu_percent: REAL • server_cpu_percent_enum: STRING • total_os_cpu_percent: REAL • total_os_cpu_percent_enum: STRING • time_since_startup: INTEGER • time_since_startup_enum: STRING • server_version: STRING • server_version_enum: STRING • current_interval: INTEGER • current_interval_enum: STRING • sample_timestamp: STRING • description: STRING • koq_value: STRING • cache_average_free_scan: INTEGER • cache_average_free_scan_enum: STRING • cache_hit_ratio: REAL • cache_hit_ratio_enum: STRING • cache_maximum_free_page_scan: INTEGER • cache_maximum_free_page_scan_enum: STRING • cache_free_buffers: INTEGER • cache_free_buffers_enum: STRING • sqlserveragent_failed_jobs: INTEGER • sqlserveragent_failed_jobs_enum: STRING • server_cpu_percent_io: REAL • server_cpu_percent_io_enum: STRING • cpu_percent_idle: REAL • cpu_percent_idle_enum: STRING <p>(Continued on the next page.)</p>

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Summary (Continued)	<ul style="list-style-type: none"> • total_blocking_lock_requests: INTEGER • total_blocking_lock_requests_enum: STRING • total_locks_remaining: INTEGER • total_locks_remaining_enum: STRING • lazy_writes_per_sec: REAL • lazy_writes_per_sec_enum: STRING • stolen_pages: INTEGER • stolen_pages_enum: STRING • stolen_pages_growth: INTEGER • stolen_pages_growth_enum: STRING • page_reads_per_sec: REAL • page_reads_per_sec_enum: STRING • readahead_pages_per_sec: REAL • readahead_pages_per_sec_enum: STRING • page_writes_per_sec: REAL • page_writes_per_sec_enum: STRING • checkpoint_pages_per_sec: REAL • checkpoint_pages_per_sec_enum: STRING • repl_distribution_delivery_latency: INTEGER • repl_distribution_delivery_latency_enum: STRING • repl_distribution_delivered_cmds_per_sec: REAL • repl_distribution_delivered_cmds_per_sec_enum: STRING • repl_distribution_delivered_trans_per_sec: REAL • repl_distribution_delivered_trans_per_sec_enum: STRING • repl_logreader_delivery_latency: INTEGER • repl_logreader_delivery_latency_enum: STRING • repl_logreader_delivered_cmds_per_sec: REAL • repl_logreader_delivered_cmds_per_sec_enum: STRING • repl_logreader_delivered_trans_per_sec: REAL • repl_logreader_delivered_trans_per_sec_enum: STRING • repl_merge_conflicts_per_sec: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Summary (Continued)	<ul style="list-style-type: none"> • repl_merge_conflicts_per_sec_enum: STRING • repl_merge_uploaded_changes_per_sec: REAL • repl_merge_uploaded_changes_per_sec_enum: STRING • repl_merge_downloaded_changes_per_sec: REAL • repl_merge_downloaded_changes_per_sec_enum: STRING • cpu_idle_ms: INTEGER • cpu_idle_sec_enum: STRING • cpu_idle_delta: INTEGER • cpu_idle_delta_enum: STRING • cpu_usage_sec: INTEGER • cpu_usage_sec_enum: STRING • cpu_usage_delta: INTEGER • cpu_usage_delta_enum: STRING • total_server_memory: INTEGER • total_server_memory_enum: STRING • sql_server_ver: STRING • sql_server_ver_enum: STRING • total_pages: INTEGER • total_pages_enum: STRING • mixed_page_allocations_per_sec: REAL • mixed_page_allocations_per_sec_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Detail	<p>MS_SQL_Server_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • server_type: STRING • server_type_enum: STRING • server_version: STRING • server_version_enum: STRING • server_status: STRING • server_status_enum: STRING • current_locks: INTEGER • current_locks_enum: STRING • max_locks_allowed: INTEGER • max_locks_allowed_enum: STRING • percent_max_locks: REAL • percent_max_locks_enum: STRING • data_cache_size: INTEGER • data_cache_size_enum: STRING • procedure_cache_size: INTEGER • procedure_cache_size_enum: STRING • procedure_cache_percent: REAL • procedure_cache_percent_enum: STRING • startup_timestamp: STRING • startup_timestamp_enum: STRING • time_since_startup: INTEGER • time_since_startup_enum: STRING • error_log_name: STRING • error_log_name_enum: STRING • error_log_size: INTEGER • error_log_size_enum: STRING • os_type: STRING • os_type_enum: STRING • os_version: STRING • os_version_enum: STRING • sample_timestamp: STRING • procedure_buffers_percent_active: REAL • procedure_buffers_percent_active_enum: STRING • procedure_buffers_total: INTEGER • procedure_buffers_total_enum: STRING • procedure_buffers_percent_used: REAL • procedure_buffers_percent_used_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Detail (Continued)	<ul style="list-style-type: none"> • procedure_cache_percent_active: REAL • procedure_cache_percent_active_enum: STRING • procedure_cache_pages: INTEGER • procedure_cache_pages_enum: STRING • procedure_cache_percent_used: REAL • procedure_cache_percent_used_enum: STRING • error_log_name_u: STRING • error_log_name_u_enum: STRING • sql_server_ver: STRING • sql_server_ver_enum: STRING
ITM_MS_SQL_Database_Summary	MS_SQL_Database_Summary attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: INTEGER • server: STRING • host_name: STRING • total_databases: INTEGER • minimum_pct_data_freespace: REAL • minimum_pct_data_freespace_enum: STRING • minimum_pct_log_freespace: REAL • minimum_pct_log_freespace_enum: STRING • total_dbs_in_error: INTEGER • total_dbs_read_only: INTEGER • total_dbs_dbo_only: INTEGER • total_dbs_single_user: INTEGER • total_dbs_no_free_space_accounting: INTEGER • current_interval: INTEGER • sample_timestamp: INTEGER • description: STRING • koq_value: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Database_Detail	<p>MS_SQL_Database_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_name: STRING • database_name_enum: STRING • db_id: STRING • db_id_enum: STRING • db_owner: STRING • error_status: STRING • error_status_enum: STRING • dbo_only_access: STRING • dbo_only_access_enum: STRING • single_user_access: STRING • single_user_access_enum: STRING • read_only_access: STRING • read_only_access_enum: STRING • select_into_bulkcopy_enabled: STRING • select_into_bulkcopy_enabled_enum: STRING • dump_tran_date: STRING • dump_tran_date_enum: STRING • free_space_accounting_suppressed: STRING • free_space_accounting_suppressed_enum: STRING • no_ckpt_after_recovery: STRING • no_ckpt_after_recovery_enum: STRING • data_size: REAL • data_size_enum: STRING • log_size: REAL • log_size_enum: STRING • data_freespace: REAL • data_freespace_enum: STRING • log_freespace: REAL • log_freespace_enum: STRING • data_freespace_percent: REAL • data_freespace_percent_enum: STRING • log_freespace_percent: REAL • log_freespace_percent_enum: STRING • total_devices: INTEGER • total_devices_enum: STRING <p>(Continued)</p>

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Database_Detail (Continued)	MS_SQL_Database_Detail attribute group <ul style="list-style-type: none"> • sample_timestamp: STRING • sample_timestamp_enum: STRING • database_space_percent_used: REAL • database_space_percent_used_enum: STRING • database_status: INTEGER • database_status_enum: STRING • log_space_percent_used: REAL • log_space_percent_used_enum: STRING • oldest_open_transaction: INTEGER • oldest_open_transaction_enum: STRING • suspect_database: INTEGER • suspect_database_enum: STRING • replicated_transaction_rate: REAL • replicated_transaction_rate_enum: STRING • replicated_transactions: INTEGER • replicated_transactions_enum: STRING • replication_latency: INTEGER • replication_latency_enum: STRING • db_owner_u: STRING • database_name_u: STRING • log_maximum_growth_size: REAL • log_maximum_growth_size_enum: STRING • database_maximum_growth_size: REAL • database_maximum_growth_size_enum: STRING • last_database_size: REAL • last_database_size_enum: STRING • database_growth_percent: REAL • database_growth_percent_enum: STRING • active_transactions: INTEGER • active_transactions_enum: STRING • transactions_per_second: REAL • transactions_per_second_enum: STRING • log_file_names: STRING • data_file_names: STRING • write_transactions_per_sec: REAL • write_transactions_per_sec_enum: STRING • long_running_transaction_process_id: INTEGER

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Database_Detail (Continued)	(Continued) <ul style="list-style-type: none"> • long_running_transaction_process_id_enum: STRING • long_running_transaction_time_per_sec: REAL • long_running_transaction_time_per_sec_enum: STRING • long_running_transaction_name: STRING • long_running_transaction_name_enum: STRING • database_state: STRING • database_state_enum: STRING • table_count: INTEGER • table_count_enum: STRING • access_validation: INTEGER • access_validation_enum: STRING • log_growths: INTEGER • log_growths_enum: STRING • log_bytes_flushed_per_sec: REAL • log_bytes_flushed_per_sec_enum: STRING • data_files_freespace: REAL • data_files_freespace_enum: STRING • aggregate_database_freespace: REAL • aggregate_database_freespace_enum: STRING • aggregate_database_freespace_pct: REAL • aggregate_database_freespace_pct_enum: STRING;

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Device_Detail	<p>MS_SQL_Device_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_version: STRING • host_name: STRING • database_name: STRING • database_name_enum: STRING • device_name: STRING • device_size: REAL • device_size_enum: STRING • device_free_space: REAL • device_free_space_enum: STRING • device_free_space_percent: REAL • device_free_space_percent_enum: STRING • device_type: STRING • device_type_enum: STRING • physical_device_name: STRING • mirror_device_name: STRING • mirror_device_name_enum: STRING • sample_timestamp: STRING • database_name_u: STRING • database_name_u_enum: STRING • device_name_u: STRING • physical_device_name_u: STRING • sql_server_ver: STRING
ITM_MS_SQL_Error_Event_Detail	<p>MS_SQL_Error_Event_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • message_timestamp: STRING • message_age: INTEGER • error_id: STRING • severity_level: INTEGER • sql_state_code: INTEGER • message_issuer: STRING • sample_timestamp: STRING • message_text_u: STRING • error_spid: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Lock_Detail	<p>MS_SQL_Lock_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • type: STRING • type_enum: STRING • process_holding_lock: INTEGER • process_holding_lock_enum: STRING • database_name: STRING • database_id: INTEGER • database_id_enum: STRING • table_name: STRING • table_name_enum: STRING • page_number: INTEGER • page_number_enum: STRING • sample_timestamp: STRING • lock_resource_type_e: INTEGER • lock_resource_type_e_enum: STRING • lock_request_status_e: INTEGER • lock_request_status_e_enum: STRING • database_name_u: STRING • table_name_u: STRING
ITM_MS_SQL_Lock_Conflict_Detail	<p>MS_SQL_Lock_Conflict_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_name: STRING • database_id: INTEGER • server_user_id: STRING • client_user_id: STRING • client_group_id: STRING • requestor_process_id: INTEGER • blocking_process_id: INTEGER • sample_timestamp: STRING • database_name_u: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Text	<p>MS_SQL_Text attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_name: STRING • client_host_name: STRING • client_process_id: STRING • process_id: INTEGER • total_cpu_time: REAL • process_status: STRING • text: STRING • sample_timestamp: STRING • database_name_u: STRING • text_u: STRING
ITM_MS_SQL_Process_Summary	<p>MS_SQL_Process_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • server_cpu_percent_system: REAL • server_cpu_percent_application: REAL • total_processes: INTEGER • total_processes_blocked: INTEGER • percent_processes_blocked: REAL • total_processes_infected: INTEGER • percent_processes_infected: REAL • total_processes_bad: INTEGER • percent_processes_bad: REAL • total_processes_stopped: INTEGER • percent_processes_stopped: REAL • total_processes_locksleep: INTEGER • percent_processes_locksleep: REAL • total_other_sleep: INTEGER • percent_processes_othersleep: REAL • percent_processes_sleeping: REAL • total_log_suspend: INTEGER • current_interval: INTEGER • sample_timestamp: STRING • description: STRING • koq_value: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Process_Detail	MS_SQL_Process_Detail attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • process_id: INTEGER • database_name: STRING • process_status: STRING • total_cpu_time: REAL • current_cpu_percent_used: REAL • total_disk_io: INTEGER • total_memory_alloc: INTEGER • blocking_process_id: STRING • os_process_id: STRING • server_user_id: STRING • command: STRING • client_host_name: STRING • client_process_id: STRING • program_name: STRING • client_user_id: STRING • client_group_id: STRING • login_name: STRING • sample_timestamp: STRING • database_name_u: STRING • command_u: STRING • program_name_u: STRING • koq_duration: INTEGER • nt_user: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Problem_Summary	<p>MS_SQL_Problem_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • age_of_last_error: INTEGER • maximum_sev_level: STRING • maximum_sev_timestamp: STRING • maximum_sev_current_interval: STRING • total_error_messages: INTEGER • total_errors_current_interval: INTEGER • total_errors_high_sev: INTEGER • total_errors_other: INTEGER • error_log_size: INTEGER • current_interval: INTEGER • sample_timestamp: STRING • description: STRING • koq_value: STRING
ITM_MS_SQL_Problem_Detail	<p>MS_SQL_Problem_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • message_timestamp: STRING • message_age: INTEGER • error_id: STRING • severity_level: STRING • sql_state_code: INTEGER • message_issuer: STRING • message_text: STRING • sample_timestamp: STRING • message_text_u: STRING • error_spid: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Statistics_Summary	<p>MS_SQL_Statistics_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • total_os_cpu_percent_busy: REAL • total_os_io_percent_busy: REAL • current_logons: INTEGER • percent_max_logons_active: REAL • total_logons_per_second: REAL • max_user_connections_allowed: INTEGER • physical_reads_per_second: REAL • physical_writes_per_second: REAL • disk_io_current_interval: INTEGER • io_errors_since_startup: INTEGER • io_errors_current_interval: INTEGER • percent_io_errors_current_interval: REAL • current_interval: INTEGER • sample_timestamp: STRING • description: STRING • koq_value: STRING • client_count: INTEGER • client_count_percent_used: REAL • network_read_rate: INTEGER • network_write_rate: INTEGER • percent_user_connections_remaining: REAL • total_logouts_per_second: REAL
ITM_MS_SQL_Statistics_Detail	<p>MS_SQL_Statistics_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • statistic_name: STRING • statistic_name_enum: STRING • current_value: INTEGER • total_since_startup: INTEGER • average_value_per_second: REAL • maximum_seen: INTEGER • minimum_seen: INTEGER • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Remote_Servers	MS_SQL_Remote_Servers attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • remote_server_id: STRING • remote_server_id_enum: STRING • remote_server_status: STRING • remote_server_status_enum: STRING • remote_server_name: STRING • remote_server_name_enum: STRING • network_name: STRING • network_name_enum: STRING • connection_level: INTEGER • connection_level_enum: STRING • current_interval: INTEGER • current_interval_enum: STRING • sample_timestamp: STRING
ITM_MS_SQL_Configuration	MS_SQL_Configuration attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • config_parameter: STRING • parm_type: STRING • parm_type_enum: STRING • config_value: STRING • run_value: STRING • minimum_value: STRING • maximum_value: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Enterprise_View	<p>MS_SQL_Server_Enterprise_View attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • collection_status: STRING • collection_status_enum: STRING • percent_max_locks: REAL • percent_max_locks_enum: STRING • server_cpu_percent: REAL • server_cpu_percent_enum: STRING • total_os_cpu_percent: REAL • total_os_cpu_percent_enum: STRING • time_since_startup: INTEGER • time_since_startup_enum: STRING • server_version: STRING • server_version_enum: STRING • total_databases: INTEGER • total_databases_enum: STRING • minimum_pct_data_freespace: REAL • minimum_pct_data_freespace_enum: STRING • minimum_pct_log_freespace: REAL • minimum_pct_log_freespace_enum: STRING • total_locks: INTEGER • total_locks_enum: STRING • total_lock_conflicts: INTEGER • total_lock_conflicts_enum: STRING • database_max_locks: STRING • database_max_locks_enum: STRING • database_max_blocks: STRING • database_max_blocks_enum: STRING • table_max_locks: STRING • table_max_locks_enum: STRING • percent_process_block: REAL • percent_process_block_enum: STRING • server_cpu_percent_system: REAL • server_cpu_percent_system_enum: STRING • server_cpu_percent_application: REAL • server_cpu_percent_application_enum: STRING • total_processes: INTEGER • total_processes_enum: STRING <p>(Continued on the next page)</p>

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Enterprise_View(Continued)	<p>MS_SQL_Server_Enterprise_View attribute group</p> <ul style="list-style-type: none"> • total_processes_blocked: INTEGER • total_processes_blocked_enum: STRING • age_of_last_error: INTEGER • age_of_last_error_enum: STRING • maximum_sev_level: STRING • maximum_sev_level_enum: STRING • maximum_sev_timestamp: STRING • maximum_sev_timestamp_enum: STRING • total_errors_high_sev: INTEGER • total_errors_high_sev_enum: STRING • error_log_size: INTEGER • error_log_size_enum: STRING • current_logons: INTEGER • current_logons_enum: STRING • percent_max_logons_active: REAL • percent_max_logons_active_enum: STRING • physical_reads_per_second: REAL • physical_reads_per_second_enum: STRING • physical_writes_per_second: REAL • physical_writes_per_second_enum: STRING • io_errors_current_interval: INTEGER • io_errors_current_interval_enum: STRING • sample_timestamp: STRING • server_status: STRING • server_status_enum: STRING • database_max_locks_u: STRING • database_max_locks_u_enum: STRING • database_max_blocks_u: STRING • database_max_blocks_u_enum: STRING • table_max_locks_u: STRING • table_max_locks_u_enum: STRING • sql_server_ver: STRING • sql_server_ver_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Table_Detail	MS_SQL_Table_Detail attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_name: STRING • database_name_enum: STRING • database_id: INTEGER • database_id_enum: STRING • table_name: STRING • table_name_enum: STRING • table_id: INTEGER • table_id_enum: STRING • index_name: STRING • index_name_enum: STRING • index_id: INTEGER • index_id_enum: STRING • optimizer_statistics_age: INTEGER • optimizer_statistics_age_enum: STRING • fragmentation: REAL • fragmentation_enum: STRING • space_used: REAL • space_used_enum: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Lock_Summary	<p>MS_SQL_Lock_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • exclusive_extent_locks: INTEGER • extent_locks: INTEGER • update_extent_locks: INTEGER • exclusive_intent_locks: INTEGER • shared_intent_locks: INTEGER • intent_locks: INTEGER • exclusive_page_locks: INTEGER • shared_page_locks: INTEGER • page_locks: INTEGER • update_page_locks: INTEGER • exclusive_table_locks: INTEGER • shared_table_locks: INTEGER • table_locks: INTEGER • exclusive_locks: INTEGER • shared_locks: INTEGER • sample_timestamp: STRING • table_lock_escalations_per_sec: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Job_Summary	<p>MS_SQL_Job_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • running_jobs_current_interval: INTEGER • running_jobs_current_interval_enum: STRING • max_running_jobs_interval: INTEGER • max_running_jobs_interval_enum: STRING • jobs_failed_current_interval: INTEGER • jobs_failed_current_interval_enum: STRING • max_jobs_failed_interval: INTEGER • max_jobs_failed_interval_enum: STRING • jobs_failed_since_startup: INTEGER • jobs_failed_since_startup_enum: STRING • sample_timestamp: STRING • active_jobs: INTEGER • active_jobs_enum: STRING • failed_jobs: INTEGER • failed_jobs_enum: STRING • job_success_rate: REAL • job_success_rate_enum: STRING • queued_jobs: INTEGER • queued_jobs_enum: STRING • successful_jobs: INTEGER • successful_jobs_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Job_Detail	<p>MS_SQL_Job_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • job_name: STRING • job_id: STRING • enabled: INTEGER • enabled_enum: STRING • current_status: INTEGER • current_status_enum: STRING • last_run_timestamp: STRING • last_run_outcome: INTEGER • last_run_outcome_enum: STRING • next_run_timestamp: STRING • sample_timestamp: STRING • job_category_id: INTEGER • job_category_id_enum: STRING • job_category_name: STRING • job_status: INTEGER • job_status_enum: STRING • job_duration: INTEGER • job_duration_enum: STRING • job_error_code: STRING • job_error_code_enum: STRING
ITM_MS_SQL_Services_Detail	<p>MS_SQL_Services_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • service_name: STRING • service_state: INTEGER • service_state_enum: STRING • service_status: INTEGER • service_status_enum: STRING • service_start_mode: INTEGER • service_start_mode_enum: STRING • service_type: INTEGER • service_type_enum: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Filegroup_Detail	MS_SQL_Filegroup_Detail attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_name: STRING • filegroup_id: INTEGER • filegroup_name: STRING • filegroup_size: REAL • filegroup_space_percent_used: REAL • filegroup_max_growth_size: REAL • sample_timestamp: STRING
ITM_MS_SQL_Lock_Resource_Type_Summary	MS_SQL_Lock_Resource_Type_Summary attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • lock_resource_type: INTEGER • lock_resource_type_enum: STRING • lock_requests_per_second: REAL • lock_timeouts_per_second: REAL • lock_wait_time: INTEGER • number_deadlocks_per_second: REAL • sample_timestamp: STRING • average_wait_time_locks: INTEGER

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Workload_Group_Stats	<p>MS_SQL_Workload_Group_Stats attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • workload_group_name: STRING • workload_group_name_enum: STRING • cpu_usage_percent: REAL • cpu_usage_percent_enum: STRING • queued_requests: REAL • queued_requests_enum: STRING • requests_completed_per_sec: REAL • requests_completed_per_sec_enum: STRING • active_requests: REAL • active_requests_enum: STRING • max_requests_cpu_time: REAL • max_requests_cpu_time_enum: STRING • blocked_tasks: REAL • blocked_tasks_enum: STRING • reduced_memory_grants_per_sec: REAL • reduced_memory_grants_per_sec_enum: STRING • max_requests_memory_grant: REAL • max_requests_memory_grant_enum: STRING • query_optimization_per_sec: REAL • query_optimization_per_sec_enum: STRING • suboptimal_plans_per_sec: REAL • suboptimal_plans_per_sec_enum: STRING • active_parallel_threads: REAL • active_parallel_threads_enum: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Resource_Pool_Stats	<p>MS_SQL_Resource_Pool_Stats attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • resource_pool_name: STRING • resource_pool_name_enum: STRING • cpu_usage_percent: REAL • cpu_usage_percent_enum: STRING • cpu_usage_target_percent: REAL • cpu_usage_target_percent_enum: STRING • cpu_control_effect_percent: REAL • cpu_control_effect_percent_enum: STRING • cache_memory_target: REAL • cache_memory_target_enum: STRING • compile_memory_target: REAL • compile_memory_target_enum: STRING • query_exec_memory_target: REAL • query_exec_memory_target_enum: STRING • memory_grants_per_sec: REAL • memory_grants_per_sec_enum: STRING • memory_grant_timeouts_per_sec: REAL • memory_grant_timeouts_per_sec_enum: STRING • active_memory_grants_count: REAL • active_memory_grants_count_enum: STRING • active_memory_grant_amount: REAL • active_memory_grant_amount_enum: STRING • pending_memory_grants_count: REAL • pending_memory_grants_count_enum: STRING • max_memory: REAL • max_memory_enum: STRING • used_memory: REAL • used_memory_enum: STRING • target_memory: REAL • target_memory_enum: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Database_Mirroring	<p>MS_SQL_Database_Mirroring attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • mirror_name: STRING • mirror_name_enum: STRING • synchronization_state: INTEGER • synchronization_state_enum: STRING • mirroring_state: INTEGER • mirroring_state_enum: STRING • mirroring_role: INTEGER • mirroring_role_enum: STRING • bytes_sent_per_sec: REAL • bytes_sent_per_sec_enum: STRING • pages_sent_per_sec: REAL • pages_sent_per_sec_enum: STRING • sends_per_sec: REAL • sends_per_sec_enum: STRING • transaction_delay: REAL • transaction_delay_enum: STRING • redo_queue: REAL • redo_queue_enum: STRING • redo_bytes_per_sec: REAL • redo_bytes_per_sec_enum: STRING • log_send_queue: REAL • log_send_queue_enum: STRING • bytes_received_per_sec: REAL • bytes_received_per_sec_enum: STRING • receives_per_sec: REAL • receives_per_sec_enum: STRING • log_bytes_received_per_sec: REAL • log_bytes_received_per_sec_enum: STRING • log_bytes_sent_per_sec: REAL • log_bytes_sent_per_sec_enum: STRING • send_receive_ack_time: REAL • send_receive_ack_time_enum: STRING • log_compressed_bytes_rcvd_per_sec: REAL • log_compressed_bytes_rcvd_per_sec_enum: STRING <p>(Continued on the next page)</p>

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Database_Mirroring (Continued)	<ul style="list-style-type: none"> • log_compressed_bytes_sent_per_sec: REAL • log_compressed_bytes_sent_per_sec_enum: STRING • mirrored_write_transactions_per_sec: REAL • mirrored_write_transactions_per_sec_enum: STRING • log_scanned_for_undo: REAL • log_scanned_for_undo_enum: STRING • log_remaining_for_undo: REAL • log_remaining_for_undo_enum: STRING • log_bytes_sent_from_cache_per_sec: REAL • log_bytes_sent_from_cache_per_sec_enum: STRING • log_bytes_redone_from_cache_per_sec: REAL • log_bytes_redone_from_cache_per_sec_enum: STRING • log_send_flow_control_time: REAL • log_send_flow_control_time_enum: STRING • log_harden_time: REAL • log_harden_time_enum: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Memory_Manager	<p>MS_SQL_Memory_Manager attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • connection_memory: REAL • connection_memory_enum: STRING • granted_workspace_memory: REAL • granted_workspace_memory_enum: STRING • lock_memory: REAL • lock_memory_enum: STRING • lock_blocks_allocated: REAL • lock_blocks_allocated_enum: STRING • lock_owner_blocks_allocated: REAL • lock_owner_blocks_allocated_enum: STRING • lock_blocks: REAL • lock_blocks_enum: STRING • lock_owner_blocks: REAL • lock_owner_blocks_enum: STRING • maximum_workspace_memory: REAL • maximum_workspace_memory_enum: STRING • memory_grants_outstanding: REAL • memory_grants_outstanding_enum: STRING • memory_grants_pending: REAL • memory_grants_pending_enum: STRING • optimizer_memory: REAL • optimizer_memory_enum: STRING • sql_cache_memory: REAL • sql_cache_memory_enum: STRING • target_server_memory: REAL • target_server_memory_enum: STRING • total_server_memory: REAL • total_server_memory_enum: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Log_Shipping_DB_Details	MS_SQL_Log_Shipping_DB_Details attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server_type: INTEGER • server_type_enum: STRING • server: STRING • database_name: STRING • database_id: INTEGER • threshold: INTEGER • threshold_exceeded: INTEGER • threshold_exceeded_enum: STRING • last_file_processed: STRING • last_file_processed_time: STRING • message: STRING • koq_source: STRING • sample_timestamp: STRING
ITM_MS_SQL_Log_Shipping_Errors	MS_SQL_Log_Shipping_Errors attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • agent_type: INTEGER • agent_type_enum: STRING • server: STRING • database_name: STRING • database_id: INTEGER • sequence_number: INTEGER • log_time: STRING • message: STRING • error_source: STRING • sample_timestamp: STRING
ITM_MS_SQL_Log_Shipping_Jobs_Detail	MS_SQL_Log_Shipping_Jobs_Detail attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • agent_type: INTEGER • agent_type_enum: STRING • server: STRING • database_id: INTEGER • database_name: STRING • session_id: INTEGER • job_status: INTEGER • job_status_enum: STRING • message: STRING • log_time: STRING • job_name: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Enterprise_ Database_Mirroring	MS_SQL_Enterprise_Database_Mirroring attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • mirror_name: STRING • mirror_name_enum: STRING • synchronization_state: INTEGER • synchronization_state_enum: STRING • mirroring_state: INTEGER • mirroring_state_enum: STRING • mirroring_role: INTEGER • mirroring_role_enum: STRING • bytes_sent_per_sec: REAL • bytes_sent_per_sec_enum: STRING • pages_sent_per_sec: REAL • pages_sent_per_sec_enum: STRING • sends_per_sec: REAL • sends_per_sec_enum: STRING • transaction_delay: REAL • transaction_delay_enum: STRING • redo_queue: REAL • redo_queue_enum: STRING • redo_bytes_per_sec: REAL • redo_bytes_per_sec_enum: STRING • log_send_queue: REAL • log_send_queue_enum: STRING • bytes_received_per_sec: REAL • bytes_received_per_sec_enum: STRING • receives_per_sec: REAL • receives_per_sec_enum: STRING • log_bytes_received_per_sec: REAL • log_bytes_received_per_sec_enum: STRING • log_bytes_sent_per_sec: REAL • log_bytes_sent_per_sec_enum: STRING • send_receive_ack_time: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Enterprise_ Database_Mirroring (Continued)	<ul style="list-style-type: none"> • send_receive_ack_time_enum: STRING • log_compressed_bytes_rcvd_per_sec: REAL • log_compressed_bytes_rcvd_per_sec_enum: STRING • log_compressed_bytes_sent_per_sec: REAL • log_compressed_bytes_sent_per_sec_enum: STRING • mirrored_write_transactions_per_sec: REAL • mirrored_write_transactions_per_sec_enum: STRING • log_scanned_for_undo: REAL • log_scanned_for_undo_enum: STRING • log_remaining_for_undo: REAL • log_remaining_for_undo_enum: STRING • log_bytes_sent_from_cache_per_sec: REAL • log_bytes_sent_from_cache_per_sec_enum: STRING • log_bytes_redone_from_cache_per_sec: REAL • log_bytes_redone_from_cache_per_sec_enum: STRING • log_send_flow_control_time: REAL • log_send_flow_control_time_enum: STRING • log_harden_time: REAL • log_harden_time_enum: STRING • sample_timestamp: STRING
ITM_MS_SQL_Enterprise_ Log_Shipping_DB_Details	<p>MS_SQL_Enterprise_Log_Shipping_DB_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server_type: INTEGER • server_type_enum: STRING • server: STRING • database_name: STRING • database_id: INTEGER • threshold: INTEGER • threshold_exceeded: INTEGER • threshold_exceeded_enum: STRING • last_file_processed: STRING • last_file_processed_time: STRING • message: STRING • koq_source: STRING • sample_timestamp: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Batch_Stats	MS_SQL_Batch_Stats attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • auto_param_attempts: REAL • auto_param_attempts_enum: STRING • batch_reuests: REAL • batch_reuests_enum: STRING • failed_auto_params: REAL • failed_auto_params_enum: STRING • safe_auto-params: REAL • safe_auto-params_enum: STRING • server_attention_rate: REAL • server_attention_rate_enum: STRING • server_compilations: REAL • server_compilations_enum: STRING • unsafe_auto_params: REAL • unsafe_auto_params_enum: STRING • server_re_compilations: REAL • server_re_compilations_enum: STRING • forced_params: REAL • forced_params_enum: STRING • server: STRING • host_name: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Properties	<p>MS_SQL_Server_Properties attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_version:STRING • server_version_enum: STRING • collation: STRING • collation_enum: STRING • collation_id: INTEGER • collation_id_enum: STRING • edition: STRING • edition_enum: STRING • edition_id: REAL • edition_id_enum: STRING • engine_edition: INTEGER • engine_edition_enum: STRING • is_clustered: INTEGER • is_clustered_enum: STRING • isintegrated_security_only: INTEGER • isintegrated_security_only_enum: STRING • is_single_user: INTEGER • is_single_user_enum: STRING • license_type: STRING • license_type_enum: STRING • num_licenses: INTEGER • num_licenses_enum: STRING • product_level: INTEGER • product_level_enum: STRING • physical_netbios_name: STRING • physical_netbios_name_enum: STRING • qualified_server_name: STRING • qualified_server_name_enum: STRING • filestream_share_name: STRING • filestream_share_name_enum: STRING • filestream_level: INTEGER • filestream_level_enum: STRING • is_hadr_enabled: INTEGER • is_hadr_enabled_enum: STRING • hadr_manager_status:INTEGER • hadr_manager_status_enum: STRING;

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Server_Transactions_Summary	<p>MS_SQL_Server_Transactions_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • tempdb_free_space: INTEGER • tempdb_free_space_enum: STRING • longest_transaction_running_time: INTEGER • longest_transaction_running_time_enum: STRING • total_non-snapshot_version_transactions: INTEGER • total_non-snapshot_version_transactions_enum: STRING • total_snapshot_transactions: INTEGER • total_snapshot_transactions_enum: STRING • total_active_transactions: INTEGER • total_active_transactions_enum: STRING • update_conflict_ratio: REAL • update_conflict_ratio_enum: STRING • total_update_snapshot_transactions: INTEGER • total_update_snapshot_transactions_enum: STRING • version_cleanup_rate: REAL • version_cleanup_rate_enum: STRING • version_generation_rate: REAL • version_generation_rate_enum: STRING • version_store_size: INTEGER • version_store_size_enum: STRING • version_store_unit_count: INTEGER • version_store_unit_count_enum: STRING • version_store_unit_creation: INTEGER • version_store_unit_creation_enum: STRING • version_store_unit_truncation: INTEGER • version_store_unit_truncation_enum: STRING;

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Service_Broker_Activation	<p>MS_SQL_Service_Broker_Activation attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • database_name_u: STRING • database_name_u_enum: STRING • stored_procedure_invoked_per_sec: REAL • stored_procedure_invoked_per_sec_enum: STRING • task_limit_reached_per_sec: REAL • task_limit_reached_per_sec_enum: STRING • total_task_limit_reached: INTEGER • total_task_limit_reached_enum: STRING • tasks_aborted_per_sec: REAL • tasks_aborted_per_sec_enum: STRING • tasks_started_per_sec: REAL • tasks_started_per_sec_enum: STRING • total_tasks_started: INTEGER • total_tasks_started_enum: STRING;
ITM_MS_SQL_Service_Broker_Statistics	<p>MS_SQL_Service_Broker_Statistics attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • broker_transaction_rollbacks: REAL • broker_transaction_rollbacks_enum: STRING • enqueued_local_messages_per_sec: REAL • enqueued_local_messages_per_sec_enum: STRING • total_enqueued_local_messages: REAL • total_enqueued_local_messages_enum: STRING • total_enqueued_messages: REAL • total_enqueued_messages_enum: STRING • total_enqueued_transport_messages: REAL • total_enqueued_transport_messages_enum: STRING • total_forwarded_messages: REAL • total_forwarded_messages_enum: STRING • forwarded_pending_messages: REAL • forwarded_pending_messages_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Service_Broker_Statistics (Continued)	(Continued) <ul style="list-style-type: none"> • total_forwarded_messages_discarded: REAL • total_forwarded_messages_discarded_enum: STRING • forwarded_messages_discarded_per_sec: REAL • forwarded_messages_discarded_per_sec_enum: STRING • enqueued_messages_per_sec: REAL • enqueued_messages_per_sec_enum: STRING • enqueued_transport_messages_per_sec: REAL • enqueued_transport_messages_per_sec_enum: STRING • sql_receives_per_sec: REAL • sql_receives_per_sec_enum: STRING • sql_sends_per_sec: REAL • sql_sends_per_sec_enum: STRING • total_sql_receives: REAL • total_sql_receives_enum: STRING • total_sql_sends: REAL • total_sql_sends_enum: STRING;

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Service_Broker_Transport	MS_SQL_Service_Broker_Transport attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • current_bytes_received_io: REAL • current_bytes_received_io_enum: STRING • current_bytes_send_io: REAL • current_bytes_send_io_enum: STRING • current_message_fragments_send_io: REAL • current_message_fragments_send_io_enum: STRING • message_fragment_received_per_sec: REAL • message_fragment_received_per_sec_enum: STRING • message_fragments_send_per_sec: REAL • message_fragments_send_per_sec_enum: STRING • avg_size_message_fragments_received: REAL • avg_size_message_fragments_received_enum: STRING • avg_size_message_fragments_send: REAL • avg_size_message_fragments_send_enum: STRING • total_open_connection_count: REAL • total_open_connection_count_enum: STRING • pending_bytes_received_io: REAL • pending_bytes_received_io_enum: STRING • pending_bytes_send_io: REAL • pending_bytes_send_io_enum: STRING • pending_message_fragments_received_io: REAL • pending_message_fragments_received_io_enum: STRING • pending_message_fragments_send_io: REAL • pending_message_fragments_send_io_enum: STRING • receive_ios_per_sec: REAL • receive_ios_per_sec_enum: STRING • send_ios_per_sec: REAL • send_ios_per_sec_enum: STRING;

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Table_Summary	<p>MS_SQL_Table_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • total_tables: INTEGER • total_tables_enum: STRING • maximum_optimizer_statistics_age: INTEGER • maximum_optimizer_statistics_age_enum: STRING • minimum_optimizer_statistics_age: INTEGER • minimum_optimizer_statistics_age_enum: STRING • maximum_fragmentation_percent: REAL • maximum_fragmentation_percent_enum: STRING • minimum_fragmentation_percent: REAL • minimum_fragmentation_percent_enum: STRING • maximum_fragmentation_percent_since_startup: REAL • maximum_fragmentation_percent_since_startup_enum: STRING • maximum_table_size_mb: REAL • maximum_table_size_mb_enum: STRING • minimum_table_size_mb: REAL • minimum_table_size_mb_enum: STRING • last_collection_duration: REAL • last_collection_duration_enum: STRING • last_collection_start_time: STRING • last_collection_start_time_enum: STRING • total_filetables: INTEGER • total_filetables_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Integration_Service_Details	<p>MS_SQL_Integration_Service_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_enum: STRING • buffer_memory: REAL • buffer_memory_enum: STRING • buffer_in_use: INTEGER • buffer_in_use_enum: STRING • buffer_spooled: INTEGER • buffer_spooled_enum: STRING • flat_buffer_memory: REAL • flat_buffer_memory_enum: STRING • flat_buffers_in_use: INTEGER • flat_buffers_in_use_enum: STRING • private_buffer_memory: REAL • private_buffer_memory_enum: STRING • private_buffers_in_use: INTEGER • private_buffers_in_use_enum: STRING • total_rows_read: REAL • total_rows_read_enum: STRING • total_rows_written: REAL • total_rows_written_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_ Data_Mining	MS_SQL_Analysis_Services_Data_Mining attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_enum: STRING • cases_per_sec: REAL • cases_per_sec_enum: STRING • current_models_processing: REAL • current_models_processing_enum: STRING • concurrent_data_mining_queries: REAL • concurrent_data_mining_queries_enum: STRING • predictions_per_sec: REAL • predictions_per_sec_enum: STRING • queries_per_sec: REAL • queries_per_sec_enum: STRING • rows_per_sec: REAL • rows_per_sec_enum: STRING • total_predictions: REAL • total_predictions_enum: STRING • total_queries: REAL • total_queries_enum: STRING • total_rows: REAL • total_rows_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Memory_Statistics	<p>MS_SQL_Analysis_Services_Memory_Statistics attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_enum: STRING • average_cache: REAL • average_cache_enum: STRING • aggregation_map_files: REAL • aggregation_map_files_enum: STRING • cleaner_balance_per_sec: REAL • cleaner_balance_per_sec_enum: STRING • cleaner_current_price: REAL • cleaner_current_price_enum: STRING • cleaner_memory: REAL • cleaner_memory_enum: STRING • cleaner_memory_nonshrinkable: REAL • cleaner_memory_nonshrinkable_enum: STRING • cleaner_memory_shrinkable: REAL • cleaner_memory_shrinkable_enum: STRING • cleaner_memory_shrunk_kb_sec: REAL • cleaner_memory_shrunk_kb_sec_enum: STRING • dimension_index_files: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
	<p>(Continued)</p> <ul style="list-style-type: none"> • dimension_index_files_enum: STRING • dimension_property_files: REAL • dimension_property_files_enum: STRING • dimension_string_files: REAL • dimension_string_files_enum: STRING • fact_aggregation_files: REAL • fact_aggregation_files_enum: STRING • fact_data_files: REAL • fact_data_files_enum: STRING • fact_string_files: REAL • fact_string_files_enum: STRING • filestore_clock_pages_examined_per_sec: REAL • filestore_clock_pages_examined_per_sec_enum: STRING • filestore_clock_pages_haveref_per_sec: REAL • filestore_clock_pages_haveref_per_sec_enum: STRING • filestore_clock_pages_valid_per_sec: REAL • filestore_clock_pages_valid_per_sec_enum: STRING • total_filestore_io_errors: REAL • total_filestore_io_errors_enum: STRING • total_filestore_io_errors_per_sec: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Memory_Statistics (Continued)	<p>(Continued)</p> <ul style="list-style-type: none"> • total_filestore_io_errors_per_sec_enum: STRING • filestore: REAL • filestore_enum: STRING • filestore_kb_reads_per_sec: REAL • filestore_kb_reads_per_sec_enum: STRING • filestore_kb_writes_per_sec: REAL • filestore_kb_writes_per_sec_enum: STRING • filestore_memory_pinned: REAL • filestore_memory_pinned_enum: STRING • filestore_page_faults_per_sec: REAL • filestore_page_faults_per_sec_enum: STRING • filestore_reads_per_sec: REAL • filestore_reads_per_sec_enum: STRING • filestore_writes_per_sec: REAL • filestore_writes_per_sec_enum: STRING • in_memory_aggregation_map_file: REAL • in_memory_aggregation_map_file_enum: STRING • in_memory_aggregation_map_file_kb_per_sec: REAL • in_memory_aggregation_map_file_kb_per_sec_enum: STRING • in_memory_dimension_index_file: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
	<p>(Continued)</p> <ul style="list-style-type: none"> • in_memory_dimension_index_file_enum: STRING • in_memory_dimension_index_file_kb_per_sec: REAL • in_memory_dimension_index_file_kb_per_sec_enum: STRING • in_memory_property_file: REAL • in_memory_property_file_enum: STRING • in_memory_property_file_kb_per_sec: REAL • in_memory_property_file_kb_per_sec_enum: STRING • in_memory_string_file: REAL • in_memory_string_file_enum: STRING • in_memory_string_file_kb_per_sec: REAL • in_memory_string_file_kb_per_sec_enum: STRING • in_memory_fact_aggregation_file: REAL • in_memory_fact_aggregation_file_enum: STRING • in_memory_fact_aggregation_file_kb_per_sec: REAL • in_memory_fact_aggregation_file_kb_per_sec_enum: STRING • in_memory_fact_data_file: REAL • in_memory_fact_data_file_enum: STRING • in_memory_fact_data_file_kb_per_sec: REAL • in_memory_fact_data_file_kb_per_sec_enum: STRING • in_memory_fact_string_file: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
	<p>(Continued)</p> <ul style="list-style-type: none"> • in_memory_fact_string_file_enum: STRING • in_memory_fact_string_file_per_sec: REAL • in_memory_fact_string_file_per_sec_enum: STRING • in_memory_map_file: REAL • in_memory_map_file_enum: STRING • in_memory_map_file_per_sec: REAL • in_memory_map_file_per_sec_enum: STRING • in_memory_other_file: REAL • in_memory_other_file_enum: STRING • in_memory_other_file_per_sec: REAL • in_memory_other_file_per_sec_enum: STRING • map_files: REAL • map_files_enum: STRING • memory_limit_low: REAL • memory_limit_low_enum: STRING • memory_limit_high: REAL • memory_limit_high_enum: STRING • memory_usage: REAL • memory_usage_enum: STRING • other_files: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Memory_Statistics (Continued)	<p>(Continued)</p> <ul style="list-style-type: none"> • other_files_enum: STRING • page_pool_1_alloc: REAL • page_pool_1_alloc_enum: STRING • page_pool_1_lookaside: REAL • page_pool_1_lookaside_enum: STRING • page_pool_64_alloc: REAL • page_pool_64_alloc_enum: STRING • page_pool_64_lookaside: REAL • page_pool_64_lookaside_enum: STRING • page_pool_8_alloc: REAL • page_pool_8_alloc_enum: STRING • page_pool_8_lookaside: REAL • page_pool_8_lookaside_enum: STRING • potential_in-memory_aggregation_map_file: REAL • potential_in-memory_aggregation_map_file_enum: STRING • potential_in-memory_dimension_index_file: REAL • potential_in-memory_dimension_index_file_enum: STRING • potential_in-memory_dimension_property_file: REAL • potential_in-memory_dimension_property_file_enum: STRING • potential_in-memory_dimension_string_file: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
	<p>(Continued)</p> <ul style="list-style-type: none"> • potential_in-memory_dimension_string_file_enum: STRING • potential_in-memory_fact_aggregation_file: REAL • potential_in-memory_fact_aggregation_file_enum: STRING • potential_in-memory_fact_data_file: REAL • potential_in-memory_fact_data_file_enum: STRING • potential_in-memory_fact_string_file: REAL • potential_in-memory_fact_string_file_enum: STRING • potential_in-memory_map_file: REAL • potential_in-memory_map_file_enum: STRING • potential_in-memory_other_file: REAL • potential_in-memory_other_file_enum: STRING • quota_blocked: REAL • quota_blocked_enum: STRING • quota: REAL • quota_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Storage_Engine_Query	<p>MS_SQL_Analysis_Services_Storage_Engine_Query attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_enum: STRING • aggregation_hits_per_sec: REAL • aggregation_hits_per_sec_enum: STRING • aggregation_lookups_per_sec: REAL • aggregation_lookups_per_sec_enum: STRING • average_time_per_query: REAL • average_time_per_query_enum: STRING • bytes_sent_per_sec: REAL • bytes_sent_per_sec_enum: STRING • calculation_cache_hits_per_sec: REAL • calculation_cache_hits_per_sec_enum: STRING • calculation_cache_lookups_per_sec: REAL • calculation_cache_lookups_per_sec_enum: STRING • current_dimension_queries: REAL • current_dimension_queries_enum: STRING • current_measure_group_queries: REAL • current_measure_group_queries_enum: STRING • current_pyramid_operations: REAL • current_pyramid_operations_enum: STRING • data_bytes_per_sec: REAL • data_bytes_per_sec_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Storage_Engine_Query (Continued)	<p>(Continued)</p> <ul style="list-style-type: none"> • data_reads_per_sec: REAL • data_reads_per_sec_enum: STRING • dimension_cache_hits_per_sec: REAL • dimension_cache_hits_per_sec_enum: STRING • dimension_cache_lookups_per_sec: REAL • dimension_cache_lookups_per_sec_enum: STRING • dimension_queries_per_sec: REAL • dimension_queries_per_sec_enum: STRING • filter_rows_excluded_per_sec: REAL • filter_rows_excluded_per_sec_enum: STRING • filter_rows_included_per_sec: REAL • filter_rows_included_per_sec_enum: STRING • filtered_rows_per_sec: REAL • filtered_rows_per_sec_enum: STRING • flat_cache_hits_per_sec: REAL • flat_cache_hits_per_sec_enum: STRING • flat_cache_lookups_per_sec: REAL • flat_cache_lookups_per_sec_enum: STRING • index_bytes_per_sec: REAL • index_bytes_per_sec_enum: STRING • index_reads_per_sec: REAL • index_reads_per_sec_enum: STRING • map_bytes_per_sec: REAL • map_bytes_per_sec_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Storage_Engine_Query (Continued)	<p>(Continued)</p> <ul style="list-style-type: none"> • map_reads_per_sec: REAL • map_reads_per_sec_enum: STRING • measure_group_cache_hits_per_sec: REAL • measure_group_cache_hits_per_sec_enum: STRING • measure_group_lookups_per_sec: REAL • measure_group_lookups_per_sec_enum: STRING • measure_group_queries_per_sec: REAL • measure_group_queries_per_sec_enum: STRING • network_round_trips_per_sec: REAL • network_round_trips_per_sec_enum: STRING • persisted_cache_hits_per_sec: REAL • persisted_cache_hits_per_sec_enum: STRING • persisted_cache_lookups_per_sec: REAL • persisted_cache_lookups_per_sec_enum: STRING • pyramid_operations_per_sec: REAL • pyramid_operations_per_sec_enum: STRING • queries_answered_per_sec: REAL • queries_answered_per_sec_enum: STRING • queries_from_cache_direct_per_sec: REAL • queries_from_cache_direct_per_sec_enum: STRING • queries_from_cache_filtered_per_sec: REAL • queries_from_cache_filtered_per_sec_enum: STRING • queries_from_files_per_sec: REAL • queries_from_files_per_sec_enum: STRING • rows_read_per_sec: REAL • rows_read_per_sec_enum: STRING • rows_sent_per_sec: REAL • rows_sent_per_sec_enum: STRING • total_bytes_sent: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_Storage_Engine_Query (Continued)	<p>(Continued)</p> <ul style="list-style-type: none"> • total_bytes_sent_enum: STRING • total_dimension_queries: REAL • total_dimension_queries_enum: STRING • total_measure_group_queries: REAL • total_measure_group_queries_enum: STRING • total_network_round_trips: REAL • total_network_round_trips_enum: STRING • total_pyramid_operations: REAL • total_pyramid_operations_enum: STRING • total_queries_answered: REAL • total_queries_answered_enum: STRING • total_queries_from_cache_direct: REAL • total_queries_from_cache_direct_enum: STRING • total_queries_from_cache_filtered: REAL • total_queries_from_cache_filtered_enum: STRING • total_queries_from_file: REAL • total_queries_from_file_enum: STRING • total_rows_sent: REAL • total_rows_sent_enum: STRING
ITM_MS_SQL_Analysis_Services_Rows_Processing	<p>MS_SQL_Analysis_Services_Rows_Processing attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • server_enum: STRING • rows_read_per_sec: REAL • rows_read_per_sec_enum: STRING • rows_written_per_sec: REAL • rows_written_per_sec_enum: STRING • total_rows_converted: REAL • total_rows_converted_enum: STRING • total_rows_read: REAL • total_rows_read_enum: STRING • total_rows_written: REAL • total_rows_written_enum: STRING • rows_converted_per_sec: REAL • rows_converted_per_sec_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_MDX_Processing	<p>MS_SQL_Analysis_Services_MDX_Processing attribute group</p> <ul style="list-style-type: none"> • hub_timestamp: STRING • server: STRING • server_enum: STRING • current_cached_evaluation_nodes: REAL • current_cached_evaluation_nodes_enum: STRING • current_evaluation_nodes: REAL • current_evaluation_nodes_enum: STRING • total_bulk_mode_evaluation_nodes: REAL • total_bulk_mode_evaluation_nodes_enum: STRING • total_cached_bulk_mode_evaluation_nodes: REAL • total_cached_bulk_mode_evaluation_nodes_enum: STRING • total_cached_other_evaluation_nodes: REAL • total_cached_other_evaluation_nodes_enum: STRING • total_cached_storage_engine_evaluation_nodes: REAL • total_cached_storage_engine_evaluation_nodes_enum: STRING • total_calculations_covered: REAL • total_calculations_covered_enum: STRING • total_cell_by_cell_evaluation_nodes: REAL • total_cell_by_cell_evaluation_nodes_enum: STRING • total_cell_by_cell_hits_in_cache_of_evaluation_nodes: REAL • total_cell_by_cell_hits_in_cache_of_evaluation_nodes_enum: STRING • total_cell_by_cell_misses_in_cache_of_evaluation_nodes: REAL

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_MDX_Processing (Continued)	<p>(Continued)</p> <ul style="list-style-type: none"> • total_cell_by_cell_misses_in_cache_of_evaluation_nodes_enum: STRING • total_evaluation_nodes_covered_single_cell: REAL • total_evaluation_nodes_covered_single_cell_enum: STRING • total_evaluation_nodes_calculations_at_same_granularity: REAL • total_evaluation_nodes_calculations_at_same_granularity_enum: STRING • total_evictions_of_evaluation_nodes: REAL • total_evictions_of_evaluation_nodes_enum: STRING • total_hash_index_hits_in_cache_of_evaluation_nodes: REAL • total_hash_index_hits_in_cache_of_evaluation_nodes_enum: STRING • total_storage_engine_evaluation_nodes: REAL • total_storage_engine_evaluation_nodes_enum: STRING • total_subcube_hits_in_cache_of_evaluation_nodes: REAL • total_subcube_hits_in_cache_of_evaluation_nodes_enum: STRING • total_subcube_misses_in_cache_of_evaluation_nodes: REAL • total_subcube_misses_in_cache_of_evaluation_nodes_enum: STRING • total_autoexist: REAL • total_autoexist_enum: STRING • total_calculation_cache_registered: REAL • total_calculation_cache_registered_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Analysis_Services_ MDX_Processing (Continued)	(Continued) <ul style="list-style-type: none"> • total_cells_calculated: REAL • total_cells_calculated_enum: STRING • total_existing_operators: REAL • total_existing_operators_enum: STRING • total_flat_cache_inserts: REAL • total_flat_cache_inserts_enum: STRING • total_nonempty_algorithms: REAL • total_nonempty_algorithms_enum: STRING • total_nonempty_algorithms_for_calculated_members: REAL • total_nonempty_algorithms_for_calculated_members_enum: STRING • total_nonempty_unoptimized_algorithms: REAL • total_nonempty_unoptimized_algorithms_enum: STRING • total_recomputes: REAL • total_recomputes_enum: STRING • total_sonar_subcubes: REAL • total_sonar_subcubes_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Running_Queries_Details	<p>MS_SQL_Running_Queries_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • cpu_time_sec: REAL • cpu_time_sec_enum: STRING • total_elapsed_time_sec: REAL • total_elapsed_time_sec_enum: STRING • query_reads: REAL • query_reads_enum: STRING • query_writes: REAL • query_writes_enum: STRING • query_start_time: STRING • query_start_time_enum: STRING • sql_handle: STRING • sql_handle_enum: STRING • command_type: STRING • command_type_enum: STRING • row_count: REAL • row_count_enum: STRING • query_status: INTEGER • query_status_enum: STRING • running_query_text: STRING • running_query_text_enum: STRING • session_id: INTEGER • session_id_enum: STRING • blocking_session_id: INTEGER • blocking_session_id_enum: STRING • client_host_name: STRING • client_host_name_enum: STRING • login_name: STRING • login_name_enum: STRING • program_name: STRING • program_name_enum: STRING • database_name_u: STRING • database_name_u_enum: STRING • percent_work_complete: REAL • percent_work_complete_enum: STRING • granted_query_memory: REAL • granted_query_memory_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Blocked_Queries_Details	<p>MS_SQL_Blocked_Queries_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • blocked_session_id: INTEGER • blocked_session_id_enum: STRING • total_wait_time_sec: REAL • total_wait_time_sec_enum: STRING • blocked_query_text: STRING • blocked_query_text_enum: STRING • blocking_query_text: STRING • blocking_query_text_enum: STRING • blocked_parent_query_text: STRING • blocked_parent_query_text_enum: STRING • blocking_reason: STRING • blocking_reason_enum: STRING • blocking_resource: STRING • blocking_resource_enum: STRING • blocking_session_id: INTEGER • blocking_session_id_enum: STRING • blocked_database_name_u: STRING • blocked_database_name_u_enum: STRING • blocking_user_name: STRING • blocking_user_name_enum: STRING • blocked_user_name_u: STRING • blocked_user_name_u_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Individual_Queries_Details	<p>MS_SQL_Individual_Queries_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • session_id: INTEGER • session_id_enum: STRING • query_hash: REAL • query_hash_enum: STRING • query_execution_count: REAL • query_execution_count_enum: STRING • query_text: STRING • query_text_enum: STRING • total_percent_cpu_used: REAL • total_percent_cpu_used_enum: STRING • total_percent_waiting: REAL • total_percent_waiting_enum: STRING • last_execution_time: STRING • last_execution_time_enum: STRING • total_cpu_time_sec: REAL • total_cpu_time_sec_enum: STRING • avg_cpu_time_sec: REAL • avg_cpu_time_sec_enum: STRING • avg_logical_reads: REAL • avg_logical_reads_enum: STRING • avg_logical_writes: REAL • avg_logical_writes_enum: STRING • avg_physical_reads: REAL • avg_physical_reads_enum: STRING • max_physical_reads: REAL • max_physical_reads_enum: STRING • max_logical_reads: REAL • max_logical_reads_enum: STRING • max_logical_writes: REAL • max_logical_writes_enum: STRING • max_cpu_time_sec: REAL • max_cpu_time_sec_enum: STRING • max_waiting_time_sec: REAL • max_waiting_time_sec_enum: STRING • max_clr_time_sec: REAL • max_clr_time_sec_enum: STRING • database_name_u: STRING • database_name_u_enum: STRING • sql_handle: STRING • sql_handle_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_FileTable_Detail	<p>MS_SQL_FileTable_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_id: INTEGER • database_id_enum: STRING • table_name: STRING • table_name_enum: STRING • table_id: INTEGER • table_id_enum: STRING • filetable_directory: STRING • filetable_directory_enum: STRING • filetable_enabled: INTEGER • filetable_enabled_enum: STRING • total_non_transactional_handles_opened: INTEGER • total_non_transactional_handles_opened_enum: STRING
ITM_MS_SQL_Additional_Database_Detail	<p>MS_SQL_Additional_Database_Detail attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • database_id: INTEGER • database_id_enum: STRING • database_name_u: STRING • database_name_u_enum: STRING • filetable_directory_name: STRING • filetable_directory_name_enum: STRING • non_transactional_access_level: INTEGER • non_transactional_access_level_enum: STRING • containment: INTEGER • containment_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Replicas_Details	<p>MS_SQL_Availability_Replicas_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • replica_id: STRING • replica_id_enum: STRING • group_id: STRING • group_id_enum: STRING • group_name: STRING • group_name_enum: STRING • replica_server_name: STRING • replica_server_name_enum: STRING • owner_sid: STRING • owner_sid_enum: STRING • endpoint_url: STRING • endpoint_url_enum: STRING • create_date: STRING • create_date_enum: STRING • modify_date: STRING • modify_date_enum: STRING • availability_mode: INTEGER • availability_mode_enum: STRING • failover_mode: INTEGER • failover_mode_enum: STRING • session_timeout: INTEGER • session_timeout_enum: STRING • primary_connection_type: INTEGER • primary_connection_type_enum: STRING • secondary_connection_type: INTEGER • secondary_connection_type_enum: STRING • backup_priority: INTEGER • backup_priority_enum: STRING • role: INTEGER • role_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Replicas_Status	<p>MS_SQL_Availability_Replicas_Status attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • replica_id: STRING • replica_id_enum: STRING • group_id: STRING • group_id_enum: STRING • group_name: STRING • group_name_enum: STRING • local: INTEGER • local_enum: STRING • role: INTEGER • role_enum: STRING • operational_state: INTEGER • operational_state_enum: STRING • connection_state: INTEGER • connection_state_enum: STRING • recovery_health: INTEGER • recovery_health_enum: STRING • synchronization_health: INTEGER • synchronization_health_enum: STRING • last_connect_error_number: INTEGER • last_connect_error_number_enum: STRING • total_databases: INTEGER • total_databases_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Replicas_Status_Summary	<p>MS_SQL_Availability_Replicas_Status_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • total_primary_replicas: INTEGER • total_primary_replicas_enum: STRING • total_secondary_replicas: INTEGER • total_secondary_replicas_enum: STRING • min_secondary_replicas: INTEGER • min_secondary_replicas_enum: STRING • max_secondary_replicas: INTEGER • max_secondary_replicas_enum: STRING • total_local_replicas: INTEGER • total_local_replicas_enum: STRING • total_connected_replicas: INTEGER • total_connected_replicas_enum: STRING • total_disconnected_replicas: INTEGER • total_disconnected_replicas_enum: STRING • min_disconnected_replicas: INTEGER • min_disconnected_replicas_enum: STRING • max_disconnected_replicas: INTEGER • max_disconnected_replicas_enum: STRING • total_local_failed_replicas: INTEGER • total_local_failed_replicas_enum: STRING • min_local_failed_replicas: INTEGER • min_local_failed_replicas_enum: STRING • max_local_failed_replicas: INTEGER • max_local_failed_replicas_enum: STRING • total_non_healthy_replicas: INTEGER • total_non_healthy_replicas_enum: STRING • min_non_healthy_replicas: INTEGER • min_non_healthy_replicas_enum: STRING • max_non_healthy_replicas: INTEGER • max_non_healthy_replicas_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Groups_Details	<p>MS_SQL_Availability_Groups_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • group_id: STRING • group_id_enum: STRING • group_name_u: STRING • group_name_u_enum: STRING • primary_replica_id: STRING • primary_replica_id_enum: STRING • primary_replica_u: STRING • primary_replica_u_enum: STRING • groups_role: INTEGER • groups_role_enum: STRING • database_count: INTEGER • database_count_enum: STRING • failure_condition: INTEGER • failure_condition_enum: STRING • wait_timeout: INTEGER • wait_timeout_enum: STRING • backup_preference: INTEGER • backup_preference_enum: STRING • local_replica_recovery_health: INTEGER • local_replica_recovery_health_enum: STRING • synchronization_health: INTEGER • synchronization_health_enum: STRING • secondary_replica_count: INTEGER • secondary_replica_count_enum: STRING • last_groups_role: INTEGER • last_groups_role_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_ Availability_Database_Details	MS_SQL_Availability_Database_ Details attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • database_id: INTEGER • database_id_enum: STRING • database_name_u: STRING • database_name_u_enum: STRING • group_id: STRING • group_id_enum: STRING • group_name_u: STRING • group_name_u_enum: STRING • replica_id: STRING • replica_id_enum: STRING • group_database_id: STRING • group_database_id_enum: STRING • role: INTEGER • role_enum: STRING • local: INTEGER • local_enum: STRING • synchronization_state: INTEGER • synchronization_state_enum: STRING • synchronized_commit: INTEGER • synchronized_commit_enum: STRING • synchronization_health: INTEGER • synchronization_health_enum: STRING • database_state: INTEGER • database_state_enum: STRING • suspend_reason: INTEGER • suspend_reason_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Additional_Availability_ Database_Details	MS_SQL_Additional_Availability_ Database_Details attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • database_name_u: STRING • database_name_u_enum: STRING • group_id: STRING • group_id_enum: STRING • group_database_id: STRING • group_database_id_enum: STRING • log_send_queue_size: REAL • log_send_queue_size_enum: STRING • log_bytes_sent_rate: REAL • log_bytes_sent_rate_enum: STRING • redo_queue_size: REAL • redo_queue_size_enum: STRING • redo_rate: REAL • redo_rate_enum: STRING • filestream_send_rate: REAL • filestream_send_rate_enum: STRING • role: INTEGER • role_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_ Database_Summary	<p>MS_SQL_Availability_Database_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • total_databases: INTEGER • total_databases_enum: STRING • total_primary_databases: INTEGER • total_primary_databases_enum: STRING • total_secondary_databases: INTEGER • total_secondary_databases_enum: STRING • total_unsynchronized_databases: INTEGER • total_unsynchronized_databases_enum: STRING • min_unsynchronized_databases: INTEGER • min_unsynchronized_databases_enum: STRING • max_unsynchronized_databases: INTEGER • max_unsynchronized_databases_enum: STRING • total_unsynchronized_commit_db_replicas: INTEGER • total_unsynchronized_commit_db_replicas_enum: STRING • min_unsynchronized_commit_db_replicas: INTEGER • min_unsynchronized_commit_db_replicas_enum: STRING • max_unsynchronized_commit_db_replicas: INTEGER • max_unsynchronized_commit_db_replicas_enum: STRING • total_non-online_databases: INTEGER • total_non-online_databases_enum: STRING • min_non-online_databases: INTEGER • min_non-online_databases_enum: STRING • max_non-online_databases: INTEGER • max_non-online_databases_enum: STRING • total_non_healthy_databases: INTEGER • total_non_healthy_databases_enum: STRING • min_non_healthy_databases: INTEGER • min_non_healthy_databases_enum: STRING • max_non_healthy_databases: INTEGER • max_non_healthy_databases_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Groups_Summary	<p>MS_SQL_Availability_Groups_Summary attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • host_name: STRING • total_groups: INTEGER • total_groups_enum: STRING • total_primary_role_group: INTEGER • total_primary_role_group_enum: STRING • total_secondary_role_group: INTEGER • total_secondary_role_group_enum: STRING • total_non_healthy_groups: INTEGER • total_non_healthy_groups_enum: STRING • min_non_healthy_groups: INTEGER • min_non_healthy_groups_enum: STRING • max_non_healthy_groups: INTEGER • max_non_healthy_groups_enum: STRING • total_non_online_local_replica_groups: INTEGER • total_non_online_local_replica_groups_enum: STRING • min_non_online_local_replica_groups: INTEGER • min_non_online_local_replica_groups_enum: STRING • max_non_online_local_replica_groups: INTEGER • max_non_online_local_replica_groups_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Replicas_Statistics	<p>MS_SQL_Availability_Replicas_Statistics attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • instance_name: STRING • instance_name_enum: STRING • bytes_received_per_sec: REAL • bytes_received_per_sec_enum: STRING • bytes_sent_per_sec: REAL • bytes_sent_per_sec_enum: STRING • bytes_sent_to_transport_per_sec: REAL • bytes_sent_to_transport_per_sec_enum: STRING • flow_control_time: REAL • flow_control_time_enum: STRING • flow_control_per_sec: REAL • flow_control_per_sec_enum: STRING • received_messages_per_sec: REAL • received_messages_per_sec_enum: STRING • resent_messages_per_sec: REAL • resent_messages_per_sec_enum: STRING • sent_messages_per_sec: REAL • sent_messages_per_sec_enum: STRING • messages_sent_to_transport_per_sec: REAL • messages_sent_to_transport_per_sec_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_ Database_Statistics	MS_SQL_Availability_Database_ Statistics attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • database_name_u: STRING • database_name_u_enum: STRING • file_bytes_received_per_sec: REAL • file_bytes_received_per_ sec_enum: STRING • log_bytes_received_per_sec: REAL • log_bytes_received_per_ sec_enum: STRING • undo_log_remaining: REAL • undo_log_remaining_enum: STRING • log_send_queue: REAL • log_send_queue_enum: STRING • mirrored_write_transaction_per_sec: REAL • mirrored_write_transaction_per_ sec_enum: STRING • recovery_queue: REAL • recovery_queue_enum: STRING • redo_blocked_per_sec: REAL • redo_blocked_per_sec_enum: STRING • redo_bytes_remaining: REAL • redo_bytes_remaining_enum: STRING • redone_bytes_per_sec: REAL • redone_bytes_per_sec_enum: STRING • total_log_requiring_undo: REAL • total_log_requiring_undo_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Groups_Details_In_Cluster	<p>MS_SQL_Availability_Groups_Details_In_Cluster attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • group_id: STRING • group_id_enum: STRING • group_name_u: STRING • group_name_u_enum: STRING • cluster_resource_id: STRING • cluster_resource_id_enum: STRING • resource_group_id: STRING • resource_group_id_enum: STRING • failure_condition: INTEGER • failure_condition_enum: STRING • wait_timeout: INTEGER • wait_timeout_enum: STRING • backup_preference: INTEGER • backup_preference_enum: STRING
ITM_MS_SQL_Availability_Replicas_Details_In_Cluster	<p>MS_SQL_Availability_Replicas_Details_In_Cluster attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • replica_id: STRING • replica_id_enum: STRING • group_id: STRING • group_id_enum: STRING • group_name_u: STRING • group_name_u_enum: STRING • cluster_node: STRING • cluster_node_enum: STRING • replica_server_name: STRING • replica_server_name_enum: STRING • join_state: INTEGER • join_state_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Database_Details_In_Cluster	MS_SQL_Availability_Database_Details_In_Cluster attribute group <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • database_name_u: STRING • database_name_u_enum: STRING • group_id: STRING • group_id_enum: STRING • group_name_u: STRING • group_name_u_enum: STRING • replica_id: STRING • replica_id_enum: STRING • group_database_id: STRING • group_database_id_enum: STRING • failover_ready: INTEGER • failover_ready_enum: STRING • secondary_suspension_pending: INTEGER • secondary_suspension_pending_enum: STRING • join_state: INTEGER • join_state_enum: STRING

Table 40. Overview of attribute groups to event classes and slots (continued)

IBM Tivoli Enterprise Console event class	Event slots
ITM_MS_SQL_Availability_Group_Listener_Details	<p>MS_SQL_Availability_Group_Listener_Details attribute group</p> <ul style="list-style-type: none"> • originnode: STRING • hub_timestamp: STRING • server: STRING • listener_id: STRING • listener_id_enum: STRING • listener_name_u: STRING • listener_name_u_enum: STRING • group_id: STRING • group_id_enum: STRING • group_name_u: STRING • group_name_u_enum: STRING • listener_ip_address: STRING • listener_ip_address_enum: STRING • listener_subnet_ip: STRING • listener_subnet_ip_enum: STRING • port: INTEGER • port_enum: STRING • is_dhcp: INTEGER • is_dhcp_enum: STRING • state: INTEGER • state_enum: STRING • is_conformant: INTEGER • is_conformant_enum: STRING

Appendix C. Discovery Library Adapter for the Microsoft SQL Server agent

The Tivoli Management Services Discovery Library Adapter (DLA) discovers resources and relationships, and creates a Discovery Library Book file for the agent.

About the DLA

The Book file follows the Discovery Library IdML schema and is used to populate the Configuration Management Database (CMDB) and Tivoli Business Service Manager products. The Tivoli Management Services DLA discovers Microsoft SQL Server resources. For all VMware systems that are active and online at the Tivoli Enterprise Portal Server, information is included in the discovery book for those resources. The Tivoli Management Services DLA discovers active resources. It is run on demand and can be run periodically to discover resources that were not active during previous discoveries.

The DLA discovers Microsoft SQL Server components.

More information about DLAs

The following sources contain additional information about using the DLA program with all monitoring agents:

- The *IBM Tivoli Monitoring Administrator's Guide* contains information about using the Tivoli Management Services Discovery Library Adapter.
- For information about using a DLA with Tivoli Application Dependency Discovery Manager (TADDM), see the TADDM Information Center (http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.taddm.doc_7.2/welcome_page/welcome.html).

DLA data model class types represented in CDM

The source application data objects map to classes in the Common Data Model (CDM) for the Microsoft SQL Server agent.

The following information is provided for each class:

CDM class name

Class name for which the agent is providing information

Relationships

CDM relationships (hierarchical) between currently identified model objects

CDM attributes, agent attributes, descriptions, and examples

CDM and agent attributes that are required to create an instance of a resource, descriptions of the attributes, and examples of the attributes

DLA data model classes for Microsoft SQL Server agent

Each agent that uses the Discovery Library Adapter has DLA data model classes defined for the agent.

The Microsoft SQL Server agent has the following Discovery Library Adapter data model classes:

- SqlServer
- SqlServerDatabase
- SqlServerTable

- SqlServerProcess
- TMSAgent

SqlServer class

The ActiveDirectory class represents an instance of the domain controller.

CDM class name

app.db.mssql.SqlServer

Relationships

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.ComputerSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-ComputerSystem

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-OperatingSystem

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerDatabase
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-SqlServerDatabase

contains

- Source: app.db.mssql.SqlServerDatabase
- Target: app.db.mssql.SqlServerTable
- Example: SQLServer:SQLInstance:MSS-1-SqlServerDatabase contains SQLServer:SQLInstance:MSS-1-0-spt_monitor-SqlServerTable

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerProcess
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-7412-SqlServerProcess

installedOn

- Source: app.TMSAgent
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-TMSAgent is installedOn 9.65.44.55-OperatingSystem
- Source: app.TMSAgent
- Target: app.db.mssql.SqlServer
- Example: SQLServer:SQLInstance:MSS-TMSAgent monitors SQLServer:SQLInstance:MSS-SqlServer

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: Id (IdML)
Agent attribute: INODESTS.HOSTADDR
Description: The host address of the workstation where the Microsoft SQL Server instance is running.
Example: INSTANCE1_2005:PS3213W2K3:MSS-SqlServer

- CDM attribute: ManagedSystemName
Agent attribute: INODESTS.NODE
Description: The name of the IBM Tivoli Monitoring component that provides data for the management of the Microsoft SQL Server instance.
Example: INSTANCE1_2005:PS3213W2K3:MSS
- CDM attribute: KeyName
Agent attribute: AppServer
Description: A user defined string that is used to identify the Microsoft SQL Server instance.
Example: AppServer
- CDM attribute: ProductVersion
Agent attribute: KOQSRVRE.SRVVERS
Description: The version of the Microsoft SQL Server.
Example: 9.00.13
- CDM attribute: Name
Agent attribute: KOQSRVRE.SERVERID
Description: The name of the Microsoft SQL Server.
Example: INSTANCE1_2005
- CDM attribute: Status
Agent attribute: KOQSRVRE.SRVSTATUS
Description: The status of the Microsoft SQL Server (Running/Stopped).
Example: Running
- CDM attribute: Label
Agent attribute: KOQSRVRE.SERVERID
Description: The label of the Microsoft SQL Server.
Example: INSTANCE1_2005

SqlServerDatabase class

The ServiceAccessPoint class represents the instance of the services that are used by the Microsoft Active Directory.

CDM class name

app.db.mssql.SqlServerDatabase

Superior CDM class

SqlServer

Relationships

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.ComputerSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-ComputerSystem

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-OperatingSystem

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerDatabase
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-SqlServerDatabase

contains

- Source: app.db.mssql.SqlServerDatabase
- Target: app.db.mssql.SqlServerTable
- Example: SQLServer:SQLInstance:MSS-1-SqlServerDatabase contains SQLServer:SQLInstance:MSS-1-0-spt_monitor-SqlServerTable

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerProcess
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-7412-SqlServerProcess

installedOn

- Source: app.TMSAgent
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-TMSAgent is installedOn 9.65.44.55-OperatingSystem
- Source: app.TMSAgent
- Target: app.db.mssql.SqlServer
- Example: SQLServer:SQLInstance:MSS-TMSAgent monitors SQLServer:SQLInstance:MSS-SqlServer

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: Id (IdML)
Agent attribute: INODESTS.HOSTADDR
Description: The host address of the workstation where the Microsoft SQL Server instance is running.
Example: INSTANCE1_2005:PS3213W2K3:MSS-2-SqlServerDatabase
- CDM attribute: ManagedSystemName
Agent attribute: INODESTS.NODE
Description: The name of the IBM Tivoli Monitoring component that provides data for the management of the Microsoft SQL Server instance.
Example: INSTANCE1_2005:PS3213W2K3:MSS
- CDM attribute: Name
Agent attribute: KOQDBD.DBNAME
Description: The name of the Microsoft SQL Server database.
Example: tempdb
- CDM attribute: DbId
Agent attribute: KOQDB.DBID
Description: The ID of the Microsoft SQL Server database.
Example: 2
- CDM attribute: Label
Agent attribute: KOQDBD.DBNAME
Description: The label of the Microsoft SQL Server database.
Example: INSTANCE1_2005:PS3213W2K3:MSS

SqlServerTable class

The BindAddress class represents an IP protocol endpoint on which a service is provided.

CDM class name

app.db.mssql.SqlServerTable

Superior CDM class

SqlServerDatabase

Relationships

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.ComputerSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-ComputerSystem

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-OperatingSystem

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerDatabase
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-SqlServerDatabase

contains

- Source: app.db.mssql.SqlServerDatabase
- Target: app.db.mssql.SqlServerTable
- Example: SQLServer:SQLInstance:MSS-1-SqlServerDatabase contains SQLServer:SQLInstance:MSS-1-0-spt_monitor-SqlServerTable

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerProcess
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-7412-SqlServerProcess

installedOn

- Source: app.TMSAgent
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-TMSAgent is installedOn 9.65.44.55-OperatingSystem
- Source: app.TMSAgent
- Target: app.db.mssql.SqlServer
- Example: SQLServer:SQLInstance:MSS-TMSAgent monitors SQLServer:SQLInstance:MSS-SqlServer

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: Id (IdML)
Agent attribute: INODESTS.HOSTADDR
Description: The host address of the workstation where the Microsoft SQL Server instance is running.
Example: INSTANCE1_2005:PS3213W2K3:MSS-0-spt_monitor-SqlServerTable
- CDM attribute: ManagedSystemName
Agent attribute: INODESTS.NODE
Description: The name of the IBM Tivoli Monitoring component that provides data for the management of the Microsoft SQL Server instance.
Example: INSTANCE1_2005:PS3213W2K3:MSS
- CDM attribute: Name
Agent attribute: KOQTBLD.UTBLNAME
Description: The name of the table in the Microsoft SQL Server database.
Example: spt_monitor
- CDM attribute: Label

Agent attribute: KOQTBLD.UTBLNAME
Description: The label of the Microsoft SQL Server table.
Example: spt_monitor

SqlServerProcess class

The IpAddress class represents the instance of an IP address.

CDM class name

app.db.mssql.SqlServerProcess

Superior CDM class

SqlServer

Relationships

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.ComputerSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-ComputerSystem

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-OperatingSystem

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerDatabase
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-SqlServerDatabase

contains

- Source: app.db.mssql.SqlServerDatabase
- Target: app.db.mssql.SqlServerTable
- Example: SQLServer:SQLInstance:MSS-1-SqlServerDatabase contains SQLServer:SQLInstance:MSS-1-0-spt_monitor-SqlServerTable

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerProcess
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-7412-SqlServerProcess

installedOn

- Source: app.TMSAgent
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-TMSAgent is installedOn 9.65.44.55-OperatingSystem
- Source: app.TMSAgent
- Target: app.db.mssql.SqlServer
- Example: SQLServer:SQLInstance:MSS-TMSAgent monitors SQLServer:SQLInstance:MSS-SqlServer

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: Id (IdML)
Agent attribute: INODESTS.HOSTADDR

Description: The host address of the workstation where the Microsoft SQL Server instance is running.

Example: INSTANCE1_2005:PS3213W2K3:MSS-1-SqlServerProcess

- CDM attribute: ManagedSystemName
Agent attribute: INODESTS.NODE
Description: The name of the IBM Tivoli Monitoring component that provides data for the management of the Microsoft SQL Server instance.
Example: INSTANCE1_2005:PS3213W2K3:MSS
- CDM attribute: Spid
Agent attribute: KOQPRCD.PROCESSID
Description: The ID of the process.
Example: 51
- CDM attribute: Status
Agent attribute: KOQPRCD.STATUS
Description: The status of the process.
Example: runnable
- CDM attribute: PhysicalIo
Agent attribute: KOQPRCD.IO
Description: The number of accesses to the hard disk since the process started.
Example: 0
- CDM attribute: MemUsage
Agent attribute: KOQPRCD.MEMORY
Description: The number of KB allocated for this process.
Example: 2
- CDM attribute: Cpu
Agent attribute: KOQPRCD.CPU
Description: The amount of CPU time (in seconds) that the process has used since it started.
Example: 0.1317
- CDM attribute: Cmd
Agent attribute: KOQPRCD.COMMANDNAME
Description: The name of the executing command.
Example: SELECT
- CDM attribute: Label
Agent attribute: KOQPRCD.PROCESSID
Description: The label of the Microsoft SQL Server process.
Example: 51

TMSAgent class

The IpAddress class represents the instance of an IP address.

CDM class name

app.TMSAgent

Relationships

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.ComputerSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-ComputerSystem

runsOn

- Source: app.db.mssql.SqlServer
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-SqlServer runsOn 9.65.44.55-OperatingSystem

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerDatabase
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-SqlServerDatabase

contains

- Source: app.db.mssql.SqlServerDatabase
- Target: app.db.mssql.SqlServerTable
- Example: SQLServer:SQLInstance:MSS-1-SqlServerDatabase contains SQLServer:SQLInstance:MSS-1-0-spt_monitor-SqlServerTable

contains

- Source: app.db.mssql.SqlServer
- Target: app.db.mssql.SqlServerProcess
- Example: SQLServer:SQLInstance:MSS-SqlServer contains SQLServer:SQLInstance:MSS-1-7412-SqlServerProcess

installedOn

- Source: app.TMSAgent
- Target: sys.OperatingSystem
- Example: SQLServer:SQLInstance:MSS-TMSAgent is installedOn 9.65.44.55-OperatingSystem
- Source: app.TMSAgent
- Target: app.db.mssql.SqlServer
- Example: SQLServer:SQLInstance:MSS-TMSAgent monitors SQLServer:SQLInstance:MSS-SqlServer

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: Id (IdML)
Agent attribute: INODESTS.HOSTADDR
Description: The host address of the workstation where the Microsoft SQL Server instance is running.
Example: INSTANCE1_2005:PS3213W2K3:MSS-TMSAgent
- CDM attribute: ManagedSystemName
Agent attribute: INODESTS.NODE
Description: The name of the IBM Tivoli Monitoring component that provides data for the management of the Microsoft SQL Server instance.
Example: INSTANCE1_2005:PS3213W2K3:MSS
- CDM attribute: ManagedObjectName
Agent attribute: INODESTS.NODE
Description: The name of the IBM Tivoli Monitoring component that provides data for the management of the Microsoft SQL Server instance.
Example: p@INSTANCE1_2005:PS3213W2K3:MSS
- CDM attribute: SoftwareVersion
Agent attribute: INODESTS.VERSION
Description: The version of the Microsoft SQL Server agent.
Example: 06.12.00
- CDM attribute: ProductCode
Agent attribute: INODESTS.PRODUCT
Description: The product code of Microsoft SQL Server agent.
Example: 0Q
- CDM attribute: Affinity
Agent attribute: INODESTS.AFFINITIES

Description: The affinity of Microsoft SQL Server agent.
Example: 00000000000W0000000000000000000000G00004w0a7

- CDM attribute: Label
Agent attribute: INODESTS.NODE
Description: The label of the Microsoft SQL Server agent.
Example: INSTANCE1_2005:PS3213W2K3:MSS-TMSAgent

Appendix D. Integration with Tivoli Business Service Manager

Microsoft SQL Server agent provides data to create, update the status of, and view IBM Tivoli Business Service Manager services.

The Tivoli Management Services Discovery Library Adapter (DLA) and Discovery Library Toolkit provides data for the Tivoli Business Service Manager service models. The Tivoli Event Integration Facility (EIF) probe updates the status of these services, and you use the Tivoli Enterprise Portal to view the data for the services. To implement the integration of the agent with Tivoli Business Service Manager, perform the integration tasks.

Components for integrating with Tivoli Business Service Manager

The data for integrating with Tivoli Business Service Manager is supplied through the following components: Tivoli Management Services Discovery Library Adapter (DLA) and Discovery Library Toolkit, Tivoli Event Integration Facility (EIF) probe, and Tivoli Enterprise Portal Server.

Tivoli Management Services Discovery Library Adapter (DLA) and Discovery Library Toolkit

By using data from the Tivoli Management Services Discovery Library Adapter, you can build Tivoli Business Service Manager service models that include resources monitored by the Microsoft SQL Server.

The DLA files can be imported directly into Tivoli Business Service Manager using the Discovery Library Toolkit or they can be loaded into IBM Tivoli Application Dependency Discovery Manager (Tivoli Business Service Manager) and then fed into Tivoli Business Service Manager using the Discovery Library Toolkit.

See the following sources for more information about the DLA and Discovery Library Toolkit:

- Resources and relationships that are discovered by the Microsoft SQL Server and included in Tivoli Management Services DLA files: Appendix C, “Discovery Library Adapter for the Microsoft SQL Server agent,” on page 385
- Using the Tivoli Management Services DLA: *IBM Tivoli Monitoring Administrator's Guide*
- Using the Discovery Library Toolkit: *Tivoli Business Service Manager Customization Guide*

Tivoli Event Integration Facility (EIF) probe

Situation events detected by the Microsoft SQL Server can update the status of services in Tivoli Business Service Manager.

The situation events are forwarded from IBM Tivoli Monitoring to the Netcool/OMNIbus Probe for the Tivoli Event Integration Facility. The Microsoft SQL Server provides a probe rules file that updates its events with information to identify the affected service in Tivoli Business Service Manager. The EIF probe then forwards the events to the Netcool/OMNIbus ObjectServer. Tivoli Business Service Manager monitors the Netcool/OMNIbus ObjectServer for new events and updates the status of affected services.

See the following sources for more information about event integration:

- Installation (using an existing EIF probe and Netcool/OMNIbus ObjectServer installation or using Tivoli Business Service Manager to install these components): Netcool/OMNIbus Information Center or the *Tivoli Business Service Manager Installation Guide*.

- Setting up event integration between IBM Tivoli Monitoring, the EIF probe, and the Netcool/OMNIBus ObjectServer: *IBM Tivoli Monitoring Installation and Setup Guide*.
- Configuring the EIF probe to use the Microsoft SQL Server rules file after the EIF probe has been installed and configured for event integration with IBM Tivoli Monitoring: “Configuring the Tivoli Event Integration Facility (EIF) probe to enrich events”

Tivoli Enterprise Portal

You can use the integration of the Tivoli Enterprise Portal Server with Tivoli Business Service Manager to view the services in the Tivoli Business Service Manager console.

For more detailed examination and analysis, you can easily link from the Tivoli Business Service Manager console to the Tivoli Enterprise Portal Server to view the data within the Microsoft SQL Server.

Tasks to integrate the agent with Tivoli Business Service Manager

To integrate the Microsoft SQL Server agent with Tivoli Business Service Manager, you must install and configure the required components. Then, you can view the data in the Tivoli Enterprise Portal.

To integrate the Microsoft SQL Server agent with Tivoli Business Service Manager and view the data, complete the following tasks:

- Install the Discovery Library Toolkit on the Tivoli Business Service Manager server.
- Configure the Tivoli Event Integration Facility (EIF) probe to enrich Microsoft SQL Server agent events.
- Create a service in the Tivoli Business Service Manager console that you want to monitor.
- Create a data source mapping for each data source that you want to access within the Tivoli Business Service Manager.
- Configure an additional IBM Tivoli Monitoring web service for each Tivoli Enterprise Portal Server.
- View data in the Tivoli Enterprise Portal for the services that you have created to monitor through Tivoli Business Service Manager.

Installing the Discovery Library Toolkit on the Tivoli Business Service Manager

You must install the Discovery Library Toolkit on the Tivoli Business Service Manager server.

The Discovery Library Toolkit imports data from the DLA files and TADDM, which includes information about the hardware and the applications that are discovered by the source.

See "Installing the Discovery Library Toolkit" in the *Tivoli Business Service Manager Installation Guide*.

Configuring the Tivoli Event Integration Facility (EIF) probe to enrich events

The Netcool/OMNIBus Probe for Tivoli Event Integration Facility (EIF) forwards the Microsoft SQL Server agent events that are received from Tivoli Business Service Manager to the Netcool/OMNIBus ObjectServer. Tivoli Business Service Manager monitors the Netcool/OMNIBus ObjectServer for new events, and updates the status of affected services. The Microsoft SQL Server agent provides a probe rules include file that updates its events with information to identify the affected service in Tivoli Business Service Manager.

Before you begin

Install and configure the Netcool/OMNIBus ObjectServer and EIF probe and set up event integration between IBM Tivoli Monitoring and Netcool/OMNIBus.

About this task

To enable event enrichment, configure the EIF probe to use the rules file for the agent.

Procedure

1. Locate the Microsoft SQL Server agent rules file (`koq_tbsm.rules`) on a computer system where the Microsoft SQL Server agent, Tivoli Enterprise Monitoring Server, or Tivoli Enterprise Portal Server is installed. The file is in the following locations:
 - On Windows systems
The file is in the *installdir*\cms\TECLIB directory of the monitoring server, in the *installdir*\cnps\TECLIB directory of the portal server, or in the *installdir*\TMAITM6\EIFLIB directory of the agent, where *installdir* is the IBM Tivoli Monitoring or ITCAM for Microsoft Applications installation directory.
 - On Linux and UNIX systems
The file is in the *installdir*/tables/cicatrsg/TECLIB directory of the monitoring server or in the *installdir*/platform/xx/TECLIB directory of the agent, where *installdir* is the IBM Tivoli Monitoring or ITCAM for Microsoft Applications directory, *platform* is the architecture directory for the agent, and *xx* is the product code for the agent.
2. Copy the `koq_tbsm.rules` file to the following directory on the computer system where the EIF probe is installed:
 - On Windows systems
`%OMNIHOME%\probes\arch`
 - On UNIX systems
`$OMNIHOME/probes/arch`Where:
 - OMNIHOME** System-defined variable defining the installation location of Netcool/OMNIbus
 - arch** Operating system directory where the probe is installed; for example, `solaris2` when running on a Solaris system, and `win32` for a Windows system.
3. Edit the `tivoli_eif.rules` file and uncomment the `include` statement for `koq_tbsm.rules`. (The `tivoli_eif.rules` file is located in the same directory as the `koq_tbsm.rules` file.) If you are using a version of the `tivoli_eif.rules` file without an `include` statement for `koq_tbsm.rules`, add the following line after the `include` statement for `itm_event.rules`:

```
include "koq_tbsm.rules"
```
4. Restart the EIF probe.

Creating a service in Tivoli Business Service Manager

You must create a service in the Tivoli Business Service Manager console for each service that you want to monitor.

To create the services that you want to monitor in the Tivoli Business Service Manager console, see "Configuring services" in the *IBM Tivoli Business Service Manager Service Configuration Guide*.

Creating a data source mapping for each data source

You can create a data source mapping for each data source that you want to access within Tivoli Business Service Manager.

Also, you can create the data fetchers and use the data to create incoming status rules that are populated in your service templates.

For more information, see "Data sources" and "Data fetchers" in the *IBM Tivoli Business Service Manager Service Configuration Guide*.

Configuring additional IBM Tivoli Monitoring web services

You can configure additional IBM Tivoli Monitoring web services for each Tivoli Enterprise Portal Server.

To configure an additional IBM Tivoli Monitoring web service for each Tivoli Enterprise Portal server, see "Configure TBSM charts" in the *IBM Tivoli Business Service Manager Scenarios Guide*.

Viewing data in the Tivoli Enterprise Portal

From Tivoli Business Service Manager, you can open the Tivoli Enterprise Portal and view the Microsoft SQL Server agent.

You can also launch Tivoli Business Service Manager from the Tivoli Enterprise Portal.

For more information about launching applications, see "Launching to and from applications" in the *Tivoli Business Service Manager Customization Guide*.

Appendix E. Documentation library

A variety of publications are relevant to the use of the IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent.

The *IBM Tivoli Monitoring, OMEGAMON XE, and Composite Application Manager products: Documentation Guide* contains information about accessing and using publications. You can find the Documentation Guide in the following information centers:

- IBM Tivoli Monitoring and OMEGAMON XE (<http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp>)
- IBM Tivoli Composite Application Manager (<http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp>)

To open the Documentation Guide in the information center, select **Using the publications** in the **Contents** pane.

To find a list of new and changed publications, click **What's new in the information center** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center.

To find publications from the previous version of a product, click **Previous versions** under the name of the product in the **Contents** pane.

IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent library

The documentation for this agent and other product components is located in the IBM Tivoli Composite Application Manager for Microsoft Applications Information Center (http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/topic/com.ibm.itcamms.doc_6.3/welcome_msapps63.html).

One document is specific to the Microsoft SQL Server agent. The IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent User's Guide provides agent-specific information for configuring, using, and troubleshooting the Microsoft SQL Server agent.

The *Offering Guide* also provides information about installing and configuring the component products in the offering.

The **Prerequisites** topic in the information center contains information about the prerequisites for each component.

Use the information in the user's guide for the agent with the *Tivoli Enterprise Portal User's Guide* to monitor Microsoft SQL Server resources.

Prerequisite publications

See the following publications to gain the required prerequisite knowledge:

- *IBM Tivoli Monitoring Troubleshooting Guide*
- *IBM Tivoli Monitoring: Upgrading from Tivoli Distributed Monitoring*
- *IBM Tivoli Monitoring: Upgrading from V5.1.2*
- *IBM Tivoli Monitoring Universal Agent User's Guide*
- *IBM Tivoli Monitoring Universal Agent API and Command Programming Reference Guide*
- *Tivoli Enterprise Portal User's Guide*

- *IBM Tivoli Monitoring Administrator's Guide*
- *IBM Tivoli Monitoring Agent Builder User's Guide*
- *IBM Tivoli Monitoring Command Reference*
- *Configuring IBM Tivoli Enterprise Monitoring Server on z/OS*
- *IBM Tivoli Monitoring Installation and Setup Guide*
- *IBM(r) Tivoli(r) Monitoring High Availability Guide for Distributed Systems*
- *IBM Tivoli Monitoring: Messages*
- *IBM Tivoli Monitoring Troubleshooting Guide*
- *IBM Tivoli Monitoring Universal Agent User's Guide*
- *IBM Tivoli Universal Agent API and Command Programming Reference Guide*
- *IBM Tivoli Monitoring: i5/OS Agent User's Guide*
- *IBM Tivoli Monitoring: Linux OS Agent User's Guide*
- *IBM Tivoli Monitoring: UNIX OS Agent User's Guide*
- *IBM Tivoli Monitoring: UNIX Logs OS Agent User's*
- *IBM Tivoli Monitoring: Windows OS Agent User's Guide*
- *Tivoli Enterprise Portal User's Guide*
- *IBM(r) Tivoli(r) Performance Analyzer User's Guide*
- *IBM(r)Tivoli(r) Warehouse Proxy Agent User's Guide*
- *IBM(r)Tivoli(r) Warehouse Summarization and Pruning Agent User's Guide*

Related publications

The publications in related information centers provide useful information.

See the following information centers, which you can find by accessing Tivoli Documentation Central (<http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>):

- IBM Tivoli Monitoring
- IBM Tivoli Netcool/OMNIBus
- IBM Tivoli Application Dependency Discovery Manager
- IBM Tivoli Enterprise Console

Other sources of documentation

See the following sources of technical documentation about monitoring products:

- Service Management Connect (SMC)

See the introductory information about SMC (<http://www.ibm.com/developerworks/servicemanagement/>).

For information about Tivoli products, see the Application Performance Management community on SMC (<http://www.ibm.com/developerworks/servicemanagement/apm/index.html>).

Connect, learn, and share with Service Management professionals. Get access to developers and product support technical experts who provide their perspectives and expertise. You can use SMC for these purposes:

- Become involved with transparent development, an ongoing, open engagement between external users and developers of Tivoli products where you can access early designs, sprint demos, product roadmaps, and pre-release code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and Integrated Service Management.
- Benefit from the expertise and experience of others using blogs.

- Collaborate with the broader user community using wikis and forums.
- IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary/>) is an online catalog that contains integration documentation as well as other downloadable product extensions.
- IBM Redbook publications (<http://www.redbooks.ibm.com/>) include Redbooks® publications, Redpapers, and Redbooks technotes that provide information about products from platform and solution perspectives.
- Technotes (<http://www.ibm.com/support/entry/portal/software>), which are found through the IBM Software Support website, provide the latest information about known product limitations and workarounds.
- Tivoli wikis

Tivoli Wiki Central (<http://www.ibm.com/developerworks/wikis/display/tivoli/Home>) is the home for interactive wikis that offer best practices and scenarios for using Tivoli products. The wikis contain white papers contributed by IBM employees, and content created by customers and business partners. Two of these wikis are of particular relevance to IBM Tivoli Monitoring:

 - Tivoli Distributed Monitoring and Application Management Wiki (<http://www-10.lotus.com/ldd/tivmonitorwiki.nsf>) provides information about IBM Tivoli Monitoring and related distributed products, including IBM Tivoli Composite Application Manager products.
 - Tivoli System z® Monitoring and Application Management Wiki (<http://www.ibm.com/developerworks/wikis/display/tivoliomegamon/Home>) provides information about the OMEGAMON XE products, Tivoli NetView® for z/OS, Tivoli Monitoring Agent for z/TPF, and other System z monitoring and application management products.

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully.

The major accessibility features in this product enable users in the following ways:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the product documentation was modified to include the following features to aid accessibility:

- All documentation is available in both HTML and convertible PDF formats to give the maximum opportunity for users to apply screen-reader software.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Navigating the interface using the keyboard

Standard shortcut and accelerator keys are used by the product and are documented by the operating system. See the documentation provided by your operating system for more information.

Magnifying what is displayed on the screen

You can enlarge information in the product windows using facilities provided by the operating systems on which the product is run. For example, in a Microsoft Windows environment, you can reduce the resolution of the screen to enlarge the font sizes of the text on the screen. See the documentation provided by your operating system for more information.

Notices

This information was developed for products and services offered in the U.S.A. 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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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 states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply 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 Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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 the names and addresses used by an actual business enterprise 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. 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 IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© IBM 2009. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2009. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

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.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

- accessibility 403
- agent
 - functions 1
 - trace logs 271
 - troubleshooting 287
- attribute groups
 - cache duration for 6
 - list of all 59
 - that gather data from operating system or MS SQL Server files 13
 - that gather data from SELECT statements 8
 - when data is collected 6
- authorization
 - additional 19
 - required 19

B

- built-in troubleshooting features 269

C

- cache duration for attribute groups 6
- calculate historical data disk space 222
- capacity planning for historical data 222
- changes
 - baroc file updated 2
- cluster environment
 - creating a resource 38
 - installing and configuring the monitoring agent 37
 - making configuration changes 39
 - requirements 37
 - setting up the monitoring agent 36
- commands
 - tacmd 25
 - tacmd addSystem 35
 - tacmd configureSystem 35
 - tacmd startAgent 36
- components 4
 - IBM Tivoli Monitoring 4
- configuration
 - local 23
 - Manage Tivoli Enterprise Monitoring Services window 23
 - parameters 25
 - procedures 18
 - remote 34
 - settings 25
 - tacmd command line 35
 - Tivoli Enterprise Portal 34
- creating user ID 19

D

- data
 - trace logs 270
- data collection
 - processes 6
- data collection for attribute groups 6

- data from operating system or MS SQL Server files 13
- data from SELECT statements 8
- data provider 271
- database agent installation problems 277
 - detailed 275
- Discovery Library Adapter 395
 - See DLA
- Discovery Library Toolkit 395
 - installing 396
- disk capacity planning for historical data 222
- DLA 385, 395
 - data model 385
 - classes 385
- documentation
 - See publications

E

- Error Log File 27
- event
 - mapping 313

F

- files
 - agent trace 271
 - installation trace 271
 - other trace log 271
 - trace logs 270

G

- gathering support information 269
- granting permission 19

H

- historical data
 - calculate disk space 222
 - disk capacity planning 222
- Home Directory 27

I

- IBM Software Support 269
- IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent
 - performance considerations 301
- IBM Tivoli Enterprise Console
 - event mapping 313
- IBM Tivoli Monitoring 4
 - overview 1
- include file 396
- information, additional
 - policies 253
 - situations 227
 - Take Action commands 239
- installation
 - log file 271

- installation (*continued*)
 - problems 277
 - silent 32
- instance name 26
- instance property 35
- interface, user 5

L

- limited user permissions, upgrading your warehouse
 - with 310
- logging
 - agent trace logs 271
 - built-in features 269
 - installation log files 271
 - location and configuration of logs 270
 - trace log files 270
- Login 26

M

- Manage Tivoli Enterprise Monitoring Services
 - reconfiguring the agent 35
 - starting the agent 36
- Manage Tivoli Enterprise Monitoring Services window
 - local configuration 23
- messages
 - built-in features 269
- MS_SQL_DB_free_space_critical policy 253

O

- overview
 - IBM Tivoli Monitoring 1

P

- Password 26
- path names, for trace logs 270
- performance considerations 301
- permission, granting 19
- permissions, upgrading your warehouse with limited user 310
- policies
 - list of all 253
 - MS_SQL_DB_free_space_critical_NAME 253
 - overview 253
 - policies
 - more information 253
 - predefined 253
- prerequisite checker 16
- prerequisites 16
- probe rules file
 - include 396
- problems and workarounds 277
- problems with monitoring SQL Server 286
- properties
 - instance 35
- publications 399
 - prerequisite 399
 - related 400
- purposes
 - troubleshooting 269

R

- reconfiguration 35
- remote configuration 34
- remote deployment
 - troubleshooting 296
- requirements
 - configuration 15
 - installation 15

S

- SELECT statements 8
- Server Name 26
- situations
 - general troubleshooting 302
 - more information 227
 - predefined 228
 - specific troubleshooting 301
- SQL Server problems 286
- starting the monitoring agent 35
 - when 19
- support
 - gathering information for 269

T

- tacmd addSystem command 35
- tacmd command line
 - reconfiguring 35
 - remote configuration 35
- tacmd configureSystem command 35
- tacmd startAgent 36
- Take Action commands
 - more information 239
 - troubleshooting 304
- target application
 - problems 286
- Tivoli Business Service Manager
 - components for integrating with 395
 - configuring additional IBM Tivoli Monitoring web services 398
 - creating a service 397
 - creating data source mapping 397
 - installing Discovery Library Toolkit 396
 - integration 395
 - launching from Tivoli Enterprise Portal 398
 - Tivoli Enterprise Portal Server
 - Tivoli Event Integration Facility (EIF) probe 395
 - viewing data in Tivoli Enterprise Portal 398
- Tivoli Business Service Managerintegration tasks 396
- Tivoli Common Reporting
 - troubleshooting 305
- Tivoli Enterprise Portal
 - configuring the agent remotely 34
 - reconfiguring agent instance 35
- Tivoli Event Integration Facility (EIF) probe
 - configuring 396
- trace logs 270
 - directories 270
- tracing 275
- transport command 286
- troubleshooting 269, 277
 - agents 287
 - built-in features 269
 - installation 277
 - installation logs 271

- troubleshooting *(continued)*
 - remote deployment 296
 - situations 301, 302
 - Take Action commands 304
 - Tivoli Common Reporting 305
 - uninstallation 277
 - uninstallation logs 271
 - workspaces 297

U

- uninstallation
 - log file 271
 - problems 277
- upgrading for warehouse summarization 309
- upgrading your warehouse with limited user
 - permissions 310
- User ID 26
- user ID, creating 19
- user interfaces options 5
- user permissions, upgrading your warehouse with
 - limited 310

W

- warehouse summarization
 - upgrading for
 - overview 309
- warehouse summarization upgrading
 - effects on summarized attributes 309
 - tables in the warehouse 309
- Windows agent installation problems 277
- workarounds 277
 - agents 287
 - remote deployment 296
 - situations 301
 - Take Action commands 304
 - Tivoli Common Reporting 305
 - workspaces 297
- workspaces
 - troubleshooting 297



Product Number: 1234-SS1

Printed in USA

SC23-8880-06

