

Spotlight® on SQL Server Enterprise 11.7.1

## Reporting and Trending Guide



## Copyright 2016 Quest Software Inc. ALL RIGHTS RESERVED.

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

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

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

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site ([www.quest.com](http://www.quest.com)) for regional and international office information.

### Patents

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

### Trademarks

Quest, Toad, Toad World, Spotlight and the Quest logo are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit [www.quest.com/legal](http://www.quest.com/legal). All other trademarks and registered trademarks are property of their respective owners.

### Legend

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

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

An information icon indicates supporting information.

Spotlight on SQL Server Enterprise Reporting and Trending Guide

Updated - December 2016

Version - 11.7.1

# Contents

- Reporting and Trending** ..... 6
- Enable collection of data in the Spotlight Statistics Repository** ..... 7
  - Create Spotlight Statistics Repository dialog ..... 8
- Spotlight collections for Reporting and Trending** ..... 9
- Reports** ..... 10
  - Sample reports ..... 11
  - Sample Reports (Converted from Spotlight Views) ..... 14
    - Cross Reference Views to Reports ..... 14
  - Locate reports ..... 15
  - Create reports ..... 15
- Query the Spotlight Statistics Repository** ..... 16
  - Schema of the Spotlight Statistics Repository ..... 17
  - Data flow into the Spotlight Statistics Repository ..... 19
  - Browse dimension tables and retrieve data ..... 21
    - Browse dimension tables ..... 22
    - Query the fact table ..... 24
    - Query Custom Counters ..... 26
    - Query alarm data ..... 27
  - Table definitions ..... 28
    - spotlight\_datasources ..... 28
    - spotlight\_domains ..... 28
    - spotlight\_monitored\_objects ..... 28
    - spotlight\_perfdata ..... 29
    - spotlight\_stat\_classes ..... 30
    - spotlight\_stat\_keys ..... 30
    - spotlight\_stat\_names ..... 30
    - spotlight\_technologies ..... 31
    - spotlight\_timestamps ..... 31
  - Stored procedures ..... 32
    - Example stored procedures ..... 34
  - Specifications ..... 34
    - spotlight\_get\_datasources ..... 34
      - Syntax ..... 35
      - Input parameters ..... 35

Returned values .....	35
spotlight_get_domains .....	35
Syntax .....	35
Input parameters .....	35
Returned values .....	35
spotlight_get_monitored_objects .....	35
Syntax .....	36
Input parameters .....	36
Returned values .....	36
spotlight_get_table_columns .....	36
Syntax .....	36
Input parameters .....	36
Returned values .....	36
spotlight_get_table_data_instance .....	37
Syntax .....	37
Input parameters .....	37
Returned values .....	37
spotlight_get_table_keys .....	37
Syntax .....	37
Input parameters .....	37
Returned values .....	37
spotlight_get_table_range .....	38
Syntax .....	38
Input parameters .....	38
Returned values .....	38
spotlight_get_table_span .....	38
Syntax .....	38
Input parameters .....	38
Returned values .....	38
spotlight_get_tables .....	39
Syntax .....	39
Input parameters .....	39
Returned values .....	39
spotlight_get_technologies .....	39
Syntax .....	39
Input parameters .....	39
Returned values .....	39
spotlight_sqlanalysis_results .....	40
Syntax .....	40
Input parameters .....	40

Returned values .....	40
spotlight_ssr_version .....	40
Syntax .....	40
Input parameters .....	40
Returned values .....	40
<b>Deployment</b> .....	<b>41</b>
Maintenance plan for Spotlight Statistics Repository and Playback Database .....	41
Database configuration .....	41
Fragmentation and index performance .....	41
Database backup .....	42
<b>About us</b> .....	<b>43</b>
Contacting Quest .....	43
Technical support resources .....	43

# Reporting and Trending

Reporting and Trending assists with long-term trend analysis. Reports show data collected over time, not real-time data.

Reports are generated in an external viewer. Reports can be customized. You can use the sample reports provided or build on them to create your own reports by querying the data yourself using stored procedures.

Collection of data for reporting and trending is disabled by default. The first step is to [Enable collection of data in the Spotlight Statistics Repository](#).

# Enable collection of data in the Spotlight Statistics Repository

Long term history for reporting and trending is stored in the Spotlight Statistics Repository. Collection of data for reporting and trending is disabled by default. Use the Spotlight Client to enable the collection of data and then create the Spotlight Statistics Repository.

## Open this screen from the Spotlight Client

1. Click **Configure | Diagnostic server**.



2. Select **Configure the Spotlight Statistics Repository**.

## Select the Diagnostic Server

For federated Spotlight Diagnostic Server (Refer to the *Spotlight on SQL Server Deployment Guide*) you will be prompted to select the Spotlight Diagnostic Server to configure. Each Spotlight Diagnostic Server is independently configured for the Spotlight Statistics Repository.

## Configure the Spotlight Statistics repository

Option	Description
Enable use of the Spotlight Statistics Repository	Select to store data in the Spotlight Statistics Repository. Clear to stop storing (not store) data in the Spotlight Statistics Repository.
Instance	The name of the SQL Server instance to host the Spotlight Statistics Repository. Click the down arrow to list available instances.

**i** Note: Due to the size of the Spotlight Statistics Repository, it is recommended that this database is not installed on a SQL Express instance. The Spotlight Statistics Repository is not supported on High Availability Always On and mirrored databases. The Spotlight Statistics Repository cannot be installed on SQL Server 2000.

Option	Description
Authentication	<p>The authentication for Spotlight to use to connect to the Spotlight Statistics Repository.</p> <p>Select <b>Windows Authentication (using Diagnostic Server credentials)</b> to use the account that runs the Spotlight Diagnostic Server. Ensure this account is trusted by the SQL Server.</p> <p>Alternatively, select <b>SQL Server authentication</b> and fill in the <b>User</b> and <b>Password</b> fields. The SQL Server user must have sufficient privileges to connect to the SQL Server instance hosting the Spotlight Statistics Repository.</p>
Database	The name of the Spotlight Statistics Repository.
Create	Click <b>Create</b> to create a new Spotlight Statistics Repository. For more information, see <a href="#">Create Spotlight Statistics Repository dialog</a> on page 8.
Store alarms in the repository for ... days	The number of days to store alarm data in the repository, by default 30.

# Create Spotlight Statistics Repository dialog

This dialog is opened from [Enable collection of data in the Spotlight Statistics Repository](#) on request to create the Spotlight Statistics Repository.

## Authentication

This authentication is used to create the Spotlight Statistics Repository. Ensure this authentication has permission to create databases on the SQL Server instance.

Option	Description
Windows authentication (Client)	Select Windows authentication to create the Spotlight Statistics Repository using the Windows credentials of the logged in user on the Spotlight Client.
SQL Server authentication	Alternatively, select SQL Server authentication and fill in the User and Password fields. This SQL Server user must have permission to create databases on the SQL Server instance.

## Configure repository data and log files...

Click **Configure repository data and log files** to change the default name of the Spotlight Statistics Repository or the location of the corresponding data and log files. By default, the repository and data log files are created in the same location as those of the model database.

---

# Spotlight collections for Reporting and Trending

Data for Reporting and Trending is stored in the Spotlight Statistics Repository.



Note:

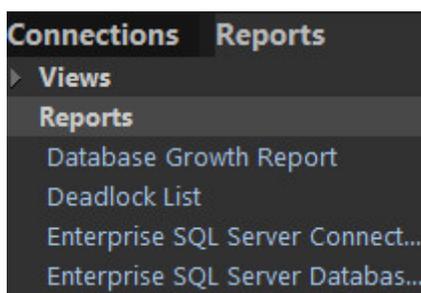
- To view / change the rate of data collection, see the online help. Collections are configurable unless otherwise stated. Collections with Cfg=N are non configurable.
- Some collections used for [Reporting and Trending](#) are also used for displays (alarms, home page components and drilldowns). Collections for [Reporting and Trending](#) are performed entirely independently of the same collections that run in the foreground / background.
- The SQL code is available for most SQL Server collections. For more information, see the online help.

# Reports

Spotlight reports are displayed in an external viewer.

## To view a Spotlight report from the Spotlight Client

1. Select **Reports | Reports**.



2. Select a report. For more information on the sample reports, see [Sample reports](#).
3. If SQL Server authentication is required, enter the user and password details.
4. Fill in the input parameters. These vary but many require the start date, end date and connection name.
5. Click **View Report**.

### **i** TIP:

- Change reports by selecting a different report from the **Report** list at the top of the Spotlight Report Viewer. Check the input fields are still correct and click **View Report**.
- If you change the report criteria (time frame, Spotlight on SQL Server or Connection Name), click **View Report** to see an updated version of the report.
- Save reports in PDF, or as an .xls (Microsoft Excel) file using the Save icon on the Report toolbar.
- Print, search, and refresh reports using the Report toolbar.
- Create your own reports or customize existing reports.

# Sample reports

To open these reports, see [Reports](#). The input parameters for these reports vary. Many require the start date, end date and connection name (monitored object).

**Table 1: Sample reports that can be run from the Spotlight Client and the Reporting Server**

Sample Report	Description						
Database Growth Report	<p>The report contains the following charts for each database:</p> <table border="1"> <thead> <tr> <th>Chart</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Data Growth Trend</td> <td> <p>Shows the following statistics for each database on the specified server:</p> <ul style="list-style-type: none"> <li>• Data size</li> <li>• Data allocated</li> </ul> </td> </tr> <tr> <td>Log Growth Trend</td> <td> <p>Shows the following statistics for the transaction log for each database on the specified server:</p> <ul style="list-style-type: none"> <li>• Log size</li> <li>• Log allocated</li> </ul> </td> </tr> </tbody> </table>	Chart	Description	Data Growth Trend	<p>Shows the following statistics for each database on the specified server:</p> <ul style="list-style-type: none"> <li>• Data size</li> <li>• Data allocated</li> </ul>	Log Growth Trend	<p>Shows the following statistics for the transaction log for each database on the specified server:</p> <ul style="list-style-type: none"> <li>• Log size</li> <li>• Log allocated</li> </ul>
Chart	Description						
Data Growth Trend	<p>Shows the following statistics for each database on the specified server:</p> <ul style="list-style-type: none"> <li>• Data size</li> <li>• Data allocated</li> </ul>						
Log Growth Trend	<p>Shows the following statistics for the transaction log for each database on the specified server:</p> <ul style="list-style-type: none"> <li>• Log size</li> <li>• Log allocated</li> </ul>						
Deadlock List	The report shows all lock requests resulting in a deadlock on the specified server for a set time period.						
SQL Server Connection Report	The report shows connection statistics for: all connections, user connections, system connections, active connections, inactive connections or blocked connections.						
SQL Server Database Space	Use this report to identify databases nearing capacity.						
SQL Server Inventory	This report shows an inventory of SQL Servers in the enterprise: version and edition, memory size and number of CPU.						
SQL Server Job Report	This report shows all jobs run on the server. Identify failed jobs. Failed jobs are colored red.						
SQL Server Performance Summary	This report shows a performance summary of servers in the enterprise. How well have resources been utilized over the selected time frame?						
Windows Disk Space	Use this report to identify Windows server disks nearing capacity.						
Windows Inventory	This report shows an inventory of Windows servers in the enterprise: operating system version and service pack, memory size and number of CPU.						
Windows Performance Summary	This report shows variation of performance over time in the areas of: CPU, Paging, Memory, Network Bytes and Queue Length.						
Windows Virtualization Summary	This report shows variation of CPU performance over time.						

Sample Report	Description
Server Alarm Report	The Server Alarms report shows all alarms that have occurred on the specified server for a set time period.
SQL Server Configuration Report	This report shows for a selected SQL Server instance: SQL Server Properties, SQL Server Instance Configuration and SQL Server xp_MSver Results.
SQL Server Health Report	<p>The Server Health report shows the following charts for the specified server for a set time period:</p> <ul style="list-style-type: none"> <li>• Number of Logged on Users</li> <li>• Batches, Compiles and Recompiles</li> <li>• Buffer Cache Hit Ratio</li> <li>• Procedure Cache Hit Ratio</li> <li>• Page Life Expectancy</li> <li>• Server Memory</li> </ul>
SQL Server IO Report	<p>This report charts logical and physical IO over time for a selected SQL Server instance.</p> <p>The Top 8 Database Files by File IO Reads Rate chart displays the file IO reads rate for the eight database files on the SQL Server that have the highest IO read rate over the selected time range.</p> <p>The Top 8 Database Files by File IO Writes Rate chart displays the file IO writes rate for the eight database files on the SQL Server that have the highest IO write rate over the selected time range.</p>
SQL Server Memory Report	This report charts page life expectancy, procedure and buffer cache hit ratios and allocated memory sizes over time for a selected SQL Server instance.
SQL Server Single Server Connection Report	This report charts the number of connections to the selected SQL Server Instance over time (active, blocked, total, inactive, system and user connections).
SQL Server Transaction Daily Range Summary	This report charts database transactions over time for a selected SQL Server Instance.
SQL Server Transaction Monthly Summary	This report shows monthly average and maximum transaction rates for each database on a selected SQL Server Instance.
Windows Memory Report	This report charts Physical Memory (free and total), Virtual Memory (free and total) and Paging over time for a selected Windows server.
Windows Server Health Report	<p>This report shows the following charts for the specified server for a set time period:</p> <ul style="list-style-type: none"> <li>• Total CPU Utilization</li> <li>• Context Switches/Sec</li> <li>• Paging Rate</li> <li>• Network Bandwidth Utilization</li> <li>• Disk Queue Length</li> </ul>

Sample Report	Description
Windows Virtualization Report	This report charts CPU Percentage Analysis and CPU Utilization over time for a selected Windows server.

**Table 2: Sample Reports that can be run only from the Reporting Server**

Sample Report	Description
SQL Server Keyed Metric Report	This report shows metric data for a selected statistic class and a selected statistic name (depending on the statistic class selected), for a selected instance over the selected time frame.
SQL Server Metric Report	This report shows metric data for a selected statistic class and a selected statistic name (depending on the statistic class selected), for a selected instance over the selected time frame.
SQL Server Transaction Single Day Summary	This report shows the daily transaction rate (per minute) for all databases on the selected instance for the selected date.
Windows Disk Metric Report	This report shows the bytes read per second, bytes written per second, IO's per second, percent busy, queue length, reads per second, service time, writes per second for the logical and physical disks on the selected server for the selected date.
Windows Keyed Metric Report	This report shows metric data for a selected statistic class and a selected statistic name (depending on the statistic class selected), for a selected server over the selected time frame.
Windows Metric Report	This report shows metric data for a selected statistic class and a selected statistic name (depending on the statistic class selected), for a selected server over the selected time frame.

# Sample Reports (Converted from Spotlight Views)

## Cross Reference Views to Reports

This table provides a Cross Reference Chart back to the Spotlight Views (from Spotlight on SQL Server versions 11.6 and earlier) and the current Spotlight Reports.

<b>Spotlight Views (from versions 11.6 and earlier)</b>	<b>Spotlight Reports (current release)</b>
Other   Custom Counters	Custom Counters Report
SQL Server Configuration   Hardware Configuration	Windows Hardware Configuration Report
SQL Server Configuration   Instance Configuration	SQL Server Configuration Report
SQL Server Configuration   Windows Configuration	Windows Configuration Report
SQL Server Environment   Large Databases	SQL Server Large Databases Report
SQL Server Health   Database File IO Statistics	The original SQL Server IO Report has been expanded to include the information from this report.
SQL Server Health   Database Files	SQL Server Health Database Files Report
SQL Server Health   Database Settings	Database Settings Report
SQL Server Health   Index Summary	SQL Server Index Summary Report
SQL Server Health   Redundant and Reverse Indexes	SQL Server Redundant and Reverse Indexes Report
SQL Server Health   SQL Server General Statistics	SQL Server General Statistics Report
SQL Server Health   Table Summary	SQL Server Health Table Summary Report
SQL Server Health   Tables Missing Clustered Indexes or Primary Keys	SQL Server Tables Missing Clustered Indexes or Primary Keys Report
SQL Server Health   Windows General Statistics	Windows General Statistics Report

### Spotlight Views (from versions 11.6 and earlier)

### Spotlight Reports (current release)

SQL Server Operations   All SQL Agent Jobs	All SQL Agent Jobs Report
SQL Server Operations   Failed SQL Agent Jobs	Failed SQL Agent Jobs Report
SQL Server Performance   Blocking	SQL Server Blocking Report
SQL Server Performance   SQL Analysis - Workload	SQL Analysis Workload Report
SQL Server Performance   Wait Statistics	SQL Server Wait Statistics Report

## Locate reports

Sample reports that are run from the Spotlight Client are located in the **Plug-ins\Trending\SSRS** folder in the Spotlight on SQL Server Client installation folder.

## Create reports

The Spotlight Report Viewer displays reports using SQL Server 2005 or 2008 Report Definition Language files (.rdl). You can build on the sample reports or create your own reports by querying the data from the Spotlight Statistics Repository.

### ***To use your own report definition files***

Set the location of your .rdl files in Spotlight.

- a. Select **Reports**.
- b. Right-click **Reports**.
- c. Select **Settings**.
- d. Navigate to the folder containing the rdl files you want to use.
- e. Click **OK**.

or

Copy your .rdl files into the **Plug-ins\Trending\SSRS** folder in the Spotlight on SQL Server Client installation folder. If you do not have write access to the Program Files folders listed above, you can create the following folder and use it as an alternative: **<user>\Documents\Spotlight Reports**

# Query the Spotlight Statistics Repository

Data is collected from the SQL Server instance then written to the Spotlight Statistics Repository.

**Table 3: What you need to know to query the Spotlight Statistics Repository**

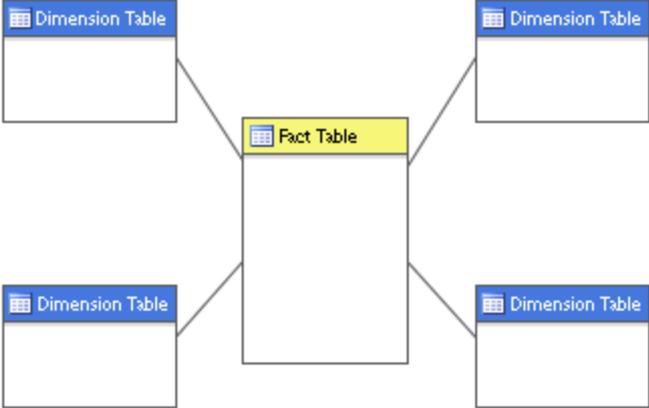
Section	Description
<a href="#">Schema of the Spotlight Statistics Repository</a>	Show how the repository is structured and the data it contains.
<a href="#">Data flow into the Spotlight Statistics Repository</a>	Collected data is inserted into the relevant dimension tables and the fact is written into the fact table (spotlight_perfdata).
<a href="#">Browse dimension tables and retrieve data</a>	Use Spotlight API stored procedures to browse dimension tables in the repository. Then use additional stored procedures or custom T-SQL to retrieve data.

**Table 4: Reference to the schema of the Spotlight Statistics Repository**

Section	Description
<a href="#">Table definitions</a>	Definitions to repository objects (tables) created by Spotlight on SQL Server in the Spotlight Statistics Repository.
<a href="#">Stored procedures</a>	Definitions to repository objects (procedures) created by Spotlight on SQL Server in the Spotlight Statistics Repository. Stored procedures return metadata about the Spotlight Statistics Repository or data from within the Spotlight Statistics Repository.
<a href="#">Specifications</a>	Specifications to the stored procedures in the Spotlight Statistics Repository.

# Schema of the Spotlight Statistics Repository

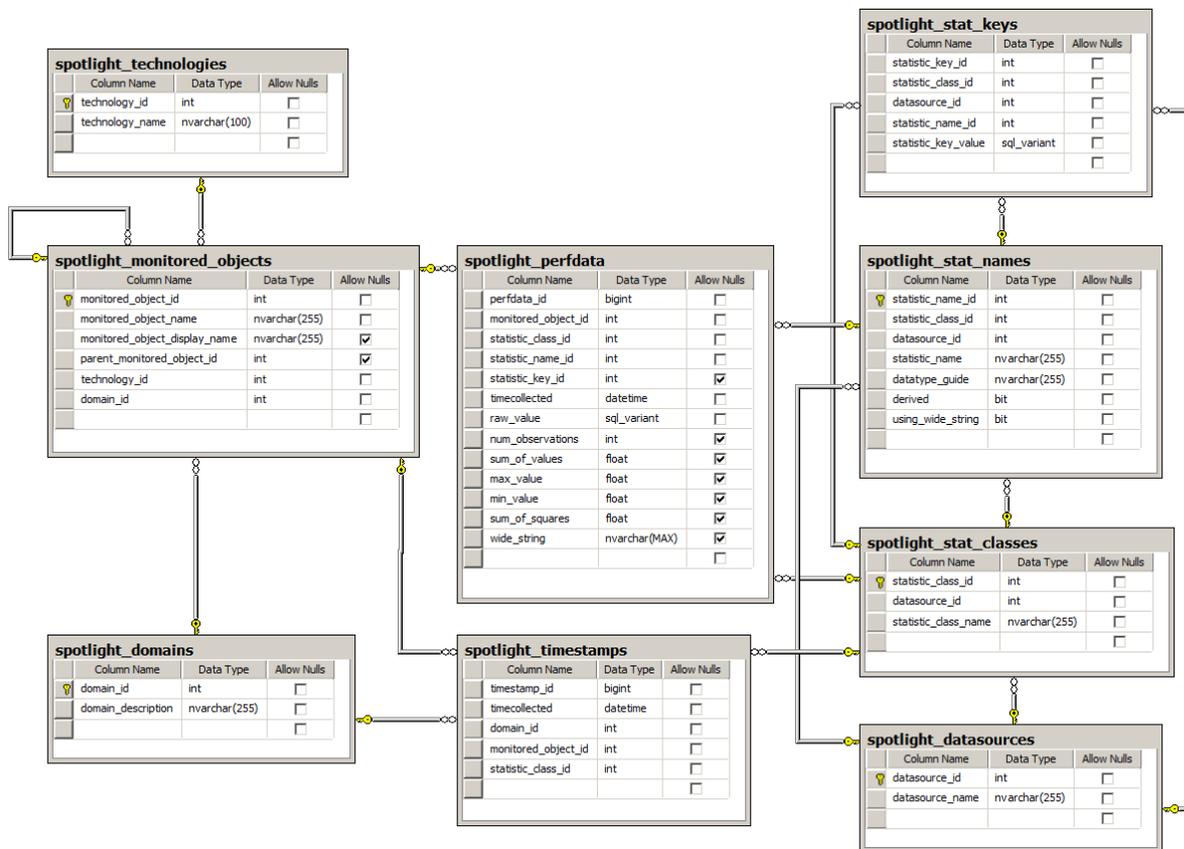
The Spotlight Statistics Repository is based on a star schema because of the simplicity it offers for data storage and retrieval. Star schemas are a simple schema consisting of a centralized fact table connected to multiple dimension tables.



Fact tables hold factual data and the dimension tables hold descriptive data. One of the benefits of a star schema is the speed of data retrieval. With star schemas, you can use relatively simple SQL queries to return particular information about, in this case, a SQL Server instance. Later in this chapter we'll show you how you can query the Spotlight Statistics Repository using Spotlight API stored procedures and T-SQL.

In the Spotlight Statistics Repository, the fact table (`spotlight_perfdata`) stores facts related to SQL Server performance, and the dimension tables hold definition data such as SQL Server instance information and attributes.

The schema of the Spotlight Statistics Repository is as follows:



The main dimension tables in the Spotlight Statistics Repository are as follows:

- **Spotlight\_datasources**

Contains data used to categorize the collections that are performed by Spotlight. There is one row in this table for each category of data collected. Examples of categories are SQL Server, Windows, and Diagnostic Server.

- **Spotlight\_domains**

Contains data about the Diagnostic Servers feeding the Spotlight Statistics Repository. In the case of multiple Diagnostic Servers feeding the Spotlight Statistics Repository, there is one row per Diagnostic Server. The Domain\_description column contains the Diagnostic Server host name.

- **Spotlight\_monitored\_objects**

Contains data about each of the SQL Server or Windows servers that a Diagnostic Server is monitoring. There is one row in this table for each server being monitored by each Diagnostic Server.

- **Spotlight\_techonologies**

Contains data that categorizes the collections that are performed. There is one row in this table for each category of data collected.

The remaining dimension tables contain data about the collections being stored in the Spotlight Statistics Repository. For more information, see [Table definitions](#) on page 28.

# Data flow into the Spotlight Statistics Repository

Data is collected from the SQL Server instance by the Spotlight Diagnostic Server then written to the Spotlight Statistics Repository. The data is inserted into the relevant dimension tables and the fact is written into the fact table (spotlight\_perfdata).

Let's use the example of data from the sqlblockinglist collection written to the Spotlight Statistics Repository. In the following chart, the Diagnostic Server, *DS123*, is collecting data about SQL Server instance, *SQLServer789*, running on Windows host, *Windows01*.

**Figure 1: How data flows into the Spotlight Statistics Repository**

If it doesn't already exist, an entry is written to spotlight\_domains about the Diagnostic Server that is monitoring the SQL Server instance.

spotlight_domains	
domain_id	domain_description
1	DS123 : 3843

If it doesn't already exist, an entry is written to spotlight\_technologies to categorize the data that has been collected.

spotlight_technologies	
technology_id	technology_name
1	database/sqlserver

If it doesn't already exist, an entry is written to spotlight\_monitored\_objects containing data about the SQL Server host that the Diagnostic Server is monitoring.

spotlight_monitored_objects					
monitored_object_id	monitored_object_name	monitored_object_display_name	parent_monitored_object_id	technology_id	domain_id
1	SQL Server 789	SQL Server 789	Windows 01	1	1

If it doesn't already exist, an entry is written to spotlight\_datasources to categorize the data that has been collected.

spotlight_datasources	
datasource_id	datasource_name
1	sqlserver_spotlight
2	windows
3	availability
4	diagnosticserver

If it doesn't already exist, an entry is written to spotlight\_stat\_classes containing the names of the collection being stored in the Spotlight Statistics Repository.

spotlight_stat_classes		
statistic_class_id	datasource_id	statistic_class_name
10	1	sqlblockinglist

If it doesn't already exist, an entry is written to spotlight\_stat\_names containing data about the columns belonging to the collection.

spotlight_stat_names					
statistic_name_id	statistic_class_id	datasource_id	statistic_name	datatype_guide	derived
102	10	1	type	string	0
103	10	1	cpu	integer	0
104	10	1	status	string	0
105	10	1	resource	string	0
106	10	1	blockedbyspid	integer	0
107	10	1	ntuser	string	0
108	10	1	command	string	0

If it doesn't already exist, an entry is written to spotlight\_stat\_keys containing key values for the collection. This allows queries to get data from the fact table based on a key value without having to scan values in the fact table itself.

spotlight_stat_keys				
statistic_key_id	statistic_class_id	datasource_id	statistic_name_id	statistic_key_value
83395	10	1	756	749356181
83396	10	1	756	-799168423
83397	10	1	756	-308268424

If it doesn't already exist, an entry is written to spotlight\_timestamps containing a summary of the times for which data was collected, qualified by monitored object.

Spotlight\_timestamps allows Spotlight to quickly determine the times for which historical data is available for individual monitored objects, without the overhead of scanning the potentially large fact table *spotlight\_perfdata*.

spotlight_timestamps				
timestamp_id	timecollected	domain_id	monitored_object_id	statistic_class_id
618930	2009-04-29 11:00:00.000	1	1	10

The last step occurs when the fact is written to spotlight\_perfdata.

This entry in the fact table (spotlight\_perfdata) shows the collection "sqlblockinglist" using dimension tables in the Spotlight Statistics Repository.

spotlight_perfdata						
perfdata_id	monitored_object_id	statistic_class_id	statistic_name_id	statistic_key_id	timecollected	raw_value
208842475	1	10	106	51483	2009-03-20 00:03:11.217	76

## Browse dimension tables and retrieve data

Use Spotlight API stored procedures to browse dimension tables in the repository. Then use additional stored procedures or custom T-SQL to retrieve data.

---

## Scenario

Your manager has asked you to provide statistics on database performance for the last quarter. You decide to query the Spotlight Statistics Repository to return file I/O statistics for the databases your manager is interested in.

To retrieve this data, you first browse the dimension tables to become familiar with the data in the Spotlight Statistics Repository. Once familiar, you write queries using the Spotlight stored procedures to get data from the fact table. You then refine the queries using custom T-SQL to retrieve more specific data from the Spotlight Statistics Repository.

---

## Browse dimension tables

First up, we're going to use some of the Spotlight API stored procedures to browse the dimension tables and get the range of values available for use in our queries. For more information on the Spotlight stored procedures used here, see [Stored procedures](#) and [Specifications](#).

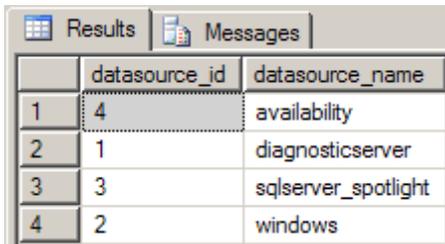
In **SQL Server Management Studio**, create a new query against the Spotlight Statistics Repository. We want to know the names of the available data sources and the names and domains of the monitored objects.

Let's start exploring the Spotlight Statistics Repository by looking at datasources. Datasources are broad categories of data.

To get a list of data sources in the repository, run the following query:

```
exec spotlight_get_datasources
```

This returns the following results:



	datasource_id	datasource_name
1	4	availability
2	1	diagnosticserver
3	3	sqlserver_spotlight
4	2	windows

We can see here that we have data from Windows, SQL Server, the Diagnostic Server and alarms (availability).

Now we want to know the name of the Diagnostic Server storing data in the Spotlight Statistics Repository. We can determine this by running:

```
exec spotlight_get_domains
```

To get the following result:



	domain_id	domain_description
1	1	DS123:3843

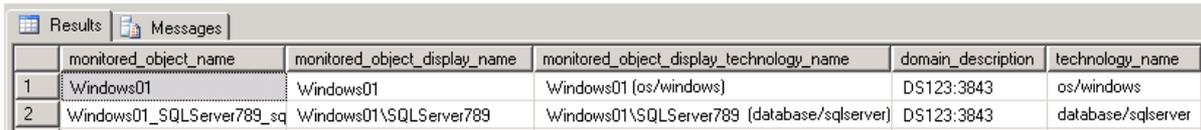
We can see here that there is one Diagnostic Server writing data to this repository. There can be multiple Diagnostic Servers, each one resulting in a row.

Finally, we want to know information about the servers being monitored.

We can do this by running the following query:

`exec spotlight_get_monitored_objects`

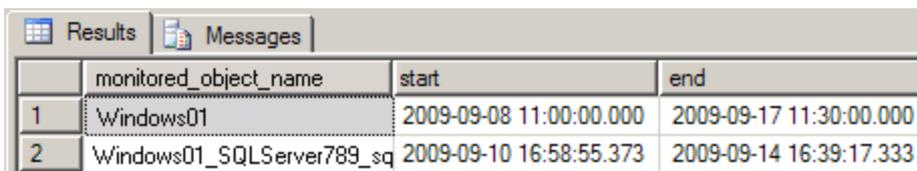
To get the following results:



	monitored_object_name	monitored_object_display_name	monitored_object_display_technology_name	domain_description	technology_name
1	Windows01	Windows01	Windows01 (os/windows)	DS123:3843	os/windows
2	Windows01_SQLServer789_sq	Windows01\SQLServer789	Windows01\SQLServer789 (database/sqlserver)	DS123:3843	database/sqlserver

A common question is “what is the range of timestamps available for a monitored server?” The following query returns the time ranges for all monitored servers:

```
select
    so.monitored_object_name,
    min(st.timecollected) as 'start',
    max(st.timecollected) as 'end'
from
    spotlight_timestamps st
join spotlight_monitored_objects so on st.monitored_object_id =
    so.monitored_object_id
group by
    so.monitored_object_name
```



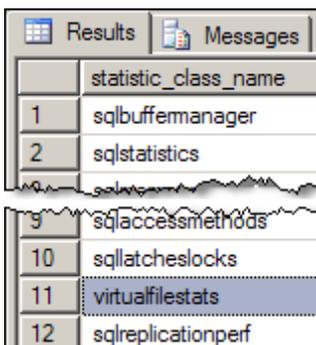
	monitored_object_name	start	end
1	Windows01	2009-09-08 11:00:00.000	2009-09-17 11:30:00.000
2	Windows01_SQLServer789_sq	2009-09-10 16:58:55.373	2009-09-14 16:39:17.333

Now that we have some information about the Diagnostic Server, monitored servers and categories of collections in the Spotlight Statistics Repository, we want to know what statistics are available to us.

The `spotlight_get_tables` stored procedure returns a list of tables in the repository for a given data source. Since we are looking for SQL Server data, from our first query above, we know that the datasource is “sqlserver\_spotlight”, so we will use that as the parameter in the following query:

`exec spotlight_get_tables 'sqlserver_spotlight'`

The following results are returned.



	statistic_class_name
1	sqlbuffermanager
2	sqlstatistics
3	sqlaccessmethods
4	sqlaccessmethods
5	sqlaccessmethods
6	sqlaccessmethods
7	sqlaccessmethods
8	sqlaccessmethods
9	sqlaccessmethods
10	sqllatcheslocks
11	virtualfilestats
12	sqlreplicationperf

We’re going to look at virtual file stats as this will give us the file IO statistics we need for our report.

To find out what columns are available to us in the virtual file stats table use the following query. With this query we need to specify the datasource and table name.

```
exec spotlight_get_table_columns 'sqlserver_spotlight','virtualfilestats'
```

We need to produce the report over a time range, but for curiosity's sake, we would like to find out the time range of all stored data for a particular domain name, monitored object, and table combination.

To do this, use the spotlight\_get\_table\_span stored procedure and specify the domain name, monitored object, and table name.

```
exec spotlight_get_table_span 'DS123:3843','Windows01_SQLServer789_sqlserver','virtualfilestats'
```

	Earliest	Latest
1	2009-03-26 11:15:16.153	2009-07-26 17:15:17.113

## Query the fact table

Now we can put it all together and query the fact table. We'll use the spotlight\_get\_table\_range stored procedure and specify the following:

- start date
- end date
- domain name
- monitored object
- table name

Spotlight\_get\_table\_range returns data for the requested time range for a table, for a particular monitored object and domain.

This gives us the following query:

```
exec spotlight_get_table_range '2009-03-26 11:15:16.153','2009-07-26 17:15:17.113', 'DS123:3843','Windows01_SQLServer789_sqlserver', 'virtualfilestats'
```

Which returns all columns for the table 'virtualfilestats':

	timecollected	monitor...	kbread	dbfilename	readrate	dbandfgname	filename	numberwrites	kbondisk	iostallwr
1	2009-05-26 11:15:16.153	2	3432	master.C:\Progr...	0.000277777	master.PRIMA...	C:\Program Files\Micro...	1	4096	0
2	2009-05-26 11:15:16.153	2	324	master.C:\Progr...	0	master.<Log>	C:\Program Files\Micro...	32	1280	643
3	2009-05-26 11:15:16.153	2	3624	model.C:\Progra...	0	model.PRIMA...	C:\Program Files\Micro...	2	2240	19
4	2009-05-26 11:15:16.153	2	264	model.C:\Progra...	0	model.<Log>	C:\Program Files\Micro...	8	1024	46
5	2009-05-26 11:15:16.153	2	6152	msdb.C:\Progra...	0	msdb.PRIMARY	C:\Program Files\Micro...	5	7616	90
6	2009-05-26 11:15:16.153	2	464	msdb.C:\Progra...	0	msdb.<Log>	C:\Program Files\Micro...	6	2048	155
7	2009-05-26 11:15:16.153	2	6992	QuestWorkData...	0	QuestWorkDa...	C:\Program Files\Micro...	28	5120	746
8	2009-05-26 11:15:16.153	2	448	QuestWorkData...	0	QuestWorkDa...	C:\Program Files\Micro...	92	2048	2084
9	2009-05-26 11:15:16.153	2	16128	tempdb.C:\Progr...	0	tempdb.PRIM...	C:\Program Files\Micro...	2161	10944	21120
10	2009-05-26 11:15:16.153	2	452	tempdb.C:\Progr...	0	tempdb.<Log>	C:\Program Files\Micro...	12552	1280	37591
11	2009-05-26 11:15:16.153	2	1752	QuestSoftware...	0	QuestSoftware...	C:\Program Files\Micro...	1	5120	8

Although there is a lot of useful data in the query results, it's not exactly what we are looking for. We need to refine the data returned from the Spotlight Statistics Repository. Following on from the example above, if we

want to retrieve only certain columns from the 'virtualfilestats' table, we can use custom T-SQL to return specific data.

To return only the 'iorate', 'disk', 'readrate', and 'writesrate' columns from the 'virtualfilestats' table, we can use the following custom T-SQL:

```
select
    sp.timecollected,
    max(case when sn.statistic_name = 'dbname' then sp.raw_value end) as
    'dbname',
    max(case when sn.statistic_name = 'dbfilename' then sp.raw_value end) as
    'dbfilename',
    max(case when sn.statistic_name = 'disk' then sp.raw_value end) as 'disk',
    max(case when sn.statistic_name = 'iorate' then sp.raw_value end) as 'iorate',
    max(case when sn.statistic_name = 'readrate' then sp.raw_value end) as
    'readrate',
    max(case when sn.statistic_name = 'writesrate' then sp.raw_value end) as
    'writesrate'
from
    spotlight_perfdata sp
    join spotlight_stat_classes sc on sp.statistic_class_id = sc.statistic_class_id
    join spotlight_stat_names sn on sp.statistic_name_id = sn.statistic_name_id
    join spotlight_monitored_objects so on sp.monitored_object_id =
    so.monitored_object_id
where
    sc.statistic_class_name = 'virtualfilestats'
    and so.monitored_object_name = 'Windows01_SQLServer789_sqlserver'
    and sp.timecollected between '2009-09-01' and '2009-09-30'
group by
    sp.timecollected, sp.statistic_key_id
order by
    sp.timecollected
```

This gives us the following results:

	timecollected	dbname	dbfilename	disk	iorate	readrate	writesrate
103	2009-09-15 06:00:00.000	tempdb	tempdb.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0.00305555555555...	0.000416666666666...	0.00263888888888...
104	2009-09-15 06:00:00.000	tempdb	tempdb.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0.0211574074074074	0	0.0211574074074074
105	2009-09-15 12:00:00.000	master	master.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0.00023148148148...	0.000231481481481...	0
106	2009-09-15 12:00:00.000	master	master.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0	0	0
107	2009-09-15 12:00:00.000	model	model.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0	0	0
108	2009-09-15 12:00:00.000	model	model.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0	0	0
109	2009-09-15 12:00:00.000	msdb	msdb.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0.00046296296296...	0.00046296296296...	0
110	2009-09-15 12:00:00.000	msdb	msdb.C:\Program Files\Microsoft SQL Server\MSSQ...	C:	0	0	0
111	2009-09-15 12:00:00.000	QuestSoftware_SoSS	QuestSoftware_SoSS.C:\Program Files\Microsoft SQL ...	C:	0.00013888888888...	0.00013888888888...	0
112	2009-09-15 12:00:00.000	QuestSoftware_SoSS	QuestSoftware_SoSS.C:\Program Files\Microsoft SQL ...	C:	0	0	0
113	2009-09-15 12:00:00.000	QuestWorkDatabase	QuestWorkDatabase.C:\Program Files\Microsoft SQL ...	C:	0.00287037037037...	0.000416666666666...	0.0024537037037037
114	2009-09-15 12:00:00.000	QuestWorkDatabase	QuestWorkDatabase.C:\Program Files\Microsoft SQL ...	C:	4.62962962962963...	0	4.62962962962963...
115	2009-09-15 12:00:00.000	QuestWorkDatabase_4.0	QuestWorkDatabase_4.0.C:\Program Files\Microsoft S...	C:	0.00023148148148...	0.000231481481481...	0
116	2009-09-15 12:00:00.000	QuestWorkDatabase_4.0	QuestWorkDatabase_4.0.C:\Program Files\Microsoft S...	C:	0	0	0
117	2009-09-15 12:00:00.000	QuestWorkDatabase_pj	QuestWorkDatabase_pj.C:\Program Files\Microsoft S...	C:	9.25925925925926...	9.25925925925926...	0
118	2009-09-15 12:00:00.000	QuestWorkDatabase_pj	QuestWorkDatabase_pj.C:\Program Files\Microsoft S...	C:	0	0	0
119	2009-09-15 12:00:00.000	QuestWorkDatabaseSH	QuestWorkDatabaseSH.C:\Program Files\Microsoft S...	C:	0.00013888888888...	0.00013888888888...	0

## Query Custom Counters

You have created a custom counter in Spotlight on SQL Server and now want to create a report that contains data from that counter.

Querying custom counters is exactly the same as querying other data in the Spotlight Statistics Repository. SQL Server custom counters are stored in the statistic class 'sqlcustomcounters' and Windows custom counters are stored in the class 'windowscustomcounters'. You can use the API stored procedures to retrieve date ranges and column names with these like any other data collection.

**Figure 2: Sample T-SQL statement that retrieves SQL Server custom counter values.**

```

select
sp.timecollected,
max(case when sn.statistic_name = 'countername' then sp.raw_value end) as 'countername',
max(case when sn.statistic_name = 'countervalue' then sp.raw_value end) as 'countervalue'
from
spotlight_perfdata sp
join spotlight_stat_classes sc on sp.statistic_class_id = sc.statistic_class_id
join spotlight_stat_names sn on sp.statistic_name_id = sn.statistic_name_id
join spotlight_monitored_objects so on sp.monitored_object_id = so.monitored_object_id
where
sc.statistic_class_name = 'sqlcustomcounters'
and so.monitored_object_name = 'Windows01_SQLServer789_sqlserver'
and sp.timecollected between '2009-09-01' and '2009-09-30'
group by
sp.timecollected, sp.statistic_key_id
order by
sp.timecollected

```

**Figure 3: Results returned for the custom counter "Pages Allocated"**

	timecollected	countername	countervalue
1	2009-09-17 12:00:00.000	Pages Allocated	7
2	2009-09-17 18:00:00.000	Pages Allocated	8
3	2009-09-18 00:00:00.000	Pages Allocated	8

## Query alarm data

Querying alarm data is the same as querying other data stored in the repository. Alarm data is stored in a statistic class called 'alarms'.

**Figure 4: Sample T-SQL statement that can be used to retrieve alarm data. Required parameters: monitored object name, start and end times.**

```
select
sp.timecollected,
max(case when sn.statistic_name = 'severity' then sp.raw_value end) as 'severity',
max(case when sn.statistic_name = 'text' then sp.raw_value end) as 'text',
max(case when sn.statistic_name = 'rule' then sp.raw_value end) as 'rule',
max(case when sn.statistic_name = 'action' then sp.raw_value end) as 'action',
max(case when sn.statistic_name = 'key' then sp.raw_value end) as 'key',
max(case when sn.statistic_name = 'raw text' then sp.raw_value end) as 'raw text'
from
spotlight_perfdata sp
join spotlight_stat_classes sc on sp.statistic_class_id = sc.statistic_class_id
join spotlight_stat_names sn on sp.statistic_name_id = sn.statistic_name_id
join spotlight_monitored_objects so on sp.monitored_object_id = so.monitored_object_id
where
sc.statistic_class_name = 'alarms'
and so.monitored_object_name = 'Windows01_SQLServer789_sqlserver'
and sp.timecollected between '2009-09-08' and '2009-09-17'
group by
sp.timecollected, sp.statistic_key_id
order by
sp.timecollected
```

**Figure 5: T-SQL statement results**

	timecollected	severity	text	rule	action	key	raw text
1	2009-09-08 12:50:50.787	5	The buffer cache page life expectancy is 57 seconds.	Memory - Page Life Expectancy	changed	NULL	The buffer cache pag
2	2009-09-08 12:55:50.800	4	The buffer cache page life expectancy is 177 secon...	Memory - Page Life Expectancy	changed	NULL	The buffer cache pag
3	2009-09-08 13:05:50.830	1	NULL	Memory - Page Life Expectancy	changed	NULL	
4	2009-09-09 11:38:44.030	4	The buffer cache page life expectancy is 120 secon...	Memory - Page Life Expectancy	changed	NULL	The buffer cache pag
5	2009-09-09 11:48:44.040	1	NULL	Memory - Page Life Expectancy	changed	NULL	
6	2009-09-10 09:52:31.957	4	The buffer cache page life expectancy is 294 secon...	Memory - Page Life Expectancy	changed	NULL	The buffer cache pag
7	2009-09-10 09:57:31.970	1	NULL	Memory - Page Life Expectancy	changed	NULL	
8	2009-09-10 12:17:32.370	4	The buffer cache page life expectancy is 296 secon...	Memory - Page Life Expectancy	changed	NULL	The buffer cache pag
9	2009-09-10 12:22:32.383	1	NULL	Memory - Page Life Expectancy	changed	NULL	
10	2009-09-10 15:34:04.153	5	Data file growths remaining could not be evaluated ...	Databases - Dynamic Data File Gro...	changed	master	Data file growths reme
11	2009-09-10 15:34:04.153	5	Log file growths remaining could not be evaluated b...	Databases - Dynamic Log File Growt...	changed	master	Log file growths rema
12	2009-09-10 15:34:24.093	2	Collection 'SQL CPU Percent' failed : Windows data...	Monitored Server - Collection Execut...	changed	SQL CPU Percent	Collection '{0}' failed
13	2009-09-10 15:35:24.170	1	NULL	Monitored Server - Collection Execut...	changed	SQL CPU Percent	
14	2009-09-10 15:39:04.200	1	NULL	Databases - Dynamic Data File Gro...	changed	master	

## Table definitions

### spotlight\_datasources

spotlight\_datasources contains data used to categorize the collections that are performed. There is one row in this table for each category of data collected.

Column	Description
datasource_id	Generated ID.
datasource_name	Name of the data source.

### spotlight\_domains

spotlight\_domains contains data about Spotlight Diagnostic Servers. If you have multiple Spotlight Diagnostic Servers feeding the Spotlight Statistics Repository, there is one row per Spotlight Diagnostic Server in this table.

Column	Description
domain_id	Generated ID.
domain_description	The name of the domain (analogous to Spotlight Diagnostic Server host name – means that this Spotlight Statistics Repository can hold data from multiple Spotlight Diagnostic Servers).

### spotlight\_monitored\_objects

spotlight\_monitored\_objects contains data about the SQL Servers and Windows hosts that a Spotlight Diagnostic Server is monitoring. There is one row in this table for each SQL Server or Windows host being monitored by each Spotlight Diagnostic Server.

Column	Description
monitored_object_id	Generated ID.
monitored_object_name	Name of monitored object (Windows host name or SQL Server Instance Name etc).
monitored_object_display_name	A user-friendly interpretation of the encoded name that the Spotlight Diagnostic Server uses to uniquely identify a monitored object. This name is suitable for display on reports or other user-visible outputs.
parent_monitored_object_id	The monitored_object_id value of another monitored object upon which this one depends, or null if no such dependency exists. For example, a SQL Server database's parent will be the Windows host upon which it runs.
technology_id	The technology of this monitored object.
domain_id	The domain (that is, Spotlight Diagnostic Server) this monitored object belongs to.

## spotlight\_perfdata

spotlight\_perfdata is the fact table into which all of the data is stored. Each fact is effectively matrixed by the hierarchy of attributes that denote which Spotlight Diagnostic Server, monitored object, collection and collection column the datum originally came from.

Column	Description
perfdata_id	Generated ID.
monitored_object_id	The monitored object that this fact came from.
statistic_class_id	The collection that this fact came from.
statistic_name_id	The column in the collection that this fact came from.
statistic_key_id	The key column of the collection that this fact came from.
timecollected	The time this fact was collected from the monitored object.
raw_value	The raw value returned by the collection for that column.
num_observations	Internal use only.
sum_of_values	Internal use only.
max_value	Internal use only.
min_value	Internal use only.

Column	Description
sum_of_squares	Internal use only.
wide_string	An alternative to raw_value that is used where the width of the data collected is too large for a variant data type. Either wide_string will contain a value or raw_value but never both.  The using_wide_string column in the spotlight_stat_names table shows whether a statistic name is using the wide_string column or not.

## spotlight\_stat\_classes

spotlight\_stat\_classes contains the names of collections being stored in the Spotlight Statistics Repository. There is one row for each table/data source combination.

Column	Description
statistic_class_id	Generated ID.
datasource_id	The data source that this collection belongs to.
statistic_class_name	The name of the collection.

## spotlight\_stat\_keys

spotlight\_stat\_keys contains key values for all of the collections stored in the Spotlight Statistics Repository. Each key value/table/data source combination has a row in this table. The purpose of this table is to allow queries to get data from the fact table based on a key value without having to scan values in the fact table itself.

Column	Description
statistic_key_id	Generated ID.
statistic_class_id	The statistic_class (Collection) this column belongs to.
datasource_id	The data source of the collection that this column belongs to.
statistic_name_id	Generated ID.
statistic_key_name	The name of the key column.

## spotlight\_stat\_names

spotlight\_stat\_names contains data about the columns belonging to collections. There is one row in this column for each column/table/data source combination.

Column	Description
statistic_name_id	Generated ID.
statistic_class_id	The statistic_class (collection) this column belongs to.
datasource_id	The data source of the collection that this column belongs to.
statistic_name	The name of the column.
datatype_guide	This column provides a broad guide to the type of data that can be expected in facts for this named statistic. Values can be "string", "numeric" or "date".
derived	Whether this is a derived column or not.
using_wide_string	Whether the value for this data source is stored in the wide_string column in the spotlight_perdata table or not. If not it is stored in the raw_value column.

## spotlight\_technologies

spotlight\_technologies contains data used to categorize the collections that are performed. There is one row in this table for each category of data collected.

Column	Description
technology_id	Generated ID.
technology_name	The name of the technology.

## spotlight\_timestamps

spotlight\_timestamps contains a summary of times for which data was collected, qualified by monitored entity. It allows Spotlight to quickly determine the times for which historical data is available for individual monitored entities, without the overhead of scanning the potentially large fact table *spotlight\_perdata*. There is one row in this table for each category of data collected.

Column	Description
timestamp_id	Generated ID.
timecollected	The time this fact was collected from the monitored object.
domain_id	The domain that this fact belongs to.
monitored_object_id	Generated ID.
statistic_class_id	The collection that this fact came from.

# Stored procedures

The following stored procedures return either metadata about the Spotlight Statistics Repository or data from within the Spotlight Statistics Repository.

Stored procedure	Description
spotlight_get_datasources	Returns a list of all data sources in the repository (effectively the contents of spotlight_datasources).
spotlight_get_domains	Returns a list of all the Spotlight Diagnostic Servers putting data into this repository.
spotlight_get_monitored_objects @domain_name, @technology_name	Returns a list of names of all the monitored objects that have data in this repository for a given domain_name. If domain_name is NULL all monitored objects are returned. Also returns the monitored object display name, domain description, and technology name.
spotlight_get_table_columns @datasource_name, @table_name	Returns a list of columns for a table belonging to a data source (effectively the contents of spotlight_stat_names).
spotlight_get_table_data_instance @date_ time, @domain_name, @monitored_ object_name, @table_name	Returns the rows and columns from a table with a timestamp that most closely matches @date_time.
spotlight_get_table_keys @datasource_ name, @table_name	Returns a list of keys for a table belonging to a data source (effectively the contents of spotlight_stat_keys).
spotlight_get_table_range @start_date, @end_date, @domain_name, @monitored_object_name, @table_name	Returns the rows and columns from a table between @start_date and @end_date.
spotlight_get_table_span @domain_ name, @monitored_object_name, @table_ name	Returns the time span of all stored data for a particular domain_name, monitored object, and table combination.
spotlight_get_tables @datasource_name	Returns a list of tables in the repository for a given data source (effectively the contents of spotlight_stat_classes).
spotlight_get_technologies	Returns a list of technologies in the repository (effectively the contents of spotlight_technologies).
spotlight_rt_get_aggregated_workload_ analysis @start_date, @end_date, @domain_name, @monitoredobject	Returns data for the requested time range for the Workload view.
spotlight_rt_get_batch_data @start_date, @end_date, @domain_name, @monitoredobject_list, @table_name, @column_list	Returns data for the requested time range for a table and associated columns for multiple monitored objects in a domain.
spotlight_rt_get_fileio_stats @start_date, @end_date, @domain_name, @monitoredobject	Returns data for the Database File IO Statistics view.

Stored procedure	Description
spotlight_rt_get_point_data @start_date, @end_date, @domain_name_list, @monitoredobject_list, @table_name, @column_name	Returns the same data as spotlight_rt_get_batch_data but for a particular point in time.
spotlight_rt_get_point_data_top @start_date, @end_date, @domain_name_list, @monitoredobject_list, @table_name, @column_list, @sortdesc_column_list, @maxrows	Returns the top n rows of spotlight_rt_get_point_data based on the column sort order passed in.
spotlight_rt_get_waitstats @start_date, @end_date, @domain_name, @monitoredobject	Returns data for the view Wait Statistics view.
spotlight_sqlanalysis_results @domain_name, @monitored_object_name, @date_time_from, @date_time_to, @maxrows	Returns the worst performing SQL recorded in the Spotlight Statistics Repository over the supplied time span for a particular SQL Server instance.
spotlight_ssr_version	Returns the version number of the Spotlight Statistics Repository.

The following stored procedures are used for the sample reports:

Stored procedure	Description
spotlight_rep_alarms	Returns a list of alarms for a chosen monitored object from the repository.  Used in the Server Alarms report. This is a further example of how to query the Spotlight Statistics Repository to return tabular information.
spotlight_rep_cache_hit_ratios	Returns hit rates for the buffer and procedure caches for all monitored objects in the repository.  Used in the SQL Server Health report. This is an example of querying two columns from a table over time from the Spotlight Statistics Repository.
spotlight_rep_cpu_utilization	Returns the amount of CPU being used by the specified server against the total amount being used by all Windows processes.  Used in the Windows Server Health report. This is an example of querying one column from a table over time from the Spotlight Statistics Repository.
spotlight_rep_disk_growth	Returns information on disk space usage over time.
spotlight_rep_disk_statistics	Returns disk statistics (percentage busy and disk queue length) for the specified server.  Used in the Windows Server Health report. This is an example of querying multiple columns with multiple rows from a table over time

Stored procedure	Description
	from the Spotlight Statistics Repository.
spotlight_rep_growth_data_files	<p>Returns the following statistics for the specified server:</p> <ul style="list-style-type: none"> <li>• Data size</li> <li>• Data allocated</li> <li>• Index size</li> <li>• Table size</li> <li>• Log size</li> <li>• Log allocated</li> </ul> <p>Used in the Database Growth report. This is an example of querying multiple columns with multiple rows from a table over time from the repository.</p>
spotlight_rep_sql_batches_rate	Returns the rate at which batches of SQL statements are being submitted to SQL Server for execution.
spotlight_rep_top_sql	Returns the top 200 SQL Statement executions that had the highest total elapsed time in each hour.
spotlight_rep_top_sql_summary	Returns summary information about the top 200 SQL Statement executions. Returns either the dates the data was collected, the times that the data was collected or a list of databases that the data was collected for.

## Example stored procedures

The following stored procedures can be used as examples of how to retrieve data from the Spotlight Statistics Repository.

- spotlight\_rt\_get\_allsqlagentjobs
- spotlight\_rt\_get\_custom\_counter
- spotlight\_rt\_get\_failedsqlagentjobs
- spotlight\_rt\_get\_mostcommonblock
- spotlight\_rt\_get\_sqlblockedmost
- spotlight\_rt\_get\_sqlblockingmost

## Specifications

The following section contains specifications for each of the stored procedures.

### spotlight\_get\_datasources

Returns a list of all data sources in the repository (effectively the contents of spotlight\_datasources).

## Syntax

spotlight\_get\_datasources

## Input parameters

There are no input parameters for this stored procedure.

## Returned values

ColumnName	ColumnType	Description
datasource_id	int	ID of data source
datasource_name	nvarchar(255)	Name of data source

## spotlight\_get\_domains

Returns a list of all the diagnostic servers putting data into this repository.

## Syntax

spotlight\_get\_domains

## Input parameters

There are no input parameters for this stored procedure.

## Returned values

ColumnName	ColumnType	Description
domain_id	int	ID of domain
domain_description	nvarchar(255)	Descriptive name of domain

## spotlight\_get\_monitored\_objects

Returns a list of names of all the monitored objects that have data in this repository for a given domain\_name. If domain\_name is NULL all monitored objects are returned.

## Syntax

spotlight\_get\_monitored\_objects @domain\_name, @technology\_name

## Input parameters

@domain\_name

@technology\_name

## Returned values

ColumnName	ColumnType	Description
monitored_object_name	nvarchar(255)	Name of monitored object
monitored_object_display_name	nvarchar(255)	A user-friendly interpretation of the encoded name that the Spotlight Diagnostic Server uses to uniquely identify a monitored object. This name is suitable for display on reports or other user-visible outputs.
monitored_object_display_technology_name	nvarchar(255)	Combines the columns monitored_object_display_name and technology_name to display a user-friendly name that includes the monitored object name and the technology.
domain_description	nvarchar(255)	Descriptive name of domain
technology_name	nvarchar(255)	Descriptive name of technology

## spotlight\_get\_table\_columns

Returns a list of columns for a table belonging to a data source (effectively the contents of spotlight\_stat\_names).

## Syntax

spotlight\_get\_table\_columns @datasource\_name, @table\_name

## Input parameters

@datasource\_name

@table\_name

## Returned values

ColumnName	ColumnType	Description
statistic_name	nvarchar(255)	Name of table column

# spotlight\_get\_table\_data\_instance

Returns the rows and columns from a table with a timestamp that most closely matches @date\_time.

## Syntax

```
spotlight_get_table_data_instance @date_time, @domain_name, @monitored_object_name, @table_name
```

## Input parameters

@date\_time

@domain\_name

@monitored\_object\_name

@table\_name

## Returned values

The returned values are dependent on the columns implicit in the table represented by the @table\_name parameter.

# spotlight\_get\_table\_keys

Returns a list of keys for a table belonging to a data source (effectively the contents of spotlight\_stat\_keys).

## Syntax

```
spotlight_get_table_keys @datasource_name, @table_name
```

## Input parameters

@datasource\_name

@table\_name

## Returned values

ColumnName	ColumnType	Description
statistic_key_value	nvarchar(255)	Key value for procedure

# spotlight\_get\_table\_range

Returns the rows and columns from a table between @start\_date and @end\_date.

## Syntax

```
spotlight_get_table_range @start_date, @end_date, @domain_name, @monitored_object_name, @table_name
```

## Input parameters

@start\_date

@end\_date

@domain\_name

@monitored\_object\_name

@table\_name

## Returned values

The returned values are dependent on the columns implicit in the table represented by the @table\_name parameter.

# spotlight\_get\_table\_span

Returns the time span of all stored data for a particular domain\_name, monitored object, and table combination.

## Syntax

```
spotlight_get_table_span @domain_name, @monitored_object_name, @table_name
```

## Input parameters

@domain\_name

@monitored\_object\_name

@table\_name

## Returned values

The returned values are dependent on the columns implicit in the table represented by the @table\_name parameter.

# spotlight\_get\_tables

Returns a list of tables in the repository for a given data source (effectively the contents of spotlight\_stat\_classes).

## Syntax

```
spotlight_get_tables @datasource_name
```

## Input parameters

@datasource\_name

## Returned values

ColumnName	ColumnType	Description
statistic_class_name	nvarchar(255)	Name of table

# spotlight\_get\_technologies

Returns a list of technologies in the repository (effectively the contents of spotlight\_technologies).

## Syntax

```
spotlight_get_technologies
```

## Input parameters

There are no input parameters for this stored procedure.

## Returned values

ColumnName	ColumnType	Description
technology_id	int	ID of technology
technology_name	nvarchar(255)	Descriptive name of technology

# spotlight\_sqlanalysis\_results

Returns the worst performing SQL recorded in the Spotlight Statistics Repository over the supplied time span for a particular SQL Server instance.

## Syntax

```
spotlight_sqlanalysis_results @domain_name, @monitored_object_name, @date_time_from, @date_time_to, @maxrows
```

## Input parameters

@domain\_name

@monitored\_object\_name

@date\_time\_from

@date\_time\_to

@maxrows

## Returned values

Returns information about the worst performing SQL recorded in the Spotlight Statistics Repository.

# spotlight\_ssr\_version

Returns the version number of the Spotlight Statistics Repository.

## Syntax

```
spotlight_ssr_version
```

## Input parameters

There are no input parameters for this stored procedure.

## Returned values

Returns the version number.

## Deployment

### Maintenance plan for Spotlight Statistics Repository and Playback Database

It is recommended that a maintenance plan for both the Spotlight Statistics Repository and the Playback Database is implemented.

Both the Spotlight Statistics Repository and Playback Database support a large number of monitored servers and automatically maintain the age of data available. Due to the nature of the automatic purging of old data, the data may become fragmented and this may decrease the efficiency of data retrieval operations over time.

This topic provides guidance on how to configure and automate the maintenance of the Spotlight Statistics Repository and the Playback Database.

### Database configuration

The Spotlight Statistics Repository and the Playback Database do not require the database to be run under the Full Recovery model. However, since Spotlight uses the model system database to create the repository, your Spotlight repository may inherit the Full Recovery setting. Unless you are required to use the Full Recovery model and are prepared to perform regularly scheduled log file backups, we recommend you change the repository recovery model to Simple. Doing so will help maintain a considerably smaller transaction log file.

To change the repository database to Simple recovery, run the following command:

```
ALTER DATABASE [SpotlightStatisticsRepository] SET RECOVERY SIMPLE;
```

### Fragmentation and index performance

The Spotlight Statistics Repository implements a star schema because of the data warehousing-style simplicity it offers for data storage and retrieval. The star schema implements a primary fact table that references a number of dimension tables. The fact table holds collections and alarm data, and the dimension tables hold definition

information on everything from instance information to collection names. Spotlight maintains the age of the data in the repository, but the continuous INSERT and DELETE operations performed for data insertion and purging can lead to fragmentation. Fragmentation will eventually lead to increased CPU and I/O resource consumption.

Although the Playback Database uses a simpler schema, its data will also become fragmented over time as new data is inserted and old data purged.

To best address performance concerns, fragmentation should be minimized and index statistics should be kept up-to-date. Spotlight's scheduled maintenance feature does this automatically. By default, the Spotlight Diagnostic Server runs maintenance procedures daily at 3am for the Playback Database and Spotlight Statistics Repository. The maintenance schedule can be changed in Spotlight Options. See the online help for more information.

## Database backup

The Spotlight Statistics Repository and Playback Database should be included in the list of important databases which have a disaster recovery plan associated with them. The implementation of this task is dependent on the policies and infrastructure of your organization.

## We are more than just a name

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

## Our brand, our vision. Together.

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

## Contacting Quest

For sales or other inquiries, visit [www.quest.com/company/contact-us.aspx](http://www.quest.com/company/contact-us.aspx) or call +1 949 754-8000.

## Technical support resources

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

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

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