Quest Spotlight® on IBM DB2 6.10.0

Help

Contents

Welcome to Quest Spotlight® on IBM DB2

Welcome to Spotlight on IBM DB2 LUW (Linux, Unix, and Windows). The help topics in this book cover features available in the interface when Spotlight is connected to a version 9 database.

Spotlight is a powerful database monitoring and diagnostic tool. Its unique user interface provides you with an intuitive, visual representation of the activity on the DB2 databases and database partitions that you choose to monitor. Graphical flows and line and fill graphs illustrate the activity on and between database components. Icons display the value of key statistics. Using the Spotlight browser, you can switch your monitoring focus between the various databases and partitions. Additionally, you can view activity at the DB2 instance level for any of these systems.

The power of Spotlight lies in its ability to provide visual and audible warnings if the performance metrics exceed acceptable thresholds. The components and dataflows on the home page change color to show you the source of the problem.

A range of reports and graphs provide you with detailed information about a DB2 database, database partition, or the instance to which the database belongs. This information can be viewed on the screen or be printed.

You can set Spotlight options to warn you when a threshold is reached. You can define a number of thresholds so that warning messages are displayed well before the traffic levels into or out of databases become critical.

Spotlight uses a number of different techniques to warn you when your DB2 instance or a database is exceeding a threshold. For example, to issue a warning, you can configure Spotlight to change a color on the home page or drilldown, provide an audible signal, or perform an action, such as send an email message.

Related topics

Understanding drilldowns

Understanding the home page

Contacting Customer Support

New in This Release

Spotlight on IBM DB2 LUW 6.10.0 is a maintenance release with fixed internal bugs. See New in This Release for more information.

Resolved Issues and Enhancements in 6.10.0

Resolved Issue or Enhancement	Defect ID
Some security enhancements	Enhancement
Spotlight on Unix: Fixed an issue where the disk metric does not show any data in some SUSE Linux versions.	TMB-2639

Required authority or privileges to run Spotlight

Spotlight either calls DB2 snapshot APIs (for DB2 for LUW 8) or invokes IBM table functions (for DB2 for LUW 9) to collect the raw processing data that it then translates into meaningful performance information. To run Spotlight, you need any one of the following authorities on each DB2 instance you are monitoring:

- SYSMAINT

Or

- SYSCTRL

Or

- SYSADM

Or

- SYSMON (or privileges equivalent to SYSMON authority)

Note: Spotlight on IBM® DB2® LUW installation does not create any database objects. Therefore, aside from the required user authority, Spotlight does not grant or require any additional privileges on objects.

Setting up a connection profile

Before Spotlight can monitor a DB2 database or database partition, you must connect to that system. If you have not connected to the database or database partition in Spotlight before, you must set up a connection profile for that system.

For an overview of the process, see To connect to systems in Spotlight basics help. For Spotlight on IBM® DB2® LUW, use the procedure below.

Note: Options in the DB2 Monitor Switch Settings section of the Properties window are not used for connections to DB2 version 9. For version 9 and later, you need to alter the status of instance-level monitor switches to collect the desired data.

To set up a connection profile

1. Click the Connect   icon in the Spotlight toolbar to open the Spotlight Connection Manager.

2. In the Connections bar on the left, click **Spotlight on IBM DB2 LUW** to display the connections currently defined for DB2 for LUW databases and database partition.

3. In the Spotlight on IBM DB2 LUW pane on the right, double-click **New connection**.

4. In the New Connection window, identify the connection:

Select connection type	Select **Spotlight on IBM DB2 LUW**.
New connection name	Enter a unique name for the connection. This name will identify the connection on the Spotlight browser.

5. Click **OK** to open the Properties window for the new connection. The Details tab on this window is in focus.

6. In the DB2 Details section, enter the information that connects Spotlight to the DB2 database or specific database partition:

DB2 Instance	The name of the DB2 instance as defined on the database server.
DB2 Database	Enter the alias of the DB2 database that you want to monitor in this connection.
DB2 Auto Monitor Nodes	For a multiple-partition database, select this option if you want the Connection Manager to set up a separate connection automatically for every partition (node) for the database. (The alternative for monitoring a multiple-partition database is to select an option in the **DB2 Node** field.) When you select **DB2 Auto Monitor Nodes**, the **DB2 Node** field is disabled.
DB2 Auto Monitor Active Databases	Select this option if you want Spotlight to monitor all active databases residing in the instance in which the database or database partition for which you are creating this connection resides. When Spotlight connects to this database or database partition, it also connects to all the other active databases in the instance. These databases are listed in the browser pane. You can monitor their home pages and drilldowns just as you can for the database whose connection you set up.
DB2 Node	For a multiple-partition database, select **Global** if you want the Connection Manager to establish a single connection that aggregates statistics for all partitions for the database. Or select the number for the specific partition to establish a connection for that database partition only. For a single-partition database, select **Global**.
DB2 User	Enter the user ID with which to connect to the DB2 database.
DB2 Password	Enter the password with which to connect to the DB2 database.

7. In the SQL Archive Options section, set the options that run the Spotlight SQL Statement Archiver for this connection; these are used for static SQL statements:

Archive Default	To turn the Archiver on or off for this connection, select

Startup	**Archive SQL statements** or **Do not archive SQL statements**, respectively.
Archive Capacity	Enter the number of distinct static SQL statements that the Archiver will collect before it starts dropping statements.
	Enter **0** (zero) if you want to set no limit on the number of distinct SQL statements that the Archiver can collect.

8. In the Method to Retrieve Data section, select an option from which to retrieve data.

Shapshot Routines	Select this option to use the snapshot function to retrieve data from the server.
Monitor Functions	Select this option to use the mon-table to retrieve data from the server.

9. In the OS Details section, provide the information that Spotlight needs to monitor the operating system on which the database or database partition runs. For detailed information about connecting to a specific operating system type, refer to the **Spotlight on Unix**, **Spotlight on Windows**, or **Spotlight on Windows Clusters** help (accessed from the **Help|Contents** option).

OS Monitor	Select this option if you want Spotlight to monitor the operating system on which the database or database partition runs. When you select this option, specific fields listed below are enabled.
OS Type	Choose the operating system type of the target machine. For Unix machines, make sure that either RECEX or SSH is installed.
OS Local Machine	(For a Windows target machine only) Select this option if the machine you are currently using is the one you want to monitor. If you select this option, you need to complete only the **OS Address** field.
OS Domain	(For a Windows target machine only) Enter the domain to which the specified user (in the **OS User** field) belongs. If you entering no user in the **OS User** field, leave this field blank.
OS Address	Enter the address of the target machine: - For a Windows machine, enter its IP address, hostname , or URL. - For a Linux or Unix machine, enter its connection string.
OS User	Enter a valid user ID or account to log on to the target machine. For Windows machines, Spotlight needs to use an account with the privileges required to retrieve server information, query the registry, and access WMI and performance monitor objects. An account with administrative rights to the target machine allows this.
OS Password	Enter the password used to log on to the target machine.
OS Port	(For a Unix SSH target machine only) Enter the port number that Spotlight will use for its SSH (secure shell) connection to the Unix host. The default value is 22.

10. Select **Save password details (for this connection)** to save the user ID and password information with this connection. Each time you open this connection in Spotlight, you will not have to provide the user ID and password.

If you unselect this option, you are prompted to provide the **DB2 User** and **DB2 password** values each time you open this connection in Spotlight. If Spotlight monitors the operating system also for this connection, you are prompted to provide the **OS User** and **OS Password** information as well.

11. Click **OK** to add the new connection to the list of DB2 for LUW connections.

About calibration in a DB2 for LUW environment

What happens during the Spotlight calibration process

The dataflow components that display on the DB2 for LUW (Linux, Unix, and Windows) database, database partition, or instance home page on your Spotlight desktop show the rate of different types of dataflows in your system over time. The speed of a home page dataflow varies based on the volume of data being transferred within the database or the instance. The greater the data traffic, the faster the flow.

Each home page dataflow corresponds to a type of dataflow in your database, partition, or instance. In order that home page dataflows represent actual dataflows in your system accurately, Spotlight must know the normal range of activity for this system. For example, if the normal range of a specific type of data flow in your database during a specific span of time is from 0 to 100, a value of 8 is slow—the corresponding home page dataflow moves slowly. However, if the range is from 0 to 10, a value of 8 is very fast—the home page dataflow moves much faster.

Spotlight provides a calibration tool that calculates the normal range of dataflows in your connection—that is, in the database or database partition to which you are connected and in the instance to which this database or partition belongs. When you start the calibration process, Spotlight measures the flow of data through your system over a set period of time. It then uses this data to set the upper limit of each dataflow on the database or database partition home page and on the corresponding instance home page. (Spotlight always uses zero as the lowest value when calculating the normal range of values for your system.)

The calibration window

The first time you open a connection you created in Spotlight, the calibration window is displayed, prompting you to enter a length of time for the calibration. After you provide the information and close the window, the calibration process begins.

Connections to multiple-partition databases

You can set up a connection to a multiple-partition DB2 for LUW database is several ways. The way the connection is configured determines how the calibration process is run:

- If the connection is set up with **Global** selected in the **DB2 Node** field, the calibration process simultaneously calibrates all of the selected database's associated partitions (nodes) for the same length of time.

- If the connection is set up with the **DB2 Auto Monitor Nodes** option selected, the Connection Manager creates a separate connection automatically for each partition for the database. (However, you connect to and disconnect from these partitions simultaneously through the single database connection set up in the connection profile.) When you open this connection for the first time, Spotlight displays a succession of calibration windows, one for each of the partitions, prompting you for the length of time for the calibration. Calibration begins on each partition when you close its calibration window.
- If your connection is set up with a specific partition selected in the **DB2 Node** field, the calibration is run on that database partition only.

During the course of a calibration the user may still cancel, change, or stop the calibration process. The action (cancel, change, stop) affects the database and its associated partitions in the following ways:

- If your connection is set up with the **Global** option selected in the **DB2 Node** field, the action is performed at the database level across all partitions simultaneously.
- If the connection is set up at with the **DB2 Auto Monitor Nodes** option selected, the action is performed at the database level and across all partition connections simultaneously.
- If your connection is set up for the specific database partition specified in the **DB2 Node** field, the action is performed on that database partition only.

When to run calibration

Spotlight automatically runs calibration when you open a open a connection for the first time in Spotlight database for the first time. However, you can choose to run the calibration at any time.

If the data traffic in your database or database partition has increased or decreased significantly since you last ran a calibration, calibrate that system again so that your home page dataflows accurately depict your system's activity.

Saving the calibration values

When you disconnect from the database or database partition connection, you are prompted to save any connection changes. This save includes the calibration values for the connection home pages. The calibration values for the connection are used until you change them or until you perform another calibration.

Related topics

Calibrating Spotlight connections

Setting up a connection profile for a DB2 for LUW database or partition

Setting SQL Archiver Options

The SQL Statement Archiver collects data in the background for static SQL statements that have executed on the DB2 instance. You can view metrics for these statements in the Static SQL tab in the Top SQL drilldown.

The Archiver starts to collect data on static SQL as soon as you open the connection. The longer the Archiver runs, the more SQL history you can view on the drilldown. However, running the Archiver requires overhead. As a result, you might not want it activated at all times.

You can use SQL Archive options in the connection profile for a database to control whether the SQL Archiver is on or off. You can also use these options to set capacity for the SQL Archiver.

To change SQL archiving options

1. If Spotlight is currently connected to the database or partition connection for which you are changing archiving options, right-click the connection in the Spotlight browser and select **Disconnect**.

2. Click the Connect icon in the Spotlight toolbar to open the Spotlight Connection Manager.

3. In the Connections bar on the left, click the category from which to select the Spotlight connection to the database or database partition.

4. Right-click the connection in the list, and select **Properties** to open the Properties window for the connection. The Details tab on this window is in focus.

5. In the SQL Archive Options section, complete the following:

Archive Default Startup	To turn the Archiver on or off, select **Archive SQL statements** or **Do not archive SQL statements**, respectively.
Archive Capacity	Enter the number of distinct static SQL statements that the Archiver will collect before it starts dropping statements.

6. Click **OK** to save the changes.

Related topics
Static SQL tab on the top SQL drilldown
Setting up a connection profile

Navigating between database, partition, and instance home pages and drilldowns

Using the Spotlight browser, you can easily navigate between the home pages and drilldowns for the various databases and database partitions that you are monitoring.

Spotlight also provides the option to view the home page and associated drilldowns for the DB2 instance to which the database or partition currently selected in the Spotlight browser belongs. At the instance-level, the home page and drilldowns display an aggregate of performance information captured across all active databases in the DB2 instance.

To open the home page or drilldown for another DB2 database or partition

- From the Spotlight browser, select the name of the database or partition that you want to monitor. The home page or drilldown is redisplayed , showing performance information for the system you selected.

Any other drilldown you open from this home page or from the toolbar will show performance information for this selected database or partition.

To open the home page for the DB2 instance to which a database or database partition belongs

- With the database or database partition selected in the Spotlight browser, click the **DB2 Instance** button in the Spotlight toolbar.

The instance home page is displayed, showing performance information for the DB2 instance to which the database or database partition selected in Spotlight browser belongs. Any drilldown you open from this home page (or from the toolbar) will show performance information captured across all active databases in the instance.

To switch from the instance home page or drilldown back to the database- or partition-level home page

- Click the button in the Spotlight tool bar.

The home page for the database or database partition that is highlighted in the Spotlight browser (and from which you launched the instance home page) is redisplayed . Any drilldown you open from this home page (or from the toolbar) will now show performance information for this database or partition.

Understanding the home page

A main feature of Spotlight is its user interface. This interface uses two major window types–home pages and drilldowns–to help you monitor performance on your IBM DB2 for LUW (Linux, Unix, and Windows) databases and database partitions. The home pages enable you to determine the status of your system at a glance, while the drilldowns provide performance details on specific areas of your system, enabling you to pinpoint and resolve problems.

Home page views

The home page is based on the Spotlight connection you are currently monitoring–either a database or a single database partition. This home page view provides a visual representation of the performance–in the form of icons, labels, and data flows–of this system. Throughout your Spotlight session, you can switch between the home pages for the different connections by selecting that connection in the Spotlight browser. Additionally, you can click **DB2 Instance** icon in the toolbar to monitor a home page at the instance level for the currently selected connection. This home page monitors activity across all databases in the DB2 instance to which the selected database or database partition belongs.

Home page components

The icons on the home page represent components–such as processing agents, areas of memory, data storage areas–in the DB2 database or database partition that you are monitoring.

On the database and partition home pages, related components are also grouped together in panels. These panels are connected to each other by dataflows, which illustrate the rate at which data traffic is moving through the system. The groupings reflect how your database or partition works.

Likewise, on the instance home page, related information is grouped into several panels that show overview and summary data concerning all databases and partitions within the instance.

Spotlight updates the component statistics and flows on all the home pages in real time and uses color and sound to indicate performance bottlenecks.

Related topics

Understanding drilldowns

Spotlight on DB2 for LUW drilldowns

Database or partition home page

The database or partition home page appears on the Spotlight desktop when you choose to monitor a specific DB2 for LUW (Linux, Unix, and Windows) database or a partition for that database. This home page provides a visual overview of activity on your database or on the partition. Related statistics and graphs are grouped together on panels, which, in turn, are connected by a series of graphical dataflows. Spotlight updates these statistics, graphs, and flows in real time so that you see how your system is performing. Visual or audible alarms on the home page help you to spot bottlenecks quickly.

The main elements of the database home page include the following components:

- **The toolbar**–The set of icons across the top of the home page that provide quick methods for navigating between the home pages, drilldowns, and help.

- **Spotlight browser**–The panel that lists the databases and their partitions that you are monitoring. The name of the specific database or partition that you are currently monitoring on the home page is highlighted. This name changes color to match the color of the component or dataflow that represents the worst severity currently existing in that database or partition. (This identifier can also match the color of a severity occurring at the instance level on the instance home page.)

- **Statistics panels**–The vertical groupings of related components that represent areas of the database being monitored. For example, the Agents panel groups the Registered, Non-Pooled, Pending Locks, and Pending Token components, each of which monitors the status or a specific activity of database agents in the database or partition.

- **Dataflows**–The flow graphs that illustrates the rate at which data, such requests for data and the I/O generated in response to these requests, is moving through the database. Dataflows change their speed and color, alerting you to performance issues.

- **Connection identifier**–The name of the connection that you are currently monitoring. (This name was defined when the connection was set up in the Spotlight Connection Manager.) Located in the top right corner of the home page, this identifier changes color to match the color of the component or dataflow that represents the worst severity detected in the DB2 database or partition you are monitoring. (This identifier can also match the color of a severity detected at the instance level for this database or partition. You can view this severity on the instance home page.)

Home page elements

Click on any of the following for details about a specific home page element:

The toolbar
Spotlight browser
Client Apps panel
Agents panel
Buffer Pool panel
I/O Agents panel
Storage panel
Active Log panel
Dataflows
Database Information panel
Connection Identifier

How do I...?

Navigate between database, partition, and instance home pages and drilldowns
View help for components on the home page
Open the Alarm Log
Perform an immediate refresh on statistics

Related topics
Database home page alarms

Instance home page

The instance home page shows performance information for the DB2 for LUW (Linux, Unix, and Windows) instance that contains the database or database partition currently selected in the Spotlight browser. This home page provides a visual overview of activity across all active databases in the instance. Related statistics, graphs, and flows are grouped together on panels that Spotlight updates in real time so that you see how your system is performing. The use of visual or audible alarms on the home page helps you to spot bottlenecks quickly.

The main elements of the instance home page include the following:

- **The toolbar**–The set of icons across the top of the home page that provide quick methods for navigating between the home pages, drilldowns, and help.

- **Spotlight browser**–The panel that lists the databases and database partitions to which Spotlight is currently connected. The database or partition from which you launched the instance home page is highlighted. This name changes color to match the color of the component or dataflow that represents the worst severity currently existing in that database or partition (or existing at the instance-level).

- **Statistics panels**–The groupings of related components in the instance being monitored, such as instance-wide I/O and SQL activity. For example, the Global SQL Request/Response panel groups the Static SQL, Dynamic SQL, and Transactions components, each of which monitors a type of SQL activity across the instance.

- **Dataflows**–A flow that illustrates the rate at which data is moving through all active databases in the instance. Dataflows change their speed and color, alerting you to performance issues.

- **Connection identifier**–The name of the database or partition connection currently selected in the Spotlight browser. (This database or partition resides in the DB2 instance whose home page you are now viewing.) Located in the top right corner of the home page, this identifier changes color to match the color of the component or dataflow that represents the worst severity detected in the DB2 instance (or at the database or database-partition level).

Home page elements

Click on any of the following for details about a specific home page element:
The toolbar
Spotlight browser
Instance Identification panel
Global Sort/Join panel
Switches panel

HADR panel

Global SQL Request/Response panel

Global Database I/O panel

Global Application/Agent panel

Connection Identifier

How do I...?

Navigate between database, partition, and instance home pages and drilldowns

View help for components on the home page

Open the Alarm Log

Perform an immediate refresh on statistics

Related topics

Instance home page alarms

Spotlight on IBM® DB2® LUW menus

You can use menus to access the functions of Spotlight on IBM® DB2® LUW (Linux, UNIX, and Windows). There are two different types of menus:

- **Standard menus displayed via the menu bar on the Spotlight console.**

The standard menus available in Spotlight on IBM® DB2® LUW are also those that are common to all Spotlight applications. For more information, see the Spotlight menus topic in **Spotlight basics** help.

- Shortcut menus displayed when you right-click an object in a Spotlight home page or in a drilldown.

Most shortcut menus in Spotlight on IBM® DB2® LUW, and most of the options within those menus, are also those that are common to all Spotlight applications.

Selecting commands from a menu

You can select a command from a menu by typing the letter that is underlined in the command name. For example, to connect to a system, press the letter **O** when the **File** menu is displayed. If the menu is not visible you can type the **File** menu access key (**ALT+F**) and then press **O**.

Some commands also have a shortcut key. If you use shortcut keys, you do not have to display the menu first.

Spotlight on IBM® DB2® LUW toolbar

The Spotlight on IBM® DB2® LUW (Linux, Unix, and Windows) toolbar is located at the top of the main application window. The toolbar provides you with quick access to commonly-used commands and functions. Click once on a toolbar button to carry out the command. To see a description of what each button does, rest your mouse pointer over the button. Information about the button appears. If a button appears dimmed, it is unavailable.

The buttons on the Spotlight toolbar are of two kinds:

- **Buttons common to all Spotlight applications.**

The buttons that are common to all Spotlight applications are covered in the Spotlight toolbar topic in the **Spotlight basics** help.

- **Buttons that are specific to Spotlight on IBM® DB2® LUW.**

The toolbar buttons that are specific to Spotlight on IBM® DB2® LUW are those used to access drilldowns. The graphic below shows the drilldown toolbar buttons for Spotlight on IBM® DB2® LUW. For a description of each drilldown, click an icon on the graphic.

About Spotlight on IBM® DB2® LUW drilldowns

Drilldowns display detailed information about the system that you are monitoring. Each drilldown contains a series of reports and graphs that provide you with specific information about the components of your system. The statistics that are available help you identify and anticipate performance problems.

Note: The drilldowns listed below are specific to Spotlight on IBM® DB2® LUW. Other Spotlight applications have their own individual drilldowns.

For information on drilldowns in general, see the topic Components and Drilldowns in the **Spotlight basics** help. For information on accessing drilldowns, see the topic To view drilldowns. (To open the **Spotlight basics** help, select **Help|Contents**.)

For further information, see the following drilldowns:

- DB2 Instance—Opens the instance home page, which shows real-time performance information for the instance that contains the database or database partition currently selected in the Spotlight browser.

- Database Manager Summary—Shows global summary information about all active databases and partitions in the DB2 instance.

- Client Application Analysis—Shows statistics about the applications currently running at the DB2 instance or database level.

- Database Analysis—Shows statistics about the databases in the DB2 instance.

- Tablespace Analysis—Shows statistics about the tablespaces to which a database is mapped.

- Top SQL—Shows statistics about the SQL statements currently executing or recently executed on the DB2 instance.

- Buffer Pool Analysis—Shows statistics about buffer pools configured for a database.

- FCM Analysis—Shows a breakdown of message buffer traffic for the instance you are monitoring.

- Workload Management Analysis—Reports information about the Workload Management (WLM) environment for your database.

- Diagnostic Log—Lists diagnostic log, notification, and event records for a database.

- Operating System—Shows statistics about the performance, processes, and storage on the server.

- Alarm Log (common to all Spotlights).

Related topics
Viewing and editing alarms
Viewing historical data

Buffer Pool Analysis drilldown

The Buffer Pool Analysis drilldown displays key performance statistics for the bufferpools used by monitored databases and partitions. At its top level, the drilldown provides an overview of performance during the current monitoring session. This is shown in the Bufferpools tab.

From the Bufferpools tab, you can drill down to detailed information on I/O activity, I/O times, and more for a certain bufferpool. This is shown in sub drilldown tabs, which are displayed in the lower half of the window when a bufferpool is selected. The additional statistics can be used to diagnose any problems detected in the Bufferpools tab.

Using the Buffer Pool Analysis drilldown

The steps below outline the basic workflow involved in using the Buffer Pool Analysis drilldown.

To use the Buffer Pool Analysis drilldown

1. View top level tabs to detect performance problems in bufferpools. Work from the Bufferpools tab to check statistics. Work from Criteria tab to filter the list in the Bufferpools tab based on criteria such as hit ratio or number of direct reads.

2. View rows in the Bufferpools tab to focus on a certain bufferpool. When a database is being monitored, rows show information for a bufferpool across all partitions and on individual partitions. Check columns in the Bufferpools tab to see statistics for a bufferpool.

3. View sub drilldown tabs to see additional statistics for a certain bufferpool. Click the bufferpool row to display the tabs. The information they contain can help you diagnose problems detected in the Bufferpools tab.

4. Open a Toad for DB2 session to alter a bufferpool or view its properties. (You must have either product licensed and installed on the same computer where the Spotlight for DB2 client resides.)

View properties for or alter a bufferpool in a Toad for DB2 session

If you have Toad for DB2 installed, you can alter a bufferpool or view it properties.
Right-click the bufferpool entry in the list of bufferpools, and proceed as follows:

To alter a tablespace

- Select **Toad** ▸ **Alter** to open the Alter window, enabling you to change the bufferpool properties.

To view tablespace properties

- Select **Toad** ▸ **Properties** to open the Toad database explorer on the bufferpool.

On the resulting window for any of these functions, press F1 for instructions to proceed with the function.

Related Topics
Top level tabs
Sub drilldown tabs
Rows in the Bufferpools tab
Columns in the Bufferpools tab

Criteria Tab for Drilldowns

As you view drilldowns, you can filter them based on criteria specified in the Criteria tab. This is available for the following drilldowns: Database Analysis, Tablespace Analysis, Buffer Pool Analysis, Client Application Analysis, and Diagnostic Log.

Filtering lets you narrow the scope of the list shown in a drilldown. You can use it to focus on objects with a certain level of activity, certain attributes, or a certain current state. For example, you might filter a drilldown to see databases with the greatest I/O, tablespaces with the most direct writes, buffer pools with the lowest hit ratio, or applications that are currently connected.

Criteria for a drilldown can include multiple filters. The more filters you create, the narrower the scope of the drilldown list. Only items that meet filter criteria are shown. All others are filtered out. If no objects meet filter criteria, "No Data" is shown in the drilldown.

Filtering applies to the current database and current session. It is not saved when the session ends.

Note: When filtering is in use, an asterisk is shown beside the name of a drilldown's list tab. For example, when filtering is in use for the Tablespace Analysis drilldown, **Tablespaces*** is shown at the top of the Tablespaces tab.

Using the Criteria tab

Follow the procedure below to specify filter criteria. Before you create or modify filters, review the columns shown in the drilldown. The columns you can select for a filter are the ones currently displayed. If you want to filter on columns that are not currently displayed, select those columns in the **Column Organizer**. To open this window, right click any column header and select **Organize Columns…** from the right-click menu.

Note: The first filter you create determines which values are shown in the **Value** list box for the next filter. For example, if your first filter specifies "Content = Large" and you select the **Extent Size** column for the next filter, the **Value** list box shows extent sizes for tablespaces with the Large content type. If you want to use an extent size that is not listed, enter it in the **Value** list box.

To specify filter criteria

1. Click the Criteria tab for a drilldown.

2. If you plan to use multiple filters in criteria, select an option from the **Result Set will include** list box as follows:

Any	Select this option if the result set should include items that meet criteria for one or more filters. Items are included even if they do not meet criteria for all filters.

All	Select this option if the result set should include items that meet criteria for all filters. Items are included only if they meet criteria for all filters.

3. Specify criteria for a filter as follows:

a.	**Column**	From this list box, select a column for the filter. Objects are checked against values in this column. The list includes all columns currently displayed in the drilldown.
b.	**Operator**	From this list box, select a comparison operator for column values. The choices are **=** (equal to), **<>** (less than and greater than), **>** (greater than), and **<** (less than). Use the **=** (equal to)operator for a column that shows an attribute such as status, name, or ID.
c .	**Value**	From this list box, select the column value for the filter. In most cases, the values shown are for entire objects. Values for partitions are not listed. If a value for an entire object is not shown, you can enter it manually. Objects are checked to see if their values are equal to, less than and/or greater than this value (depending on the operator selected).

4. Click **Add** to add the filter to the **Active Filters** list.

5. Repeat steps 3 and 4 for each filter you want to create.

6. To remove a filter from the **Active Filters** list, select it in the list, then click **Remove**. To remove all filters from the list, click **Remove All**.

Tip: To change the value for an active filter, remove the filter from the **Active Filters** list. Then create a new filter with the same column but a different value.

7. When you finish specifying filter criteria, click the information tab for a drilldown. The tab lists objects that meet filter criteria.

Buffer Pool Analysis drilldown Details tab

The Details tab on the Buffer Pools drilldown provides the following information about the selected buffer pool:

• Configuration information

• Use of extended storage

• The amount of time required to perform I/O that involved this buffer pool

Extended storage utilization statistics

Extended storage is a secondary level of storage for buffer pools, allowing the user to access memory beyond the maximum allowed for each process. When the database has buffer pools configured for extended storage, the extended storage counters are useful. They show the number of data and index pages belonging to the tablespace that are copied in and out of extended storage. If the number of pages copied out of extended storage back into the buffer pools is higher than the number of pages copied into extended storage from the buffer pools, your database at this tablespace level is probably benefiting from extended storage.

Buffer pool identification and configuration information

The Buffer pool Details list shows current identification, extended storage, and configuration settings for the selected buffer pool.

Buffer Pool Name	Page Size (bytes)	Extended Data Pages Out
Buffer Pool ID	Total Size	Extended Index Pages Out
Node Group	Uses Extended Storage	Extended Data Pages In
Database	Files Closed	Extended Index Pages In
Pages		

Graph

The Buffer pool Timings bar graph compares the amount of time (in seconds) required to perform each of these types of I/O that involve the selected buffer pool.

Direct Write

Direct Read

Pool Async Write

Pool Write

Pool Async Read

Pool Read

Use the graph to get an overall sense of the time required to perform I/O involving this buffer pool.

Related topics

Buffer Pool Analysis drilldown

Buffer Pool Analysis drilldown I/O Activity tab

The I/O Activity tab on the Buffer Pools drilldown provides a summary of I/O activity involving the selected buffer pool. The graphs show rates for the following:

- Reads and writes that use the buffer pool

- Direct reads and writes, which do not go through the buffer pool, but access the database directly

The bottom of the tab lists average I/O times (in milliseconds) for both buffered and direct reads and writes. The average times are shown in milliseconds.

Graphs

The graphs on the I/O Activity tab keep track of the rates for I/O activity that involves the buffer pool. These rates are plotted over consecutive monitoring intervals.

Data Read Rate

Index Read Rate

Index and Data Write Rate

Direct Read and Write Rate

Troubleshooting high I/O activity on the database

On the Databases drilldown, you might notice consistently high rates for the following activities on specific databases:

- Physical read rates (compared to logical read rates)

- Synchronous read or write rates

- Direct I/O rates for a database

These high rates might be indications of performance degradation since the activities involve disk I/O and slow-downs in query processing. You can use the graphs on this I/O Activity tab to pinpoint those buffer pools that are not being used effectively to offset I/O on the database.

Related topics

Database Analysis drilldown

Buffer Pool Analysis drilldown

Buffer Pool Analysis drilldown I/O Times tab

The I/O Times tab on the Buffer Pool Analysis drilldown contains graphs that let you visually monitor average I/O times and buffer pool hit ratios for a buffer pool. The four graphs track the following:

- Average amount of time spent on direct I/O.

- Average amount of time spent on asynchronous I/O.

- Average amount of time spent on physical I/O.

- Buffer pool hit ratios.

All four graphs use counter data to show statistics as they change over time. This allows you to see when average I/O times are high and buffer pool hit ratios are low for a buffer pool.

Consistently high I/O times can be a symptom of I/O conflict. Consistently low hit ratios can be a sign that a buffer pool is not operating efficiently. Both problems degrade performance. To identify applications that are contributing to I/O conflict, work from the Client Application Analysis drilldown. To investigate buffer pool conditions, work from the Buffer Pool Analysis drilldown.

Graphs in the I/O Times tab

Click the links below to see descriptions of each graph in the I/O Times tab:

Average Direct I/O Times

Average Async I/O Times
Average Physical I/O Times
Read Hit Ratios

Functions for graphs

As you view graphs in the I/O Times tab, you can click icons in graph toolbars to:

Copy graph details to the clipboard

Change the graph view

Filter a graph

Maximize or restore a graph

Pan over a zoomed graph

Zoom a graph

Buffer Pool Analysis drilldown Statistics tab

The Statistics tab on the Buffer Pools drilldown lists performance statistics for database events associated with the selected buffer pool. Each row in the list identifies a specific event and provides counts for the event based on these time measurements:

- **Interval**—Displays the number of occurrences of the event for the most recent monitoring interval.

- **Interval Rate/Sec**—Displays the rate per second during the most recent monitoring interval.

- **Lifetime**—Displays the number of occurrences during the life of the event (since the database was activated or performance statistics were reset).

Related topics

Buffer Pool Analysis drilldown

Buffer Pool Statistics

Buffer Pool Analysis drilldown Tables tab

The Tables tab on the Buffer Pools drilldown shows a list of the tables that use the selected buffer pool. The list provides the following information for each table in a grid format. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not display in your current list view, or be in the order in which they are listed here.

Table Schema	Rows Read
Table Name	Rows Written
Tablespace	Overflow Accesses
Index Tablespace	Freeze First Column
Columns	

Related topics

Buffer Pool Analysis drilldown

Buffer Pool Analysis drilldown Tablespaces tab

The Tablespaces tab on the Buffer Pools drilldown breaks down the activity for the selected buffer pool by each tablespace associated with the buffer pool. The tab shows a list of these tablespaces and provides the following information about each in a grid format. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not display in your current list view, or be in the order in which they are listed here.

Tablespace Name	Physical Reads	Async Writes
Type	Async Reads	Avg. Async Write Time
Logical Reads	Bufferpool Hit Rate	Freeze First Column
Avg. Pool Read Time	Pool Writes	
Avg. Async Read Time	Avg. Pool Write Time	

Related topics
Tablespace Analysis drilldown
Buffer Pool Analysis drilldown

Client Application Analysis drilldown

The Client Application Analysis drilldown displays key performance statistics for all applications connected to monitored databases and partitions. At its top level, the drilldown provides an overview of performance during the current monitoring session. This is shown in the Client Applications tab.

From the Client Applications tab, you can drill down to detailed information on caching, sorting, locks, lock chains, units of work, and more for a certain application. This is shown in sub drilldown tabs, which are displayed in the lower half of the window when an application is selected. The additional statistics can be used to diagnose any problems detected in the Client Applications tab.

Using the Client Application Analysis drilldown

The steps below outline the basic workflow involved in using the Client Application Analysis drilldown.

To use the Client Application Analysis drilldown

1. View top level tabs to detect performance problems in client applications. Work from the Client Applications tab to check statistics. Work from Criteria tab to filter the list in the Client Applications tab based on criteria such as number of lock waits or timeouts.

   Tip: To reduce the overhead incurred by this drilldown, show only active applications on the Client Applications tab. Right-click within the list of applications on the tab, and select **Show only 'active' applications**. This feature is not available on IBM DB2 10.5 and later if Monitor Functions Method is used to retrieve data.

2. View rows in the Client Applications tab to focus on a application. Check columns in the Client Applications tab to see statistics for an application.

3. View sub drilldown tabs to see additional statistics for a certain application. Click the row for the application to display the tabs. The information they contain can help you diagnose problems detected in the Client Applications tab.

4. If the drilldown reveals problems with certain SQL statements, you can tune the statements using SQL Optimizer for DB2. (You must have either product licensed and installed on the on the same computer where the Spotlight for DB2 client resides.) SQL Tuning sessions can be opened from the Unit of Work - SQL sub drilldown.

5. If the drilldown reveals that certain applications are having a major impact on a database, you can force the applications to disconnect from the DB2 instance.

Related topics

Top level tabs

Sub drilldown tabs

Columns in the Client Applications tab

Criteria Tab for Drilldowns

As you view drilldowns, you can filter them based on criteria specified in the Criteria tab. This is available for the following drilldowns: Database Analysis, Tablespace Analysis, Buffer Pool Analysis, Client Application Analysis, and Diagnostic Log.

Filtering lets you narrow the scope of the list shown in a drilldown. You can use it to focus on objects with a certain level of activity, certain attributes, or a certain current state. For example, you might filter a drilldown to see databases with the greatest I/O, tablespaces with the most direct writes, buffer pools with the lowest hit ratio, or applications that are currently connected.

Criteria for a drilldown can include multiple filters. The more filters you create, the narrower the scope of the drilldown list. Only items that meet filter criteria are shown. All others are filtered out. If no objects meet filter criteria, "No Data" is shown in the drilldown.

Filtering applies to the current database and current session. It is not saved when the session ends.

Note: When filtering is in use, an asterisk is shown beside the name of a drilldown's list tab. For example, when filtering is in use for the Tablespace Analysis drilldown, **Tablespaces*** is shown at the top of the Tablespaces tab.

Using the Criteria tab

Follow the procedure below to specify filter criteria. Before you create or modify filters, review the columns shown in the drilldown. The columns you can select for a filter are the ones currently displayed. If you want to filter on columns that are not currently displayed, select those columns in the **Column Organizer**. To open this window, right click any column header and select **Organize Columns…** from the right-click menu.

Note: The first filter you create determines which values are shown in the **Value** list box for the next filter. For example, if your first filter specifies "Content = Large" and you select the **Extent Size** column for the next filter, the **Value** list box shows extent sizes for tablespaces with the Large content type. If you want to use an extent size that is not listed, enter it in the **Value** list box.

To specify filter criteria

1. Click the Criteria tab for a drilldown.

2. If you plan to use multiple filters in criteria, select an option from the **Result Set will include** list box as follows:

Any	Select this option if the result set should include items that meet criteria for one or more filters. Items are included even if they do not meet criteria for all filters.
All	Select this option if the result set should include items that meet criteria for all filters. Items are included only if they meet criteria for all filters.

3. Specify criteria for a filter as follows:

a.	**Column**	From this list box, select a column for the filter. Objects are checked against values in this column. The list includes all columns currently displayed in the drilldown.
b.	**Operator**	From this list box, select a comparison operator for column values. The choices are = (equal to), <> (less than and greater than), > (greater than), and < (less than). Use the = (equal to)operator for a column that shows an attribute such as status, name, or ID.
c.	**Value**	From this list box, select the column value for the filter. In most cases, the values shown are for entire objects. Values for partitions are not listed. If a value for an entire object is not shown, you can enter it manually. Objects are checked to see if their values are equal to, less than and/or greater than this value (depending on the operator selected).

4. Click **Add** to add the filter to the **Active Filters** list.

5. Repeat steps 3 and 4 for each filter you want to create.

6. To remove a filter from the **Active Filters** list, select it in the list, then click **Remove**. To remove all filters from the list, click **Remove All**.

Tip: To change the value for an active filter, remove the filter from the **Active Filters** list. Then create a new filter with the same column but a different value.

7. When you finish specifying filter criteria, click the information tab for a drilldown. The tab lists objects that meet filter criteria.

Client Application Analysis drilldown Details tab

The Details tab on the Client Applications Analysis drilldown provides the following information about the selected application:

• Identification details about the application

• Application status

• The amount of time the application has spent on various activities during the last monitoring interval

Application information

The Application Information section lists information describing the application:

App OS User ID	Database Path	Comm Protocol
App Code Page	Status Current	Comm DRDA AS Token
App Country Code	Status Change	TP Client User ID
App Assoc Agents	Client Name	TP Client Workstation
App ID	Client Product ID	TP Client Application
App Auth ID	Client Process ID	TP Account String
App Name	Client Agent ID	App Type
Database	Client Platform	Workload ID
Database Alias		

Application details

The Application Details sections provides statistics about the application:

App Agent ID	Last UOW Start	Prvt Wrkspc Section Lookups
App Assoc Max Agents	Last UOW Stop	Prvt Wrkspc Section Inserts
App Agents Stolen	Max Shared Workspace Size	OLAP Total Functions
Status Priority	Shrd Wrkspc Overflows	OLAP Function Overflows
Connect Start	Shrd Wrkspc Section Lookups	
Connect Complete	Shrd Wrkspc Section Inserts	
Comm Inbound Address	Max Private Workspace Size	
Statistics Reset	Prvt Wrkspc Overflows	

Related topic
Client Application Analysis drilldown

Client Application Analysis drilldown Caching tab

The Caching tab on the Client Application Analysis drilldown measures the impact of the database's resource caches (memory areas) on the selected application. This tab graphically depicts memory utilization statistics for package cache, catalog cache, the SQL workarea , and buffer pools for the application.

In general, caches enhance the performance of a database when their respective hit rates are high. Usually the hit rate is the ratio of retrieval requests that were satisfied by data already in the cache to the number of all requests–those satisfied by the cache plus those that required a disk read. (This ratio is expressed as a percentage in the various graphs on this tab.)

If you discover a cache utilization issue for a specific database on the Databases drilldown, use the information on this Caching tab to help determine whether the selected application is contributing to the problem.

Graphs

Use the following graphs to follow cache performance over consecutive monitoring intervals:

- Package Cache
- Catalog Cache
- SQL Workarea
- Buffer Pool

Related topics
Database Analysis drilldown
Client Application Analysis drilldown

Client Application Analysis drilldown Sorting tab

The Sorting tab on the Client Application Analysis drilldown uses graphs to show sort activity performed by the selected application and contrasts this activity to overall sorting activity in the database.

About sorts

An application query often requires that retrieved rows be put in a specific order–that is, be *sorted* –to satisfy a SQL sort request (as with the ORDER BY clause), to construct an index, or to perform a hash join. *Sort heap* is a block of memory that the database manager allocates to an application as workspace whenever the application requires a sort. An important factor in the performance of any sort is the size of its sort heap, which is determined by one of these parameters:

- SORTHEAP database configuration parameter
- SHEAPTHRES database manager configuration parameter

A sort consists of these steps:

- **The actual sort process**–If the information being sorted cannot fit into the sort heap, some information is written to temporary database tables. This incident is called a *sort overflow*. Sorts that do not overflow perform better than those that do.

- **Return of the sorted rows**–If sorted information can be returned directly without requiring a temporary table to store the final, sorted list of data, the sort is referred to as a *piped sort*. If the sorted information requires a table to temporarily store the data before it is returned, the sort is referred to as a *non-piped sort*. A piped sort always performs better than a non-piped sort.

Sorting problems–overflows , non-piped sorts, less-than-optimal sort heap allocation size—might indicate a need to adjust the SORTHEAP and SHEAPTHRES parameters. However, these problems might signal a need to tune the SQL within the application. Use the Appl Sort & Hash Join and the Appl Sort & Join Overflow Pct graphs on this tab to identify those applications that are generating sort and join overflows and hash loops. Then use the Top SQL drilldown to investigate problems with specific SQL in the application.

Graphs

Use the following graphs to determine sort activity and problems over consecutive monitoring intervals:

App Sort and Hash Join
App Sort and Join Overflow Pct
Database Sort & Hash Join
Database Sort & Join Overflow Pct

Related topics
Database Analysis drilldown
Top SQL drilldown
Client Application Analysis drilldown

Client Application Analysis drilldown I/O Activity tab

The I/O Activity tab on the Client Application Analysis drilldown provides a summary of I/O activity generated by the selected application. The graphs show rates for the following:

- Reads and writes that use the buffer pool

- Direct reads and writes, which do not go through the buffer pool, but access the database directly

The bottom of the tab lists average I/O times (in milliseconds) for both buffered and direct reads and writes. The average times are shown in milliseconds.

Graphs

The graphs on the I/O Activity tab keep track of the rates for I/O generated by the application over consecutive monitoring intervals.
Data Read Rate
Index Read Rate
Index and Data Write Rates
Direct Read and Write Rates

Troubleshooting high I/O activity on the database

On the Databases drilldown, you might notice consistently high rates for the following activities on specific databases:

- Physical read rates (compared to logical read rates)

- Synchronous read or write rates

- Direct I/O rates for a database

These high rates might be indications of performance degradation since the activities involve disk I/O and slow-downs in query processing. Use the graphs on this I/O Activity tab to identify the applications that are contributing to these rates.

Related topics
Database Analysis drilldown
Client Application Analysis drilldown

Client Application Analysis drilldown Locks tab

The Locking tab on the Client Application Analysis drilldown provides details about each lock that the application currently holds and each lock on which the application is waiting. This information can help you pinpoint deadlocks that this application is causing or experiencing. You can also determine whether this application is reaching the maximum number of locks available to it.

Application identification

These details identify the application:

Auth ID

Appl Name

Agent ID

If you want Spotlight to collect and display information on the Locking tab for *all* applications, select **Collect lock information for all applications**. Otherwise, the Locking tab shows lock information for the currently selected application only.

Locks Held information

The Locks Held list grid shows the total of number of locks the application is holding and provides the following information about each lock. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not appear in your current list view, or be in the order in which they are shown here.

Lock Mode	Object ID
Object Type	Lock Status
Table Schema	Escalated
Table Name	Freeze First Column
Tablespace	

Use this information to determine the locations and types of locks being held and whether the application is reaching the maximum number of locks available to it.

Locks Waiting information

The Locks Waiting list grid shows the total number of locks on which the application is waiting and provides the following information about each of these locks. Since you can choose to hide or display any column in the grid, some columns below might not appear in your current list view.

Lock Mode	Tablespace
Object Type	Holder's Agent ID
Table Schema	Holder's App ID
Table Name	Freeze First Column

Use this information to determine both the sources of contention in the database and the applications in contention for these sources.

Related topics

Client Application Analysis drilldown

Client Application Analysis drilldown Lock Chains tab

The Lock Chains tab on the Client Application Analysis drilldown provides details about an application's lock relationships with other applications. It shows:

- Which applications are holding locks that the selected application is waiting for.

- Which applications are waiting for locks held by the selected application.

- Whether locks are held by indoubt transactions.

You can use this information to determine which applications are in contention for resources.

Note: Before information can be shown in the Lock Chains tab, the lock monitor switch must be turned on.

Columns in the Lock Chains tab

All columns available for the Lock Chains tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Holding Agent ID	The application handle of the agent that currently holds a lock an application is waiting for. If the column is blank or **0** (zero) is shown, this indicates that the lock held by the agent is the top of the lock chain.
Waiting Agent ID	The application handle of the agent waiting for the lock.
Waiting Appl Name	The name of the application waiting for the lock.
Waiting Applid	The ID of the application waiting for the lock.
Waiting Authid	The authorization ID of the user who is running the application that is waiting for a lock. This is the ID that was used to connect to the database.

Related Topics

Client Application Analysis drilldown

Client Application Analysis drilldown Unit of Work tab

The Unit of Work tab on the Client Application Analysis drilldown provides details and statistics about the currently executing unit of work. When the execution time for an application's current unit of work is high (Curr Uow Sec column in the list grid), use this tab to investigate this unit of work.

Optionally from this tab, you can start a tuning session on the SQL statement you are investigating.

Application information

These details identify the selected application whose current unit of work you are investigating:

Auth ID

 App Name

Agent ID

Statement

This area shows the SQL statement that is currently executing for the unit of work.

Statement Details

These details describe the SQL statement that is currently executing for the unit of work, provide a breakdown of the SQL statement's execution time, and show the most recently completed unit of work for the application:

Statement Type	Cursor Name	Execution Time, CPU User
Statement Operation	Agents/Max Agents	Execution Time, CPU Sys
Statement Section Number	Degree of Parallelism	Execution Time, Sort
Statement Package Name	Estimated Cost (timerons)	Previous UOW, Elapsed time, Stmt
Application Creator	Estimated Rows	

Statement Activity

These statistics provide activity counts for the SQL statement that is currently executing for the unit of work and describe how the current unit of work is impacting database resources:

Sorts, Statement	Rows Deleted Internally
Sort Overflows	Rows Updated Internally
Rows Read	Rows Inserted Internally
Rows Written	Rows Fetched

Procedures

Tuning an SQL statement

Related topics
Client Application Analysis drilldown

Client Application Analysis drilldown Authorities tab

Note: This feature is not available on IBM DB2 10.5 and later if Monitor Functions Method is used to retrieve data.

The Authorities tab on the Client Application Analysis drilldown provides details about the authorities granted to the application. The authorities can be granted explicitly or implicitly to the application.

Explicit Authorities information

This section lists the authorities granted explicitly to the application. The list can include any of these authorities:

SysAdm	Connect	SysMaint
DbAdm	Create Not Fenced	Implicit Schema

CreateTab	SysCtrl	Load
BindAdd		

Implicit Authorities information

This section lists the authorities granted implicitly (indirectly) to the application from a group or from PUBLIC. The list can include any of these authorities:

SysAdm	Connect	SysMaint
DbAdm	Create Not Fenced	Implicit Schema
CreateTab	SysCtrl	Load
BindAdd		

Related topics
Client Application Analysis drilldown

Client Application Analysis drilldown Memory Pools tab

The Memory Pools tab on the Client Application Analysis drilldown helps you keep track of system memory usage as it relates to the particular application you are monitoring. It is important to remember that, with the notable exception of buffer pools, most memory is allocated dynamically on an as needed basis by the Database Manager.

The memory pools list provides the following information about each memory pool in a grid format. Since you can choose to hide or display any column in the grid, some columns below might not display in your current list view.

Pool	HWM Util Pct
Curr Size (kb)	Max Size (kb)
Curr Util Pct	Node
Size HWM (kb)	Freeze First Column

Related topics
Client Application Analysis drilldown
Memory Pool types

Client Application Analysis drilldown Statistics tab

The Statistics tab on the Client Application Analysis drilldown lists performance statistics for database events associated with the selected application. Each row in the list identifies a specific event and provides counts for the event based on these time measurements:

- **Interval**–Displays the number of occurrences of the event for the most recent monitoring interval.

- **Rate**–Displays the interval rate per second during the most recent monitoring interval.

- **Lifetime**–Displays the number of occurrences during the life of the event (since the database was activated or performance statistics were reset).

Related topics

Client Application Analysis drilldown

Application statistics

Client Applications Analysis drilldown WL Activities tab

The following information shows for each activity currently running in the selected application connection. This information is gathered from the Workload Management environment set up for your database.

Note: This information is available for DB2 9.5 and later databases only.

Partition	Parent Activity ID	Rows Read (DB2 9.7 and later)	DB Work Class ID
Coord. Partition	Nesting Level	Rows Ret. (DB2 9.7 and later)	Svc. Class Work Action Set ID
State	Invocation ID	Query Cost Est. (DB2 9.7 and later)	Svc. Class Work Class ID
Type	Routine ID	Direct Reads (DB2 9.7 and later)	Start Time
UOW ID	Utility ID	Direct Writes (DB2 9.7 and later)	
Parent UOW ID	Executable ID (DB2 9.7 and later)	Svc. Class ID	
Activity ID	Total CPU (DB2 9.7 and later)	DB Work Action Set ID	

Related topics

Client Application Analysis drilldown

Forcing applications

You can use the procedure below to force one or more applications to disconnect from the DB2 database. Consult your IBM DB2 documentation to review the implications involved in forcing an application.

To force an application

1. Open the Client Application Analysis drilldown.

2. In the list grid, right-click the applications you want to force to disconnect.

3. From the right-click menu, select **Force**. A confirmation dialog appears.

4. Click **Yes**. The applications are disconnected from the instance.
   Note: It is recommended that you pause your Spotlight session before forcing multiple applications.

Related topics

Database Analysis drilldown

The Database Analysis drilldown displays key performance statistics for all active databases and database partitions in an instance. At its top level, the drilldown provides an overview of performance during the current monitoring session. This is shown in the Databases tab.

From the Databases tab, you can drill down to detailed information on caching, sorting, locking, I/O activity, and more for a certain database or partition. This is shown in sub drilldown tabs, which are displayed in the lower half of the window when a database or partition is selected. The additional statistics can be used to diagnose any problems detected in the Databases tab.

Using the Database Analysis drilldown

The steps below outline the basic workflow involved in using the Database Analysis drilldown.

To use the Database Analysis drilldown

1. View top level tabs to detect problems in a database. Work from the Databases tab to check statistics. Work from Criteria tab to filter the list in the Database tab based on criteria such as level of physical I/O or number of locks held.

2. View Rows in the Databases tab in the Tablespaces tab to focus on a certain database or one of its partitions. Check columns in the Databases tab to see statistics for a database or partition.

3. View sub drilldown tabs to see additional statistics for a certain database. Click the database row to display the tabs. The information they contain can help you diagnose problems detected in the Databases tab.

4. Open a Toad for DB2 or session to view properties, update configuration parameters, or identify storage problems for a specific database or database partition. (You must have either product licensed and installed on the same computer where the Spotlight for DB2 client resides.)

Related Topics
Top level tabs
Sub drilldown tabs
Rows in the Databases tab
Columns in the Databases tab

Criteria Tab for Drilldowns

As you view drilldowns, you can filter them based on criteria specified in the Criteria tab. This is available for the following drilldowns: Database Analysis, Tablespace Analysis, Buffer Pool Analysis, Client Application Analysis, and Diagnostic Log.

Filtering lets you narrow the scope of the list shown in a drilldown. You can use it to focus on objects with a certain level of activity, certain attributes, or a certain current state. For example, you might filter a drilldown to see databases with the greatest I/O, tablespaces with the most direct writes, buffer pools with the lowest hit ratio, or applications that are currently connected.

Criteria for a drilldown can include multiple filters. The more filters you create, the narrower the scope of the drilldown list. Only items that meet filter criteria are shown. All others are filtered out. If no objects meet filter criteria, "No Data" is shown in the drilldown.

Filtering applies to the current database and current session. It is not saved when the session ends.

Note: When filtering is in use, an asterisk is shown beside the name of a drilldown's list tab. For example, when filtering is in use for the Tablespace Analysis drilldown, **Tablespaces*** is shown at the top of the Tablespaces tab.

Using the Criteria tab

Follow the procedure below to specify filter criteria. Before you create or modify filters, review the columns shown in the drilldown. The columns you can select for a filter are the ones currently displayed. If you want to filter on columns that are not currently displayed, select those columns in the **Column Organizer**. To open this window, right click any column header and select **Organize Columns…** from the right-click menu.

Note: The first filter you create determines which values are shown in the **Value** list box for the next filter. For example, if your first filter specifies "Content = Large" and you select the **Extent Size** column for the next filter, the **Value** list box shows extent sizes for tablespaces with the Large content type. If you want to use an extent size that is not listed, enter it in the **Value** list box.

To specify filter criteria

1. Click the Criteria tab for a drilldown.

2. If you plan to use multiple filters in criteria, select an option from the **Result Set will include** list box as follows:

Any	Select this option if the result set should include items that meet criteria for one or more filters. Items are included even if they do not meet criteria for all filters.
All	Select this option if the result set should include items that meet criteria for all filters. Items are included only if they meet criteria for all filters.

3. Specify criteria for a filter as follows:

a.	**Column**	From this list box, select a column for the filter. Objects are checked against values in this column. The list includes all columns currently displayed in the drilldown.
b.	**Operator**	From this list box, select a comparison operator for column values. The choices are **=** (equal to), **<>** (less than and greater than), **>** (greater than), and **<** (less than). Use the **=** (equal to)operator for a column that shows an attribute such as status, name, or ID.
c .	**Value**	From this list box, select the column value for the filter. In most cases, the values shown are for entire objects. Values for partitions are not listed. If a value for an entire object is not shown, you can enter it manually. Objects are checked to see if their values are equal to, less than and/or greater than this value (depending on the operator selected).

4. Click **Add** to add the filter to the **Active Filters** list.

5. Repeat steps 3 and 4 for each filter you want to create.

6. To remove a filter from the **Active Filters** list, select it in the list, then click **Remove**. To remove all filters from the list, click **Remove All**.

Tip: To change the value for an active filter, remove the filter from the **Active Filters** list. Then create a new filter with the same column but a different value.

7. When you finish specifying filter criteria, click the information tab for a drilldown. The tab lists objects that meet filter criteria.

Database Analysis drilldown Details tab

The Details tab on the Database Analysis drilldown provides the following:

- Identification information for the selected database

- Key statistics that provide an overview of database activities.

Database details and statistics

The following lists the details and statistics that appear on the Details tab for the selected database or database partition. Click a specific detail for its description.

Database Name	Connections Current	I/O Threshold Cleaners Triggered
Database Status	Connection Max Concurrent	OLAP Total Functions
Database Server	Locks Held	OLAP Active Functions
Database Location	Locks Waiting	OLAP Function Overflows

Database Path	Sort Heap Allocated	RTS - Stats Cache Size
Catalog Node Number	Lock List Used (kb)	RTS - Stats Fabrications
Catalog Node Name	Database First Connect	RTS - Synchronous Runstats
Oldest Application Agent	Database Last Backup	RTS - Asynchronous Runstats
Log Secondaries Allocated	Database Statistics Reset	RTS - Fabrication Time (msec)
Log Max Secondaries Used	Log Space Cleaners Triggered	RTS - Sync Runstats Time (msec)
Log Usage(%)	I/O Victim Cleaners Triggered	Threshold Violations

Database Times Graph

The Database Times graph provides a visual overview of the amount of time spent on various processing events for the database or database partition. The times are displayed in seconds.

Direct Write
Direct Read
Buffered Write
Buffered Read
I/O Time
Prefetch Wait
Lock Wait
Sort Time
System CPU
User CPU

Related topics
Database Analysis drilldown

Database Analysis drilldown Caching tab

The Caching tab on the Database Analysis drilldown graphically depicts the utilization of the various resource caches (memory areas) maintained for each database. These caches include package cache, catalog cache, the SQL workarea , and buffer pools–all of which enhance the performance of the database when their respective hit rates are high. Usually the hit rate is the ratio of retrieval requests that were satisfied by data already in the cache to the number of all requests–those satisfied by the cache plus those that required a disk read. (This ratio is expressed as a percentage in the various graphs on this tab.)

If you identify a cache utilization problem for a specific database on this tab, you can use the Caching tab on the Client Application Analysis drilldown to identify those applications that are contributing to the problem.

Graphs

Use the following graphs to follow cache performance over consecutive monitoring intervals:

- Package Cache
- Catalog Cache
- SQL Workarea
- Buffer Pool

Related topics

Client Application Analysis drilldown

Database Analysis drilldown

Database Analysis drilldown Sorting tab

The Sorting tab on the Database Analysis drilldown uses graphs to show sort activity within the selected database and to contrast this activity to overall sorting activity in the instance.

About sorts

An application query often requires that retrieved rows be put in a specific order–that is, be *sorted*–to satisfy a SQL sort request (as with the ORDER BY clause), to construct an index, or to perform a hash join. *Sort heap* is a block of memory that the database manager allocates to an application as workspace whenever the application requires a sort. An important factor in the performance of any sort is the size of its sort heap, which is determined by one of these parameters:

- SORTHEAP database configuration parameter

- SHEAPTHRES database manager configuration parameter

A sort consists of these steps:

- **The actual sort process**–If the information being sorted cannot fit into the sort heap, some information is written to temporary database tables. This incident is called a *sort overflow*. Sorts that do not overflow perform better than those that do.

- **Return of the sorted rows**–If sorted information can be returned directly without requiring a temporary table to store the final, sorted list of data, the sort is referred to as a *piped sort*. If the sorted information requires a table to temporarily store the data before it is returned, the sort is referred to as a *non-piped sort*. A piped sort always performs better than a non-piped sort.

Sorting problems, such as overflows, non-piped sorts, less than optimal sort heap allocation size, might indicate a need to adjust the SORTHEAP and SHEAPTHRES parameters.

Graphs

Use the following graphs to determine sort activity and problems over consecutive monitoring intervals:

Database Sort and Hash Join

Database Sort and Join Overflow Pct

Instance Sort and Join Exceptions

Instance Sort and Join Exception Pct

Related topics

Client Application Analysis drilldown

Top SQL drilldown

Database Analysis drilldown

Database Analysis drilldown I/O Activity tab

The I/O Activity tab on the Database Analysis drilldown provides a summary of I/O activity on the database or partition you are monitoring. The graphs track of rates for the following:

- Reads and writes that involve the buffer pool

- Direct reads and writes, which do not go through the buffer pool, but access the database directly

Graphs

The I/O Activity tab contains the following graphs, which keep track of database I/O rates over consecutive monitoring intervals:

Data Read Rate
Index Read Rate
Index and Data Write Rates
Direct Read and Write Rates

Troubleshooting high I/O activity on the database

High rates depicted on the graphs for the following activities might be indications of performance degradation since these activities involve disk I/O and slow-downs in query processing:

- Physical read rates (compared to logical read rates)

- Synchronous read or write rates

- Direct I/O rates for a database

If you need to investigate high rates for any of these activities, refer to the graphs on the I/O Activity tab in the following drilldowns:

- **Client Application Analysis drilldown**—To identify the applications that are contributing to high I/O activity rates

- **Tablespaces drilldown**—To pinpoint the tablespaces on which the high I/O activity is occurring

- **Buffer Pools drilldown**—To identify those buffer pools that are not being used effectively to offset I/O on the database

Related topics
Client Application Analysis drilldown
Buffer Pool Analysis drilldown
Tablespace Analysis drilldown
Database Analysis drilldown

Database Analysis drilldown Locking tab

The Locking tab on the Database Analysis drilldown gives you an overview of lock activity on the database. Scanning the graphs and statistics on this drilldown, you can quickly detect the following locking problems on

the database:

- High rates of lock holding or waiting

- High numbers of deadlocks, escalations, or timeouts

- Near-or full-capacity utilization of the lock list space.

If lock problems are evident, use the Client Applications drilldown to investigate the lock activity of individual applications to pinpoint those applications that are contributing to the locking issues.

Graphs

Monitor the following graphs to detect abnormal lock activity on the database:

Average Lock Wait Time (msec)

Lock List Utilization (%)

Lock Status

Lock statistics

These lock activity counts provide more detail about any lock problems depicted in the graphs. For each activity type, you can view the count for the lifetime of the statistic (the number of occurrences since the database was activated or performance statistics were reset).

Lock Waits	Lock Escalations
Lock Timeouts	Lock Escalations, Exclusive
Lock Deadlocks Detected	Lock Wait Time, Total (ms)

Additionally, the Statistics tab shows rates based on these counts. The names for these statistics begin with *Lock*.

Related topics

Client Application Analysis drilldown

Database Analysis drilldown

Database Analysis drilldown Log Status tab

The Log Status tab on the Database Analysis drilldown provides current log utilization statistics for the selected database or partition. Use this information to help you pinpoint log utilization problems.

The following statistics describe the current status of the log files on the database. If you detect log utilization problems, such as a log-full condition, you might need to adjust configuration parameter settings for the log files. For information on current parameter settings, go to the Configuration tab for the Databases drilldown.

Log Secondaries Allocated	Log Usage (%)
Log Max Secondaries Used (Kb)	Log Total Transactions
Log Max Total Used (Kb)	Log Avg Log Space/Transaction (Kb)
Log Total Used (Kb)	Node with Least Available Log Space
Log Total Available (Kb)	

Graphs

These graphs provides an overview of log activity and space utilization over consecutive monitoring intervals:

- **Total Transaction Rate** -- Keeps track of the number of commits and rollbacks taking place in the log files per second.

- **Avg Log Space per Transaction** -- Keeps track of the average amount of log file space being used per transaction.

Related topics
Database Analysis drilldown

Database Analysis drilldown Statistics tab

The Statistics tab on the Database Analysis drilldown lists performance statistics for events that have occurred in the selected database. Each row in the list identifies a specific event and provides countsfor the event based on these time measurements:

- **Interval**–Displays the number of occurrences of the event for the most recent monitoring interval.

- **Rate**–Displays the rate per second during the most recent monitoring interval.

- **Lifetime**–Displays the number of occurrences during the life of the event (since the database was activated or performance statistics were reset).

Related topics
Database Analysis drilldown
Database statistics

Database Analysis drilldown Tables tab

The Tables tab in the Database Analysis drilldown shows a list of tables in the database that are currently experiencing I/O activity. The list provides the following information about each table in a grid format. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not appear in your current list view, or be in the order shown here.

Table Schema	Rows Written
Table Name	Overflow Accesses
Table Type	Page Reorgs
Row Accesses	Freeze First Column
Rows Read	

Related topics
Database Analysis drilldown

Database Analysis drilldown Memory Pools tab

The Memory Pools tab on the Database Analysis drilldown helps you keep track of system memory usage as it relates to the database you are monitoring. It is important to remember that, with the notable exception of buffer

pools, most memory is allocated dynamically on an as needed basis by the Database Manager.

The memory pools list provides the following information about each memory pool in a grid format. Since you can choose to hide or display any column in the grid, some columns below might not display in your current list view.

Pool	HWM Util Pct
Curr Size (kb)	Max Size (kb)
Curr Util Pct	Node
Size HWM (kb)	Freeze First Column

Related topics

Memory Pool types

Database Analysis drilldown

Database Analysis drilldown Configuration tab

The Configuration tab on the Database Analysis drilldown contains all the database level parameters available to your DB2 system.

Right-click any column heading to view the configuration parameters you can display on the Configuration tab. A checkmark appears in front of the parameters that are currently displayed. Select or de-select parameters to change the display.

The parameter list grid contains the following information about the parameters on the system:

- **Database Parameter Description**–A brief description of the parameter

- **Database Parameter**–The parameter name

- **Node nn** –The value of the parameter on each partition

Database Analysis drilldown History tab

The History tab on the Database Analysis drilldown shows the command history for all partitions of a database. The information is taken from the database history file. You can check this tab to see when database objects were created, altered, or dropped; when tables were reorganized; when a database was backed up or restored; when tablespaces or tables were quiesced; and more.

When a database row is selected, the tab shows the command history for all partitions. When a partition row is selected, the tab shows the history for that partition.

Note: By default, commands are listed in chronological order, from oldest to most recent. You can reverse that order by clicking the **Start Time** column header.

Columns in the History tab

All columns available for the History tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description

Command Text	The SQL statement for an operation.
Name	The name of the tablespace or table for which an operation was performed.
Operation	The operation that was performed by a SQL statement. One of the following can be shown: • Add table space • Alter tablespace • Archive logs • Backup • Drop tablespace • Load • Load copy • Quiesce • Reorganize table • Rename tablespace • Restore • Rollforward • Unload
Operation Type	The action identifier for an operation. This identifies how an operation was performed. For example, action identifiers for the reorganize operation indicate whether it was online or offline. Action identifiers for the quiesce operation indicate whether it was quiesce share, quiesce update, quiesce, or quiesce reset.
Partition	The partition number of the database partition in which an operation was performed.
Schema	The schema of the table for which an operation was performed.
SQL Code	The SQL return code for an operation. This is taken from the SQLCODE field of the SQLCA (SQL communications area). One of the following can be shown: • **0** (zero) is shown for successful execution • A positive integer is shown for successful execution with a warning condition. • A negative integer is shown for execution with an error condition.
SQL State	The SQL return code that indicates the outcome of the most recently executed operation statement. This is taken from the SQLSTATE field of the SQLCA (SQL

	communications area).
Start Time	The timestamp for the start of an operation. This includes year, month, date, and hour, in that order.

Database Analysis drilldown Utilities tab

The Utilities tab on the Database Analysis drilldown shows the status of running utilities for a database or partition. You can use it to monitor the status and progress of backup, load, rebalance, restore, runstats, crash recovery, and rollforward recovery operations. You can also use it to determine whether a utility is degrading database performance.

Columns in the Utilities tab

All columns available for the Utilities tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Completed Units	The number of work units that have been completed for the current phase. The value is a running count of completed work units. If the value does not increase over a long period of time, this might indicate that a utility is stalled. (A value is not shown for some utilities.) Check the **Work Metric** column to see the type of unit used, for example bytes or extents. Check the **Total Units** column to see the total amount of work for a phase.
Curr Seq	The number of the current phase of a multiphase utility. A number is only shown if a utility contains multiple sequential phases.
ID	The unique identifier for a utility job. An example of the format that might be used for an ID is **userid.jobname** .
Invoker	The invoker type for a utility. One of the following is shown: • **User** is shown if a utility was invoked by a user. • **Auto** is shown if a utility was invoked automatically by DB2.
Phase Description	A description of the work performed during a utility phase. Phase descriptions include **Delete**, **Load**, and **Redo**.
Phase Start Time	The timestamp for the start of a utility phase. This includes year, month, date, and hour, in that order.
Priority	Utility priority specifies the amount of relative importance of a throttled utility with respect to its throttled peers. A priority of 0 implies that a utility is

	executing unthrottled . Non-zero priorities must fall in the range of 1-100, with 100 representing the highest priority and 1 representing the lowest.
Start	The timestamp for the time a utility was invoked. This includes year, month, date, and hour, in that order.
State	The current state of a utility. One of the following is shown: • **Execute** is shown when a utility is executing • **Wait** is shown when a utility is waiting for an event to occur before continuing. • **Error** is shown when an error has occurred.
Total Units	The total number of work units in a phase. A value is not shown for some utilities. It is continuously updated for utilities that cannot quantify the total amount of work. Check the **Work Metric** column to see the type of unit used, for example, bytes or extents.
Type	The class of a utility. One of the following is shown: • Rebalance • Backup • Runstats • Reorg • Restore • Crash recovery • Rollforward recovery • Load • Restart recreate index
Utility Description	A description of the work being performed by a utility. The format used for a description depends on a utility's class. For example, the description for rebalancer might read "Tablespace ID: 3" to indicate that a rebalancer is working on a tablespace with an ID of 3.
Work Metric	The type of unit used to measure completed work and total work for a utility. Possible units include bytes and extents. A value is not shown for some utilities.

Database Analysis drilldown Reorgs tab

The Reorgs tab on the Database Analysis drilldown allows you to monitor table reorganizations for a database. The reorganizations include those that are currently in progress as well as those that are paused, stopped, or

completed. For reorganizations that are in progress, the tab shows how much work has been completed so far and the phase currently being performed. For all reorganizations, the tab shows reorg type, status, start and end times, number of rows compressed, and the names of the table spaces involved.

Columns in the History tab

All columns available for the Reorgs tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Completion	The completion status of a table reorganization: • **Success** is shown if a reorganization completed successfully. • **Fail** is shown if a reorganization did not complete successfully. If a reorganization failed, check the History tab for diagnostic information. This is shown in the **SQL Code** and **SQL State** columns.
Counter	A measure that indicates how much of a reorganization has been or was completed. Check the **Max Counter** column to see the total amount of work for a reorganization.
End	The timestamp for the end of a table reorganization. This includes year, month, date, and hour, in that order.
Index ID	The ID of the index used for ordering table rows during reorganization. If an index was not specified for sorting, **None** is shown.
Long Tbsp ID	The identifier of the tablespace used for reorganizing long objects (those with LONG VARCHAR or LOB data). For partitioned tables, this is the tablespace used for each partition's LONG VARCHAR or LOB.
Max Counter	A measure that indicates the total amount of work for a reorganization. Check the **Counter** column to see the progress of a reorganization that is currently being performed.
Max Phase	The maximum number of phases for a reorganization. This value only applies to classic (offline) reorganizations.
Name	The name of the table that is being reorganized or was reorganized.
Page Reorgs	The number of page reorganizations performed for a table.
Partition	The ID of the data partition for which information is shown in other columns. This value only applies to partitioned tables.
Phase	The phase of a table reorganization currently being performed or last performed. One of the following can be shown:

	- Sort - Build - Replace - Index Recreate This value only applies to classic (offline) table reorganizations.
Phase Start	The timestamp for the start of a certain phase of a table reorganization. This includes year, month, date, and hour, in that order.
Reorg Type	The method and settings specified for a reorganization. Multiple settings can be shown. The settings determine whether a reorganization is in-place (online) or classic (offline). They also determine whether write or read access is allowed during reorganization; whether data is reclustered in the order of a certain index; whether LOB data is reorganized; and whether space is reclaimed.
Rows Compressed	The number of rows compressed in a table during reorganization. If a reorganization is currently in progress, the value shown is a running count of rows that have been compressed. If a reorganization is complete, the value might not be equal to the total number of rows in a table. This is because some rows might have been rejected for compression. For more information, see the description for **Rows Rejected**.
Rows Rejected	The number of rows rejected for compression. Rows are not compressed during reorganization when their record length less than or equal to the minimum record length. If a reorganization is currently in progress, the value shown is a running count of rows that do not meet the compression requirement.
Schema	The schema of the table that is being reorganized or was reorganized.
Start	The timestamp for the start of a table reorganization. This includes year, month, date, and hour, in that order.
Status	The status of a table reorganization. One of the following states is shown: - Started/Resumed - Paused - Stopped - Completed - Truncate This value only applies to in-place (online) table reorganizations.
Tbsp ID	The identifier of the tablespace used for reorganizing a table. For

partitioned tables, this is the tablespace used for each partition.

Database Manager Summary drilldown

The Database Manager Summary drilldown displays global summary information about all active databases and partitions in the DB2 for LUW (Linux, Unix, and Windows) instance you are monitoring. In a single-partition environment, only the **All Nodes** data is presented.

Note: For detailed information about a particular database or database partition, use the Databases drilldown.

When you select a specific database or database partition from the list, a set of tabs appear in the lower half of the window. These tabs contain additional details and activity statistics about the specific database, allowing you to drill down for more information about a problem.

Monitor heap considerations

If monitor heap failure becomes a problem, use this formula, taken from the *IBM System Monitor Guide and Reference,* to adjust the MON_HEAP_SZ parameter.

(*number of monitoring applications* + 1) * (*number of databases* * (800 + (*number of tables accessed* * 20) + ((*number of applications connected* + 1) * (200 + (*number of tablespaces* * 100))))) / 4096

The main drilldown grid

You can click any column header at the top of the grid to sort the database list by the contents of that column. Sorting the list by the various status and interval measurement columns can help identify problem databases. You can also rearrange the order of the columns, and hide columns in the grid to put focus on the information you want to view.

The summary information that the Database Manager Summary drilldown displays is as follows:

Agent Requests	Coord Agents HWM	Piped Sorts Requested
Agents Created due to Empty Pool	Gateway Connections	Post Threshold Joins
Agents from Pool	Gateway Connections HWM	Post Threshold Sorts
Agents Idle	Gateway Conns Switches	Private Memory
Agents Overflows	Gateway Conns Waiting Client	Remote Connections
Agents Registered	Gateway Conns Waiting Host	Remote Conns in Execution
Agents Registered HWM	Local Connections	Service Level
Agents Stolen	Local Conns in Execution	Sort Heap Allocated
Agents Waiting	Partition	Sort Heap HWM
Agents Waiting HWM	Piped Sorts Accepted	Freeze First Column

Sub drilldown tabs

The following tabs appear on the Database Manager Summary drilldown to provide additional information about a selected database or database partition:

Sorting

Database Manager Summary drilldown Agent Activity tab

The Agent Activity tab on the Database Manager Summary drilldown provides a summary of agent activity on each database and partition in the system.

Graphs

The Agent Activity tab contains the following graphs, which keep track of agent activity over consecutive monitoring intervals:

- Local Connections

- Active Agents

- Remote Connections

- Waiting and Idle Agents

Troubleshooting high agent activity

High rates depicted on the graphs for the following activities might be indications of performance degradation and inefficient use of system resources.

In the Active Agents graph, comparing the number of agents from the agent pool to the number of agents created from the empty pool can help you determine if the agent pool is being used efficiently.

A high number of dynamically created agents from the empty pool can indicate two things:

- The NUM_POOLAGENTS parameter should be increased.

- The overall workload for this partition is too high and should be corrected by lowering the number of coordinating agents or reducing the workload.

A low number of dynamically created agents from the empty pool can indicate that NUM_POOLAGENTS is set too high. What this means is there are seldom used agents residing in the agent pool that are unnecessarily using system resources.

Both the local connections and remote connections statistics can help you gauge a need to adjust MAX_COORDAGENTS. If there are more connections than MAX_COORDAGENTS then applications that do not have a coordinating agent are inactive until there is a coordinating agent to service that connection. Because of this, the MAX_COORDAGENTS can be used to fine tune the load on the database server.

The idle agents graph can be used to fine tune the value of the NUM_POOLAGENTS parameter. Having idle agents available to service requests for agents can increase performance of the database server. In environments where there are many applications connected, it is generally recommended to increase the NUM_POOLAGENTS value to avoid the unnecessary overhead of creating and destroying agents.

Related topics
Database Manager Summary drilldown

Database Manager Summary drilldown Configuration tab

The Configuration tab on the Database Manager Summary drilldown contains all the instance level parameters available to your DB2 system.

Parameter list

The parameter list grid contains the following information about the instance level parameters on the system:

- **Parameter Name**–The parameter name as it is defined on your system

- **Parameter Value**–The value of the parameter as it is defined in the instance

- **Database Manager Parameter Description**–A brief description of the parameter

Related topics
Database Manager Summary drilldown

Database Manager Summary drilldown Contacts tab

The Contacts tab on the Database Manager Summary drilldown lists individual contacts and contact groups. Contacts are users that receive email or pager notifications from the DB2 administration server (DAS) after a task completes. Individual contacts are shown in the Contacts pane. Groups are shown in the Contact Groups pane. Both panes can show information that is defined locally or in a global list.

Columns in the Contacts pane

All columns available in the Contacts pane are listed below in alphabetic order. Initially, the pane contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Address	The SMTP mailbox address of a contact that receives notifications via email.
Description	The description of a contact.
Max Page Length	The maximum message length for a contact that receives notifications via a pager. A value is shown if the contact's paging service imposes a message-length restriction.
Name	The name of a contact.
Type	The type of notification used for a contact. One of the following is shown:

	• **Email** is shown if a contact is notified via email.
	• **Page** is shown if a contact is notified via a pager.

Columns in the Contact Groups pane

All columns available in the Contact Groups pane are listed below in alphabetic order. Initially, the pane contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Description	The description of a contact group.
Member Name	The name of a member in a contact group. This can be for an individual contact or another contact group.
Member Type	The type for a member in a contact group. One of the following is shown: • **Contact** is shown for an individual contact. • **Group** is shown for another contact group.
Name	The name of a contact group.

Database Manager Summary drilldown Memory Pools tab

The Memory Pools tab on the Database Manager Summary drilldown helps you keep track of system memory usage as it relates to the instance you are monitoring. It is important to remember that, with the notable exception of buffer pools, most memory is allocated dynamically on an as needed basis by the Database Manager.

The memory pools list provides the following information about each memory pool in a grid format. Since you can choose to hide or display any column in the grid, some columns below might not display in your current list view.

Pool	HWM Util Pct
Curr Size (kb)	Max Size (kb)
Curr Util Pct	Node
Size HWM (kb)	Freeze First Column

Related topics

Memory Pool types

Database Manager Summary drilldown

Database Manager Summary drilldown Paths tab

The Paths tab on the Database Manager Summary drilldown shows the database paths used for database manager files and directories. Information for all database partitions is shown when either a partition or the database is selected.

Columns in the Paths tab

All columns available for the Paths tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Partition	The partition number of a database partition.
Path	The path to a database manager file or directory. Paths to directories end with a file system delimiter. This is '/' on UNIX(R) and '\' on Windows(R).
Path Type	The type of database object that a path is for. Examples of the types shown are LOGPATH , DBPATH , DB_STORAGE_PATH , and LOCAL_DB_DIRECTORY .

Database Manager Summary drilldown Registry tab

The Registry tab on the Database Manager Summary drilldown shows the current DB2 registry settings for all partitions of a database. Information for all partitions is shown when either a partition or the database is selected.

Columns in the Registry tab

All columns available for the Registry tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Aggregate Name	The name of the aggregate registry group a DB2 registry variable belongs to. A name is only shown if a variable is configured through an aggregate registry variable. The column is blank if registry variable is not set from an aggregate or is set from an aggregate but has been overridden.
Level	The level at which a DB2 registry variable is given its value. One of the following levels is shown: • Database partition • Instance

	• Global • Environment
Parameter	The name of the DB2 registry variable.
Partition	The partition number of a database partition a registry variable applies to.
Value	The current setting of a DB2 registry variable.

Database Manager Summary drilldown Sorting tab

The Sorting tab on the Database Manager Summary drilldown uses graphs to show sort activity within the selected database and to contrast this activity to overall sorting activity in the instance.

About sorts

An application query often requires that retrieved rows be put in a specific order–that is, be *sorted*–to satisfy a SQL sort request (as with the ORDER BY clause), to construct an index, or to perform a hash join. *Sort heap* is a block of memory that the database manager allocates to an application as workspace whenever the application requires a sort. An important factor in the performance of any sort is the size of its sort heap, which is determined by one of these parameters:

- SORTHEAP–Database configuration parameter

- SHEAPTHRES–Database manager configuration parameter.

A sort consists of these steps:

- **The actual sort process**–If the information being sorted cannot fit into the sort heap, some information is written to temporary database tables. This incident is called a *sort overflow*. Sorts that do not overflow perform better than those that do.

- **Return of the sorted rows**–If sorted information can be returned directly without requiring a temporary table to store the final, sorted list of data, the sort is referred to as a *piped sort*. If the sorted information requires a table to temporarily store the data before it is returned, the sort is referred to as a *non-piped sort*. A piped sort always performs better than a non-piped sort.

Sorting problems, such as overflows, non-piped sorts, less than optimal sort heap allocation size, might indicate a need to adjust the SORTHEAP and SHEAPTHRES parameters.

Graphs

Use the following graph to determine sort activity and problems over consecutive monitoring intervals:

Sort and Hash Join Exceptions

Related topics

Database Manager Summary drilldown
Client Application Analysis drilldown
Top SQL drilldown
Database Analysis drilldown

Database Manager Summary drilldown Statistics tab

The Statistics tab on the Database Manager Summary drilldown lists performance statistics for events that have occurred in the selected database. Each row in the list identifies a specific event and provides counts for the event based on these time measurements:

- **Interval**–Displays the number of occurrences of the event for the most recent monitoring interval.

- **Interval Rate/Sec**–Displays the rate per second during the most recent monitoring interval.

- **Lifetime**–Displays the number of occurrences during the life of the event (since the database was activated or performance statistics were reset).

Related topics

Database Manager Summary drilldown
Database Manager Summary statistics

Database Manager Summary drilldown System Resource tab

The System Resource tab on the Database Manager Summary drilldown shows system resource information for all partitions of a database.

Review the following for additional information:

Column	Description
Partition	Partitions in the database.
Resource Attribute	Attribute assigned to resource.
Attribute Value	Value you assign to the resource attribute

Related topics

Database Manager Summary drilldown

Diagnostic Log drilldown

The Diagnostic Log drilldown captures diagnostic log, event, and notification messages (also called *records*) generated for the current database or for databases within the current DB2 instance. Because a database can generate a large volume of records, you can define filters that limit the number of records that Spotlight retrieves. Once Spotlight retrieves the records, you can manage the scope of the record list by defining filters on column values.

The Diagnostic Log drilldown is based on the PD_GET_DIAG_HIST table function.

Click one of these links for more information about this drilldown:

Open the Diagnostic Log drilldown

Use the Diagnostic Log drilldown

Limit the number of records retrieved

Filter the information displayed on the drilldown

Diagnostic information available on the drilldown

Open the Diagnostic Log drilldown

To open the Diagnostic Log drilldown

On the home page for the selected DB2 database or on the home page for the DB2 instance, click ![icon] in the Spotlight toolbar.

Use the Diagnostic Log drilldown

The steps below outline the basic workflow involved in using the Diagnostic Log drilldown.

To use the Diagnostic Log drilldown

1. On the Diagnostic Log tab of the drilldown, define predicates that limit the number of records that Spotlight retrieves.

2. Once Spotlight retrieves the records and displays them in the lower pane of the Diagnostic Log tab, use the Criteria tab to refine the list of records by defining filters based on column values.

3. On the Diagnostic Log tab, browse details about the diagnostic records captured on the drilldown. For a description of the information shown for each record, see Diagnostic Log information available on the drilldown.

Limit the number of records retrieved

Databases can generate a high volume of diagnostic records at a given time. To limit the number of diagnostic records that the PD_GET_DIAG_HIST table function retrieves, use the procedure below to define predicates for the function. These predicates control the type of diagnostic information retrieved and help limit the amount of time required for retrieval.

To define predicates

1. In the **Diagnostic Log Fetching Predicates** pane on the Diagnostics Log tab, complete the one or more of the following fields to define predicates that filter the diagnostic records:

Facility	Select the option that determines the general record grouping to retrieve: • **ALL** - Retrieve records containing optimizer statistics and records from all DB2 general diagnostic logs and

	• **MAIN** - Retrieve records from all DB2 general diagnostic logs only--db2daig.log, admin notification log, and event logs. • **OPTSTATS** - Retrieve records containing optimizer statistics only.
Event types	Select one or both types of event records to retrieve: • **Internal** - All internal event records. IBM support uses these records to troubleshoot issues. • **External** - All external event records. These records are directed to users.
Diagnostic types	Select one or both types of diagnostic records to retrieve: • **Internal** - All internal diagnostic records. IBM support uses these records to troubleshoot issues. • **External** - All diagnostic notification records, used to provide information to the user.
Minimum impact	To retrieve records containing messages that have a degree of impact on the user, specify the minimum level of this impact. For example, select *POTENTIAL* to retrieve records that contain messages of *POTENTIAL*, *IMMEDIATE*, or *CRITICAL* value. Select *NONE* to retrieve all records.
From, To	Select the **To** or **From** fields or both to specify a finite range of time in which to capture log information. Then enter the time range in which to capture log information. All records generated after the **To** timestamp and before the **From** timestamp are captured. To change the date or time, use the arrows to move the times backward or forward; or select and type over any element in the date or time entry.

2. Click **Fetch** to retrieve the records and display them in lower pane of the Diagnostic Log tab.

Filter information displayed on the drilldown

Once Spotlight retrieves the records, you can filter the list of records as needed.

To filter the list of diagnostic records

1. On the Criteria tab, define the filters based on the column values of the retrieved rows.

2. Return to the Diagnostic Log tab, and click **Fetch**. The list of records is refreshed, based on the current Criteria tab filters.

3. Repeat steps 1 and 2 to define or remove filters to continue to adjust the record display, depending on the information you want to browse.

Diagnostic information available on the drilldown

The following information can display for each record on the Diagnostic Log drilldown. For a complete description of the record information, refer to the IBM DB2 for Linux, Unix, and Windows documentation describing the PD_GET_DIAG_HIST table function.

Initially, by default, the drilldown shows a subset of all columns. available for this drilldown However, you can right-click in the column header area to perform a variety of column-customization activities. For example, you can display hidden columns and hide currently displayed columns as needed. Additionally, you can reorganize the displayed columns. For more information, see the Organize Columns topic.

Column	Description
TimeStamp	The time that the diagnostic message was generated.
DBName	The database for which the diagnostic message was generated.
Facility	The general group of records from which Spotlight retrieved the diagnostic message: • **ALL** - Records retrieved from all DB2 general diagnostic logs and records containing information about optimizer statistics. • **MAIN** - Records retrieved from all DB2 general diagnostic logs only, including db2diag.log, admin notification log, and event logs. • **OPSTATS** - Records containing optimizer statistic information
Type	The type of record: • DI -Internal diagnostic record. IBM support uses this record type to troubleshoot issues. • DX - External diagnostic record, intended to provide user information. • EI - Internal event record. IBM support uses this record type to troubleshoot issues. • EX - External event record, intended to provide user information.
Partition	The number of the partition on which the record was generated.

Level	The severity level of message: C -Critical E - Error I - Information S - Severe W - Warning
Impact	From the user's point of view, the degree to which the condition denoted by the message affects the DB2 business process.
EDU ID	The ID of the engine dispatched unit that generated the message.
EDU Name	The name of the engine dispatched unit that generated the message.
PID	ID of the operating system process that generated the message.
Process Name	The name of the operating system process that generated the message.
Thread ID	The ID of the thread that generated the message.
Appl Name	The name of the application that connected to the database.
Appl ID	The ID of the application that connected to the database.
Appl Handle	The agent ID assigned to the application that connected to the database.
Auth ID	The system authorization ID for the process.
Product	The name of the product that generated the record.
Component	The name of the component that generated the record.
Function	The name of the function that generated the record.
Probe	The number of the probe point at which the function generated the message.
Called Prod	The name of the product that is the error source (if the error source is not the product the generated the message).
Called Comp	The name of the component that is the error source (if the error source is not the component that the generated the message).
Called Func	The name of the function that is the error source (if the error source is not the function that the generated the message).
OS Error	The number identifying the operation system error.
Ret Code	The code returned by the product.
Msg Num	The number identifying the message.
Msg Type	The type of message associated with the message number.
Message	An abbreviated version of the message text.
Obj Type	The type of object to which the event applies.
Obj Name	The name of the object to which the event applies.
Obj Qual	Qualifying information for the object to which the event applies.
Event Type	The action associated with the event that caused the message.
Event Desc	An abbreviated list of the key fields associated with the event.
1st Event Qual Type	The type of the event's first qualifier. (A qualifier helps to describe what

	the event affects.) If **Facility** is *OPSTATS* , this value must be *AT*.
1st Event Qual	The event's first qualifier.
2nd Event Qual Type	The type of the event's second qualifier. (A qualifier helps to describe what the event affects.) If **Facility** is *OPSTATS* , this value must be *BY*.
2nd Event Qual	The event's second qualifier.
3rd Event Qual Type	The type of the event's second qualifier. (A qualifier helps to describe what the event affects.) If **Facility** is *OPSTATS* , this value must be *DUE TO*.
3rd Event Qual	The event's third qualifier.
Event State	The result of the event on the object or action. A percentage value shows the progress that the effect of the event has had on the object or action.
Event Attr	One or more attributes relating to the event.
Event Stack	The event stack that was current when the message was added to the log.
Call Stack	The operating system stack dump that was current when the thread generated the message.
Dump File	The full path and file name of the secondary dump file connected to the logged record.

Criteria Tab for Drilldowns

As you view drilldowns, you can filter them based on criteria specified in the Criteria tab. This is available for the following drilldowns: Database Analysis, Tablespace Analysis, Buffer Pool Analysis, Client Application Analysis, and Diagnostic Log.

Filtering lets you narrow the scope of the list shown in a drilldown. You can use it to focus on objects with a certain level of activity, certain attributes, or a certain current state. For example, you might filter a drilldown to see databases with the greatest I/O, tablespaces with the most direct writes, buffer pools with the lowest hit ratio, or applications that are currently connected.

Criteria for a drilldown can include multiple filters. The more filters you create, the narrower the scope of the drilldown list. Only items that meet filter criteria are shown. All others are filtered out. If no objects meet filter criteria, "No Data" is shown in the drilldown.

Filtering applies to the current database and current session. It is not saved when the session ends.

Note: When filtering is in use, an asterisk is shown beside the name of a drilldown's list tab. For example, when filtering is in use for the Tablespace Analysis drilldown, **Tablespaces*** is shown at the top of the Tablespaces tab.

Using the Criteria tab

Follow the procedure below to specify filter criteria. Before you create or modify filters, review the columns shown in the drilldown. The columns you can select for a filter are the ones currently displayed. If you want to filter on columns that are not currently displayed, select those columns in the **Column Organizer**. To open this window, right click any column header and select **Organize Columns…** from the right-click menu.

Note: The first filter you create determines which values are shown in the **Value** list box for the next filter. For example, if your first filter specifies "Content = Large" and you select the **Extent Size** column for the next filter,

the **Value** list box shows extent sizes for tablespaces with the Large content type. If you want to use an extent size that is not listed, enter it in the **Value** list box.

To specify filter criteria

1. Click the Criteria tab for a drilldown.

2. If you plan to use multiple filters in criteria, select an option from the **Result Set will include** list box as follows:

Any	Select this option if the result set should include items that meet criteria for one or more filters. Items are included even if they do not meet criteria for all filters.
All	Select this option if the result set should include items that meet criteria for all filters. Items are included only if they meet criteria for all filters.

3. Specify criteria for a filter as follows:

a.	**Column**	From this list box, select a column for the filter. Objects are checked against values in this column. The list includes all columns currently displayed in the drilldown.
b.	**Operator**	From this list box, select a comparison operator for column values. The choices are **=** (equal to), **<>** (less than and greater than), **>** (greater than), and **<** (less than). Use the **=** (equal to)operator for a column that shows an attribute such as status, name, or ID.
c .	**Value**	From this list box, select the column value for the filter. In most cases, the values shown are for entire objects. Values for partitions are not listed. If a value for an entire object is not shown, you can enter it manually. Objects are checked to see if their values are equal to, less than and/or greater than this value (depending on the operator selected).

4. Click **Add** to add the filter to the **Active Filters** list.

5. Repeat steps 3 and 4 for each filter you want to create.

6. To remove a filter from the **Active Filters** list, select it in the list, then click **Remove**. To remove all filters from the list, click **Remove All**.

Tip: To change the value for an active filter, remove the filter from the **Active Filters** list. Then create a new filter with the same column but a different value.

7. When you finish specifying filter criteria, click the information tab for a drilldown. The tab lists objects that meet filter criteria.

FCM Analysis drilldown

The Fast Communication Manager (FCM) Analysis drilldown provides a breakdown of message buffer traffic for the instance you are monitoring. FCM provides communication support between partitions for agent requests, and message buffers in DB2 environments.

Note: Setting `DB2_FORCE_FCM_BP` to `YES` in an MLN environment lets DB2 create FCM buffers in a separate memory segment. Once this is done, communications between FCM agents belonging to different logical partitions (on the same physical partition) occurs through shared memory.

This drilldown consists of two tabs:

- **Internode Summary tab**–Breaks down the traffic between source partition and target partition and displays summary totals for data traffic.

- **Resources tab**–Provides low water mark information for different types of traffic and displays utilization statistics.

Internode Summary tab contents

The Internode Summary tab provides the following columns. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not display, or be in this order, in your current list view.

Source Node	Buffers Total
Target Node	Percent of Traffic
Buffers Sent	Connection Status
Buffers Received	Freeze First Column

Resource tab contents

The Resource tab provides the following columns. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not display, or be in this order, in your current list view.

Node	Conn Entries
Buffers	Conn Entry LWM
Buffer LWM	RQ Blocks
Msg Anchors	RQ Block LWM
Msg Anchor LWM	Freeze First Column

Intra-partition parallelism

Intra-partition parallelism is the name for the breaking up of a query into several parts such as index creation, database loading, or SQL queries, operations usually considered as parts of a single task. When intra-partition parallelism is turned on, FCM helps these tasks run in parallel within a single database partition.

Note: In some environments, the `intra_parallel` parameter must be active before it can be used.

Tablespace Analysis drilldown

The Tablespace Analysis drilldown displays key performance statistics for the tablespaces in monitored databases and partitions. At its top level, the drilldown provides an overview of performance during the current monitoring session. This is shown in the Tablespaces tab.

From the Tablespaces tab, you can drill down to detailed information on I/O activity, I/O times, ranges, quiescers, and more for a certain tablespace. This is shown in sub drilldown tabs, which are displayed in the lower half of the window when a tablespace is selected. The additional statistics can be used to diagnose any problems detected in the Tablespaces tab.

Using the Tablespace Analysis drilldown

The steps below outline the basic workflow involved in using the Tablespace Analysis drilldown.

To use the Tablespace Analysis drilldown

1. View top level tabs to detect performance problems in tablespaces. Work from the Tablespaces tab to check statistics. Work from Criteria tab to filter the list in the Tablespaces tab based on criteria such as level of I/O activity or amount of unusable space.

2. View rows in the Tablespaces tab to focus on a certain tablespace. When a database is being monitored, rows show information for a tablespace across all partitions and on individual partitions. Check columns in the Tablespaces tab to see statistics for a tablespace.

3. View sub drilldown tabs to see additional statistics for a certain tablespace. Click the tablespace row to display the tabs. The information they contain can help you diagnose problems detected in the Tablespaces tab.

4. Open a Toad for DB2 session to alter a tablespace or view its properties, storage exceptions, or historical and hypothetical growth requirements. (You must have either product licensed and installed on the same computer where the Spotlight for DB2 client resides.)

Opening a Toad for DB2 session

If you have Toad for DB2 installed, you can alter a tablespace, view it properties, or determine which tables and indexes in the tablespace need to be reorganized.

Right-click the tablespace entry in the list of tablespaces, and proceed as follows:

To alter a tablespace

- Select **Toad** ▸ **Alter** to open the Alter window, enabling you to change the tablespace properties.

To view tablespace properties

- Select **Toad** ▸ **Properties** to open the Toad database explorer on the tablespace.

To determine which tables and indexes in the tablespace need to be reorganized

- Select **Toad** ▸ **Storage** ▸ **Identify Exceptions** to open the Space Exceptions window, enabling you to identify those tables and indexes in the tablespace that need to be reorganized or for which statistics need to be collected.

On the resulting window for any of these functions, press F1 for instructions to proceed with the function.

Related Topics

Top Level tabs

Sub drilldown tabs

Rows in the Tablespaces tab

Columns in the Tablespaces tab

Criteria Tab for Drilldowns

As you view drilldowns, you can filter them based on criteria specified in the Criteria tab. This is available for the following drilldowns: Database Analysis, Tablespace Analysis, Buffer Pool Analysis, Client Application Analysis, and Diagnostic Log.

Filtering lets you narrow the scope of the list shown in a drilldown. You can use it to focus on objects with a certain level of activity, certain attributes, or a certain current state. For example, you might filter a drilldown to see databases with the greatest I/O, tablespaces with the most direct writes, buffer pools with the lowest hit ratio, or applications that are currently connected.

Criteria for a drilldown can include multiple filters. The more filters you create, the narrower the scope of the drilldown list. Only items that meet filter criteria are shown. All others are filtered out. If no objects meet filter criteria, "No Data" is shown in the drilldown.

Filtering applies to the current database and current session. It is not saved when the session ends.

Note: When filtering is in use, an asterisk is shown beside the name of a drilldown's list tab. For example, when filtering is in use for the Tablespace Analysis drilldown, **Tablespaces*** is shown at the top of the Tablespaces tab.

Using the Criteria tab

Follow the procedure below to specify filter criteria. Before you create or modify filters, review the columns shown in the drilldown. The columns you can select for a filter are the ones currently displayed. If you want to filter on columns that are not currently displayed, select those columns in the **Column Organizer**. To open this window, right click any column header and select **Organize Columns…** from the right-click menu.

Note: The first filter you create determines which values are shown in the **Value** list box for the next filter. For example, if your first filter specifies "Content = Large" and you select the **Extent Size** column for the next filter, the **Value** list box shows extent sizes for tablespaces with the Large content type. If you want to use an extent size that is not listed, enter it in the **Value** list box.

To specify filter criteria

1. Click the Criteria tab for a drilldown.

2. If you plan to use multiple filters in criteria, select an option from the **Result Set will include** list box as follows:

Any	Select this option if the result set should include items that meet criteria for one or more filters. Items are included even if they do not meet criteria for all filters.

All	Select this option if the result set should include items that meet criteria for all filters. Items are included only if they meet criteria for all filters.

3. Specify criteria for a filter as follows:

a.	**Column**	From this list box, select a column for the filter. Objects are checked against values in this column. The list includes all columns currently displayed in the drilldown.
b.	**Operator**	From this list box, select a comparison operator for column values. The choices are **=** (equal to), **<>** (less than and greater than), **>** (greater than), and **<** (less than). Use the **=** (equal to)operator for a column that shows an attribute such as status, name, or ID.
c .	**Value**	From this list box, select the column value for the filter. In most cases, the values shown are for entire objects. Values for partitions are not listed. If a value for an entire object is not shown, you can enter it manually. Objects are checked to see if their values are equal to, less than and/or greater than this value (depending on the operator selected).

4. Click **Add** to add the filter to the **Active Filters** list.

5. Repeat steps 3 and 4 for each filter you want to create.

6. To remove a filter from the **Active Filters** list, select it in the list, then click **Remove**. To remove all filters from the list, click **Remove All**.

Tip: To change the value for an active filter, remove the filter from the **Active Filters** list. Then create a new filter with the same column but a different value.

7. When you finish specifying filter criteria, click the information tab for a drilldown. The tab lists objects that meet filter criteria.

Tablespace Analysis drilldown Details tab

The Details sub drilldown provides additional information for the tablespace selected in the Tablespace Analysis drilldown. The sub drilldown contains two panes:

- **Tablespace Details pane**—This pane shows information such as configuration settings, extent size, and I/O times for a tablespace. When an aggregate tablespace row is selected, the information is for a tablespace as it exists across all database partitions in which it is defined. When an individual tablespace row is selected, the information is for the tablespace as it exists on a certain partition.

- **Tablespace Partition Information pane**—This pane shows information such as size, total number of pages, number of used pages, and number of quiescers. The information is for a tablespace as it exists on a certain database partition. When an aggregate table space row is selected, a "Data only available for partitions" message is shown in this pane.

Related topics

Tablespace Analysis drilldown Containers tab

The Containers tab on the Tablespaces drilldown provides information about the physical storage devices on which the selected tablespace is defined.

Container list contents

The container list provides the following information for each container in a grid format. Since you can choose to hide, display, and rearrange any column in the grid, some columns below might not appear in your current list view, or be in the order they are shown here:

Column	Description
Name	Name of the device, directory, or file that is the container.
Partition	Partition number of the database partition on which the tablespace or buffer pool is defined.
ID	Internal identifier of a container.
Accessible	Identifies accessible containers.
Type	Specifies how the container is managed: • SMS directory path –System-managed space • DMS file spec –Database-managed space
Size	Size of a container in megabytes.
Pages	Size of a container in pages.
Useable	Percentage of the container space that the database manager can use to store data. **Note:** Each device can have only one container defined. To avoid wasted space, you should make the size of the device and the size of the container equivalent. For example, if the device is allocated with 5000 pages and the device container is defined to allocate 3000 pages, then 2000 pages on the device are not useable.
File System ID	Unique identifier for a file system.
File System Size	Total size of a file system
File System Used	Amount of space used on a file system.
Stripe Set	Number of the stripe set to which a container belongs.

Related topics

Tablespace Analysis drilldown I/O Activity tab

The I/O Activity tab on the Tablespace Analysis drilldown provides a summary of I/O activity on the selected tablespace. The graphs show rates for the following:

- Reads and writes that use the buffer pool

- Direct reads and writes, which do not go through the buffer pool, but access the tablespace directly

The bottom of the drilldown lists average I/O times (in milliseconds) for both buffered and direct reads and writes. The average times are shown in milliseconds.

Troubleshooting high I/O activity

On the Database Analysis drilldown, you might notice consistently high rates for the following activities on specific databases:

- Physical read rates (compared to logical read rates)

- Synchronous read or write rates

- Direct I/O rates for a database

These high rates might be indications of performance degradation since the activities involve disk I/O and slow-downs in query processing. Use the graphs listed above to pinpoint those tablespaces on which the high I/O activity is occurring.

Graphs

The following graphs on the I/O Activity tab keep track of the rates for I/O activity on the selected tablespace. These rates are plotted over consecutive monitoring intervals.

Data Read Rate

Index Read Rate

Index and Data Write Rate

Direct Read and Write Rate

Related topics

Tablespace Analysis drilldown I/O Times tab

The I/O Times tab on the Tablespace Analysis drilldown contains graphs that let you visually monitor average I/O times and buffer pool hit ratios for a tablespace. The four graphs track the following:

- Average amount of time spent on direct I/O.

- Average amount of time spent on asynchronous I/O.

- Average amount of time spent on physical I/O.

- Buffer pool hit ratios.

All four graphs use counter data to show statistics as they change over time. This allows you to see when average I/O times are high and buffer pool hit ratios are low for a tablespace.

Consistently high I/O times can be a symptom of I/O conflict. Consistently low hit ratios can be a sign that a buffer pool is not operating efficiently. Both problems degrade performance. To identify applications that are contributing to I/O conflict, work from the Client Application Analysis drilldown. To investigate buffer pool conditions, work from the Buffer Pool Analysis drilldown.

Graphs in the I/O Times tab

Click the links below to see descriptions of each graph in the I/O Times tab:

Average Direct I/O Times

Average Async I/O Times

Average Physical I/O Times

Read BP Hit Ratios

Functions for graphs

As you view graphs in the I/O Times tab, you can click icons in graph toolbars to:

Copy graph details to the clipboard

Change the graph view

Filter a graph

Maximize or restore a graph

Pan over a zoomed graph

Zoom a graph

Related topics

Tablespace Analysis drilldown

Tablespace Analysis drilldown Quiescers tab

The Quiescers tab on the Tablespaces drilldown provides information about quiescer activity that is currently in progress. It shows the type of quiesce being performed. It also shows information that can help identify who is responsible for quiescing a tablespace. This information is shown in the **Agent ID**, **Authorization ID**, **OBID**, **Tablespace ID** columns.

Columns in the Quiescers tab

All columns available for the Quiescers tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Agent ID	The ID of the agent holding a quiesce state.
Authorization ID	The authorization ID of the user holding a quiesce state. This is the ID that was used to connect to the database.
OBID	The object ID of the object that caused a tablespace to be quiesced.
State	The quiescer state. This indicates the type of quiesce being performed. For example, the type might be SHARE, INTENT TO UPDATE, or EXCLUSIVE.
Tablespace ID	The tablespace ID of the object that caused a table space to be quiesced.

Tablespace Analysis drilldown Ranges tab

The Ranges tab on the Tablespaces drilldown shows how ranges are defined for range partitioning in DMS tablespaces. When an aggregate tablespace row is selected, the tab shows partitioning information for a tablespace across all database partitions. When the row for an individual tablespace is selected, the tab shows partitioning information for the partition on which the tablespace is defied.

Note: Information in the Ranges tab applies only to DMS tablespaces. It does not apply to SMS tablespaces.

Columns in the Ranges tab

All columns available for the Ranges tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Adjustment	Any offset into the container array in which a range starts.
Container ID	The ID that uniquely identifies a container within a range.
Containers	The number of containers in a range.
End Stripe	The number of the last stripe in a range.
Max Extent	The maximum extent number that is mapped by a range.
Max Page	The maximum page number that is mapped by a range.
Number	The number of a range in the tablespace map.
Offset	The offset from stripe 0 of the beginning of the stripe set to which a range belongs.

Start Stripe	The number of the first stripe in a range.
Stripe Set	The number of the stripe set to which a range belongs.

Tablespace Analysis drilldown Statistics tab

The Statistics tab on the Tablespaces drilldown lists performance statistics for database events associated with the selected tablespace. Each row in the list identifies a specific event and provides counts for the event based on these time measurements:

- **Interval**–Displays the number of occurrences of the event for the most recent monitoring interval.

- **Rate**–Displays the rate per second during the most recent monitoring interval.

- **Lifetime**–Displays the number of occurrences during the life of the event (since the database was activated or performance statistics were reset).

Related topics

Tablespace Analysis drilldown
Tablespace statistics

Tablespace Analysis drilldown Tables tab

The Tables tab in the Tablespaces drilldown shows a list of all tables in the selected tablespace. The list provides the following information about each table in a grid format. Since you can choose to hide, display, or rearrange any column in the grid, some columns below might not appear in your current list view, or be in the order shown here.

Column	Description
Schema	Schema name of a table. Only catalog and user tables use schema names.
Name	Name of the table. Only catalog and user tables have table names.
Rows Read	The number of rows read from this table since Spotlight on IBM® DB2® LUWconnected to the DB2 for LUW database or since statistics were reset. If this value is high, consider creating additional indexes on the table.
Rows Written	The total number of rows that were changed (inserted, deleted, or update) in the table since Spotlight connected to the database or since statistics were reset. If this value is high, consider executing the DB2 RUNSTATS utility on this table to keep its statistics up to date for efficient SQL access to the table data.
Index Pages	Name of the tablespace that holds all indexes created on the table.
Overflow Accesses	The number of times that accesses to overflowed rows on this table were required since Spotlight connected to the database or since statistics were reset. An overflowed row is an updated row that no longer fits on the data page where the row was originally written. The overflow is usually the result

	of one of these events:
	• An update to a VARCHAR value
	• An ALTER TABLE statement that adds columns to the table
	Overflowed rows indicate that data fragmentation has occurred. If the number overflowed row accesses for the table is high, consider reorganizing the table using the DB2 REORG utility to clean up the fragmentation.
Page Reorgs	Total number of page reorganizations for a table.
Data Pages	Total number of pages used by data in this table.
Lob Pages	Total number of pages used by LOBS related to this table.
Long Pages	Total number of pages used by long data in this table.
XML Pages	Total number of XML pages from tablespace containers for regular and large spaces.
Data Partition ID	Identifies the data partition for which information is returned.
Type	Identities the type of table for which information is returned.

Related topics
Tablespace Analysis drilldown

Top SQL drilldown

The Top SQL drilldown shows performance metrics for the SQL statements executed in an active database or partition. It is designed to reveal statements that are the "top SQL" in terms of being the most resource intensive and high cost. For example, you can use it to identify statements that consume the most CPU, statements with the lowest hit ratio, or statements with the most sort overflows.

Metrics are shown in tabs for current, dynamic, and static SQL statements. From each top-level tab, you can drill down to detailed information on problematic statements. This is provided in sub drilldowns, which are displayed in the lower half of the window when a statement is selected. The additional information can help you diagnose why a statement is performing poorly.

Using the Top SQL drilldown

The steps below outline the basic workflow involved in using the Top SQL drilldown.

To use the top SQL drilldown

1. View top-level tabs to detect poorly performing SQL statements problems in client applications. The tabs and the SQL they show are as follows:

- Current SQL tab—This tab shows metrics for statements that are currently executing in a database or partition. The statements are those issued by applications that currently connected to the database or partition.
- Dynamic SQL tab—This tab shows metrics for dynamic SQL statements that are currently executing or have executed in the past for a database or partition. Past statements are those that have been executed since been executed since a database was started. The statements are those issued by applications that are currently connected or that have been connected at some point since the database was started.
- Static SQL tab—This tab shows metrics for static SQL statements that are currently executing or have been executed in the past for a database or partition. Past statements are those that have been executed since the current session was started. The statements are those issued by applications that currently connected to the database or partition.

   Tip: You can use various display features to focus on the most problematic statements. You can organize columns from left to right, hide or show columns, and sort the tab by values in a certain column. To sort by a certain column, click the column header. Being able to sequence and sort columns in different ways allows you to quickly identify different performance problems.

2. When you detect a problematic statement in a top-level tab, view its sub drilldown tab to see additional information and view the statement. Click the row for the statement to display the sub drilldown. The additional information can help you diagnose why a statement is performing poorly. (To close the sub drilldown, click the **X** button in the top right corner of the drilldown.)

3. After you detect and diagnose problematic statements, you can tune them using SQL Optimizer for DB2. (You must have either product licensed and installed on the on the same computer where the Spotlight for DB2 client resides.) SQL tuning sessions can be opened from the sub drilldown for each tab.

Related topics
Tuning an SQL statement
Printing the Top SQL drilldown

Top SQL drilldown Current SQL tab

The Current SQL tab on the Top SQL drilldown shows metrics for SQL statements that are currently executing against a monitored database or partition. The statements can be either dynamic or static.

Note: Spotlight captures information for statements that are in flight when it is refreshing. If a statement is not in flight during a refresh, information is not captured. This means a statement might not be shown or that values might not reflect all executions of a statement.

From this tab, you can drill down for more details on problematic statements and tune them.

Drill down on a statement

When the Current SQL tab shows a problematic statement, you can drill down to view additional performance information specific to that statement.

To drill down on a specific statement

Click the statement in the list of currently executing statements.

The Unit of Work sub drilldown is displayed at the bottom of the Current SQL tab. This sub drilldown shows the selected statement and provides additional metrics for this statement. Use these metrics to diagnose why the

statement is performing poorly. Then, from within this sub drilldown, you can open a tuning session on the statement.

Tune a statement

From the Unit of Work sub drilldown, you can open a tuning session on the currently displayed statement if you have either of the following products licensed and installed on the same computer where your Spotlight on IBM® DB2® LUW client is installed:

- Spotlight on IBM DB2 LUW

To tune the SQL statement currently displayed in the Unit of Work sub-drilldown

1. Click ![icon] on the left side of the Unit of Work sub drilldown.

SQL Optimizer for DB2 opens, showing the SQL statement that you were examining in Spotlight in the product's main window.

2. If you need instructions to proceed with the tuning session, click F1 for context-sensitive help.

Columns in the Current SQL tab

All columns available for the Current SQL tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Agent ID	The application handle of the agent for the application that issued an SQL statement.
Database	The name of the database on which an SQL statement is executing.
Data Hit Ratio	A statement's ratio of physical reads to logical reads for data pages.
Elapsed Time (microsec)	The amount of execution time elapsed so far for a statement in microseconds. This changes when the interface is refreshed.
Index Hit Ratio	A statement's ratio of physical reads to logical reads for index pages.
Operation	The operation currently being processed for an SQL statement. One of the following can be shown: • CALL • CLOSE • COMPILE • DESCRIBE • EXECUTE

	- EXECUTE IMMEDIATE - FETCH - FREE LOCATOR - OPEN - PREPARE - PREP_COMMIT - PREP_EXEC - PREP_EXEC - SELECT - STATIC COMMIT - STATIC ROLLBACK
Temp Data Hit Ratio	A statement's ratio of physical reads to logical reads for data pages read from temporary tablespaces.
Temp Index Hit Ratio	A statement's ratio of physical reads to logical reads for index pages read from temporary tablespaces.
Temp XML Hit Ratio	A statement's ratio of physical reads to logical reads for XML pages read from temporary tablespaces.
Text	The text of an SQL statement. For dynamic SQL statements, this is the SQL text associated with a package.
Type	The type for an SQL statement. Either DYNAMIC or STATIC is shown.
XML Hit Ratio	A statement's ratio of physical reads to logical reads for XML pages.

Related Topics

Top SQL drilldown
Dynamic SQL drilldown
Static SQL drilldown
Unit of Work sub drilldown

Top SQL drilldown Dynamic SQL tab

The Dynamic SQL tab on the Top SQL drilldown shows performance metrics for the dynamic SQL statements run against a monitored database or partition. The statements include those that are currently executing and those that were executed earlier in the current session. They can also include statements that were executed

prior to the current session. Up to 200 statements that have been executed since a database was started can be shown. Past statements are shown even if the issuing application is not connected.

Note: Dynamic SQL statements can be issued interactively by users and built at run time. Unlike static SQL statements, they are not hard coded into applications.

From this tab, you can drill down for more details on problematic statements and tune them.

Drill down on a statement

When the Dynamic SQL tab shows a problematic statement, you can drill down to view additional performance information specific to that statement.

To drill down on a specific statement

Click the statement in the list of dynamic statements.

The Dynamic SQL Details sub drilldown is displayed at the bottom of the Dynamic SQL tab. This sub drilldown shows the selected statement and provides additional metrics for this statement. Use these metrics to diagnose why the statement is performing poorly. Then, from within this sub drilldown, you can open a tuning session on the statement.

Tune a statement

From the Dynamic SQL Details sub drilldown, you can open a tuning session on the currently displayed statement if you have either of the following products licensed and installed on the same computer where your Spotlight on IBM DB2 LUW client is installed:

- Spotlight on IBM DB2 LUW

To tune the SQL statement currently displayed in the Dynamic SQL Details sub-drilldown

1. Click on the left side of the Dynamic SQL Details sub drilldown.

SQL Optimizer for DB2 opens, showing the SQL statement that you were examining in Spotlight in the product's main window.

2. If you need instructions to proceed with the tuning session, click F1 for context-sensitive help.

Columns in the Dynamic SQL tab

All columns available for the Dynamic SQL tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Average Execution Time (microsec)	The average execution time for a SQL statement. This value is calculated from the execution time and executions for a statement. A high average time indicates a statement should be tuned.
Compilations	The number of different compilations for a SQL statement.

Data Hit Ratio	A statement's ratio of physical reads to logical reads for data pages.
Execution Time (microsec)	The total time spent executing an SQL statement. The value is shown in microseconds. If a statement is currently executing, the value is the amount of execution time elapsed so far.
Executions	The number of times an SQL statement has been executed. This value can help you identify the most frequently executed statements in your system.
Index Hit Ratio	A statement's ratio of physical reads to logical reads for index pages.
Prep Best (microsec)	The best preparation time for an SQL statement. This is the least amount of time needed to prepare the statement. The value is shown in microseconds. By comparing best preparation to worst preparation time, you can determine whether a statement is expensive to compile.
Prep Worst (microsec)	The worst preparation time for an statement. This is the greatest amount of time needed to prepare the SQL statement. The value is shown in microseconds. By comparing worst preparation to best preparation time, you can determine whether a statement is expensive to compile.
Temp Data Hit Ratio	A statement's ratio of physical reads to logical reads for data pages read from temporary tablespaces.
Temp Index Hit Ratio	A statement's ratio of physical reads to logical reads for index pages read from temporary tablespaces.
Temp XML Hit Ratio	A statement's ratio of physical reads to logical reads for XML pages read from temporary tablespaces.
Text	The text of an SQL statement. For dynamic SQL statements, this is the SQL text associated with a package.
XML Hit Ratio	A statement's ratio of physical reads to logical reads for XML pages.

Related Topics
Top SQL drilldown
Current SQL drilldown
Static SQL drilldown
Dynamic SQL sub drilldown

Top SQL drilldown Static SQL tab

The Static SQL tab shows performance metrics for the static SQL statements run against a monitored database or partition. The statements include those that are executing and those that were executed earlier in the current

session.

Note: Static SQL statements are hard-coded in an application and do not change each time the application is run.

From this tab, you can drill down for more details on problematic statements and tune them.

How static SQL is captured

Information in the Static SQL tab is collected by the SQL Archiver. This is a Spotlight utility that captures and accumulates statistics for static SQL statements over time. Information is captured for statements that are in flight during a Spotlight refresh. Because only in-flight statements are captured, the information represents a sampling of static SQL statistics. It does not include information about every static SQL statement run against a database or all information about a certain statement. This means that a statement might not be shown in the Static tab or that the values shown might not reflect all executions of a statement.

The SQL Archiver is turned on by default. You can turn it on or off using SQL Archive options in the connection profile for a database or partition. You can also use these options to set capacity for the SQL Archiver.

Drill down on a statement

When the Static SQL tab shows a problematic statement, you can drill down to view additional performance information specific to that statement.

To drill down on a specific statement

Click the statement in the list of static statements.

The Static SQL Details sub drilldown is displayed at the bottom of the Static SQL tab. This sub drilldown shows the selected statement and provides additional metrics for this statement. Use these metrics to diagnose why the statement is performing poorly. Then, from within this sub drilldown, you can open a tuning session on the statement.

Tune a statement

From the Static SQL Details sub drilldown, you can open a tuning session on the currently displayed statement if you have either of the following products licensed and installed on the same computer where your Spotlight on IBM® DB2® LUW client is installed:

- Spotlight on IBM DB2 LUW

- SQL Tuning component of Quest Central for DB2

To tune the SQL statement currently displayed in the Static SQL Details sub-drilldown

1. Click ![icon] on the left side of the Static SQL Details sub drilldown.

Spotlight on IBM DB2 LUW opens, showing the SQL statement that you were examining in Spotlight in the product's main window.

2. If you need instructions to proceed with the tuning session, click F1 for context-sensitive help.

Columns in the Static SQL tab

All columns available for the Static SQL tab are listed below in alphabetic order. Initially, the tab contains Spotlight's default column selection. You can choose which columns to display and organize them in any order

from left to right using the Organize Columns function. This can be selected from the right-click menu for column headers.

Column	Description
Agents HWM Max	The highest number of agents used during execution of a statement.
Card Est Max	The estimated number of rows a statement will return. The estimate is made by the SQL compiler.
Consistency Token	The consistency token for the package that contains a statement. This helps to identify the version of the package. It is useful in situations where there are multiple versions for a certain combination of package name and creator.
Cost Est Max	The maximum estimated execution cost for a statement. Cost is shown in timerons. It is determined by the SQL compiler.
Cursor Name	The name of the cursor opened for a statement. If the statement does not use a cursor, a value is not shown.
Elapsed Time Max (msec)	The total amount of time spent executing a statement. The time is shown in milliseconds. If a statement is currently executing, the value is the amount of execution time elapsed so far.
Fetches Max	The total number of successful fetches performed during all observed executions of a statement.
Int Rows Deleted Max	The total number of rows deleted from the database due to internal activity caused by a statement. The activity can occur when the statement activates triggers on the table being referenced. It can also occur due to a cascading delete that enforces an ON CASCADE DELETE referential constraint. If the number of internal deletes is high, check to see whether the triggers and constraints defined for the table are necessary.
Int Rows Inserted Max	The total number of rows inserted in a database due to internal activity caused by a statement. The activity occurs when the statement activates triggers on the table being referenced. If the number of internal inserts is high, check to see whether the triggers defined for the table are necessary.
Int Rows Updated Max	The total number of rows updated in the database due to internal activity caused by a statement. The activity can occur when the statement activates triggers on the table being referenced. It can also occur due to a SET NULL row update that enforces an ON DELETE SET NULL referential constraint. If the number of internal updates is high, check to see whether the triggers and constraints defined for the table are necessary.
Last Agent ID	The unique identifier of the agent that was last used during execution of a statement.

Last Node	The number of the database partition on which a statement was last executed.
Last Sampled	The time a database or partition was last sampled by a statement. The timestamp includes year, month, date, and hour, in that order.
Num Agents Max	The maximum number of agents used during execution of a statement.
Overflows Max	The total number of sort overflows for a statement. These are sorts that ran out of sort heap and might have required disk space for temporary storage. The I/O required in writing to disk imposes additional overhead.
Package	The name of the package that contains a statement.
Parallelism Max	The maximum degree of intra-partition parallelism set for the database manager. This determines the maximum number of parallel operations a statement can use in a database partition. A number is shown if the maximum is user specified. ANY is shown if maximum is set to the default value. In this case, degree of parallelism is determined by the optimizer.
Rows Read Max	The total number of rows read during all executions of a statement.
Rows Written Max	The total number of rows changed (inserted, deleted, or updated) during all executions of a statement.
Samples	The number of times the SQL Archiver has sampled static SQL statements. (This value corresponds to the number of Spotlight refreshes that have occurred since the SQL Archiver was started.)
Schema	The authorization ID of the user who pre-compiled the application that contains a statement.
Section	The internal section number in the package that contains a statement.
Sort Time Max (msec)	The amount of time needed to execute sorts during executions of a statement. Time is shown in milliseconds.
Sorts Max	The total number of sorts performed for all executions of a statement.
Sys CPU Time Max (msec)	The total system CPU time used for all executions of a statement. The time is shown in milliseconds. A long system CPU time could mean that a statement would benefit from tuning.
User CPU Time Max (msec)	The total user CPU time used for all executions of a statement. The time is shown in milliseconds. A long user CPU time could mean that a statement would benefit from tuning.
Version	The version identifier for the package that contains a statement. This value is useful in situations where the version of a package was determined at

	precompile time for an embedded SQL program with the VERSION keyword. If a version was not specified at precompile time, a value is not shown.

Related topics
Top SQL drilldown
Current SQL drilldown
Dynamic SQL drilldown

Dynamic SQL Details drilldown for the Dynamic SQL tab

The Dynamic SQL Details sub drilldown provides additional information for the SQL statement selected in the Dynamic SQL tab of the Top SQL drilldown. The sub drilldown contains three panes:

- **Statement pane**—This pane shows the selected statement. If you want to tune the statement, you can open a tuning session in SQL Optimizer for DB2 by clicking the **SQL Tuning** icon.

- **Dynamic SQL Details pane**—This pane shows details that provide an indication of the impact of the selected statement on the database and system resources. The details include the number of rows deleted, updated, or inserted internally due to activity caused by the statement. They also include number of statement sorts and sort overflows, sort time, and rows read and written.

- **Dynamic SQL Reads pane**—This pane shows counts for the different types of reads performed during execution of the selected statement.

From this sub drilldown, you can also open a tuning session on the statement.

Related Topics
Details in the Dynamic SQL Details pane
Read types in the Dynamic SQL Reads pane

Unit of Work sub drilldown for the Current SQL tab

The Unit of Work sub drilldown provides additional information for the SQL statement that is currently executing for the unit of work. This is the statement that was selected in the Current SQL tab of the Top SQL drilldown. The sub drilldown contains three panes:

- **Statement pane**—This pane shows the SQL statement that is currently executing for the unit of work. If you want to tune the statement, you can open a tuning session in SQL Optimizer for DB2 by clicking the **SQL Tuning** icon.

- **UOW Details pane**—This pane shows details for the SQL statement that is currently executing for the unit of work. The details include statement start and stop times, number of concurrent agents, and number of blocking cursors.

- **UOW Activity pane**—This pane shows activity counts for the SQL statement that is currently executing for the unit of work. The counts provide a measure of how the current unit of work is impacting database resources.

From this sub drilldown, you can also open a tuning session on the statement.

Related Topics
Details in the UOW Details pane
Details in the UOW Activity pane

Printing the Top SQL drilldown

You can print the list of SQL statements on the SQL tab in the Top SQL drilldown; or you can print details for a specific statement on that tab.

To print the SQL statement list

1. If the Detail pane for a specific SQL statement is displayed on the drilldown, click ☒ in the upper right corner of the tab. The Detail pane is hidden.

2. Select **File ▶ Print**.

3. On the Print dialog, click **OK**.

To print the details for a specific statement

1. If the Detail pane that you want to print for a specific SQL statement is not displayed, select that statement from the list. The Detail pane is displayed.

2. Select **File ▶ Print**.

3. On the Print dialog, click **OK**.

Operating System drilldown

The Operating System drilldown provides information about the performance, processes, and storage on the server.

When you click the OS drilldown icon on the toolbar, the Operating System home page for the system where your database is running opens. For help with the OS drilldowns for Windows, Linux, or Unix, press F1 on the appropriate OS home page.

For specific information about monitoring a Unix or Linux Operating System, refer to help for Spotlight on Unix, and for monitoring a Windows Operating System, refer to help for Spotlight on Windows, or Spotlight on Windows Clusters. To open this help, select **Help | Contents**.

Workload Management Analysis drilldown

Note: This drilldown is available for DB2 9.5 and later databases only.

The Workload Management Analysis drilldown reports the information gathered by the WLM_GET table functions about the Workload Management (WLM) environment for the currently selected database.

Opening the Workload Management Analysis drilldown

To open the Workload Management Analysis drilldown

On the home page for the selected DB2 database or on the home page for the DB2 instance, click ![icon] in the toolbar.

Using the Work Management Analysis drilldown

From the following tabs on the drilldown, view both summary information and drill-down details about how workloads are being managed and how they are performing in the database:

Service Classes
Workloads
Work Classes
Work Actions
Activities
Threshold Queues

Workload Management Analysis drilldown - Service Classes tab

The Service Classes tab on Workload Management Analysis drilldown shows real-time summary and drill-down information about each service class defined on the current database.

A DB2 *service class* is a database object, assigned to one or more workloads, which defines the execution environment in which the activities in a workload run. That is, the service class is the primary point from which resources for database requests are assigned, work flow is managed, and activities are monitored in a workload.

A workload is assigned to one service class in the workload definition. However, a service class can be assigned to multiple workloads.

Service Class Summary Details

The tab lists the service classes currently defined on the database. These include both user-defined and system default service classes. Each top-level node in the list is the service superclass , which serves as the top tier for managing work environments. Expand a service superclass node to view its subclasses. The different service subclasses provide granularity in mapping resources and applying controls. A subclass is the actual environment where workload activities run.

The following summary statistics are shown for each subclass at database-partition level:

Superclass	Con. Act. HWM	Coord Act. Queue Time Std . Dev .
Subclass	Coord Act. LT HWM	Active Requests
Partition	Coord Act. LT Avg	Total Requests
Connection HWM	Coord Act. LT Std . Dev .	Req . Time Avg.
Coord Act. Compl .	Coord Act. Exec Time Avg.	Req . Time Std . Dev .
Coord Act. Aborted	Coord Act. Exec Time Std . Dev .	Req . Time Total
Coord Act. Rejected	Coord Act. Queue Time Avg.	Last Reset

Information About Specific Service Classes

To access in-depth information about what is happening across all workloads in your database at the service class level, click a service class entry and view any of the following sub-drilldowns. (Click a superclass to view information across all its subclasses; click a subclass to view information specific to the subclass.)

Workload Occurrences
Shows details about occurrences (that is, database connections assigned to a workloads) currently running across all workloads assigned to the service class. For each occurrence, view the attributes used to assign the connection to its workload, the current status of the occurrence, and the volume and success and failure rates of its activities.

Agents
Provides details about the agents handling activities associated with occurrences across all workloads in the service class.

Threshold Queues
Lists statistics for queued activities in all workloads in the service class. These statistics include queue times and sizes.

Class Activity
Graphically depicts activity totals for all workloads in the service class.

Related Topics
Workload Management Analysis drilldown

Workload Management Analysis drilldown - Workloads tab

The Workloads tab on Workload Management Analysis drilldown shows real-time information about each workload defined on the current database. DB2 *workloads* are objects to which database connections are automatically assigned. Connection attributes determine to which workload a specific connection is assigned.

Once a connection is assigned to a workload, DB2 Workload Manager can apply thresholds and resource controls to manage the execution of activities in the workload. The Workload Manager can also collect data about individual database activities and about the workload as whole to help you troubleshoot performance issues currently occurring and to avoid similar issues in the future.

Workload Summary Details

The tab lists the workloads currently defined in your Workload Manager environment on the database. These include both user-defined and system default workloads. Expand a workload node to view the following summary statistics collected for the workload at each database partition:

Name	Completed	Coord Act. Queue Time Avg (DB2 9.7 and later)
Partition	Last Reset	Coord Act. Queue Time Stddev (DB2 9.7 and later)
Concurrent HWM	Coord Act. Lifetime HWM (DB2 9.7 and later)	
Concurrent Act. HWM	Coord Act. Lifetime Avg (DB2 9.7 and later)	
CoordAct. Completed	Coord Act. Lifetime Stddev (DB2 9.7 and later)	

| CoordAct. Aborted | Coord Act. Exec Time Avg (DB2 9.7 and later) |
| CoordAct. Rejected | Coord Act. Exec Time Stddev (DB2 9.7 and later) |

Information About Specific Workloads

To access in-depth information about a specific workload, click its entry and view any of the following sub-drilldowns:

Occurrences

Shows details and statistics for each occurrence in the workload (that is, each database connection assigned to this workload). View the attributes used to assign the connection to the workload, the current status of the occurrence, and the volume and success and failure rates of its activities.

Agents

Provides details about the agents handling activities associated with occurrences in this workload.

Threshold Queues

Lists statistics for queued activities across all occurrences in the workload. These statistics include queue times and sizes.

Workload Activity

Graphically depicts activity totals for the workload across all occurrences in the workload.

Related Topics

Workload Management Analysis drilldown

Workload Management Analysis drilldown - Work Classes tab

The Work Classes tab on Workload Management Analysis drilldown shows real-time summary and drill-down information about each work class defined on the current database.

A DB2 *work class* is a database object that groups activities by their work type, such as READ, WRITE, CALL, DML, DDL, LOAD, or ALL. Work classes help the Workload Manager identify activities in a workload. However, to activate a work class, you must assign at least one work action to it. (See the help for Work Actions tab.) *Work actions* associated with the work class actually manage the identified activities.

Work classes are grouped in *work class sets* so that you can isolate multiple types of activities related to a certain area of work (for example READs and WRITEs involved large queries).

Work Class Summary Details

The tab lists the work class sets currently defined in your Workload Manager environment on the database. Expand a work class node to view the work classes assigned to it. The following information displays for each work class:

Set Name -- The name of work class set to which the selected work class belongs.

Class Name -- The name of the selected work class.

Information About Specific Work Classes

To access in-depth information about a specific work class, click its entry and view any of the following sub-drilldowns:

Activities

Shows details about each activity currently associated with this work class.

Threshold Queues

Shows details about each threshold queue containing activities associated with this work class.

Related Topics

Workload Management Analysis drilldown

Workload Management Analysis drilldown - Work Actions tab

The Work Actions tab on Workload Management Analysis drilldown shows real-time summary and drill-down information about each work action set defined on the current database.

A *work action set* contains a set of *actions*, or tasks, that you can perform on activities assigned to a given work class. (The work actions in a work action set are mapped to work classes in a given work class set.) These actions can include applying thresholds, collecting data, preventing execution, setting restrictions, and other actions.

Work Action Summary Details

The tab lists the work action sets currently defined in your Workload Manager environment. Expand the node for a work action set to view the work actions assigned to the set. The following information displays for each work action:

Set Name

Action Name

Enabled

Work Class Set Name

Work Class Name

Information About Specific Work Action Sets

To access in-depth information about a specific work action set, click a work action or work action set entry, and view any of the following sub-drilldowns:

Statistics

Lists the number of activities performed for each work class belonging to the work class set assigned to the work action set.

Activities

Provides details about each activity currently assigned to selected work action.

Threshold Queues

Provides details about each threshold queue currently in existence for the selected work action set.

Related Topics

Workload Management Analysis drilldown

Workload Management Analysis drilldown - Activities tab

The Activities tab on Workload Management Analysis drilldown shows real-time information about each activity running in applications connected to the current database. DB2 *workloads* are objects to which database connections are automatically assigned. Connection attributes determine to which workload a specific connection is assigned. Type topic text here.

Activity Summary Details

The tab lists the activities currently running in your Workload Manager environment. The activities are organized by application. Expand the node for an application to view its activities. The following information displays for each activity:

App . Handle	Activity ID	DB Work Action Set ID
Partition	Parent Activity ID	DB Work Class ID
Coord . Partition	Nesting Level	Svc . Class Work Action Set ID
State	Invocation ID	Svc . Class Work Class ID
Type	Routine ID	Start Time
UOW ID	Utility ID	
Parent UOW ID	Svc . Class ID	

Information About Specific Activities

To access in-depth information about a specific activity, click its entry. Any information returned for the WLM_ GET_ACTIVITY_DETAILS table function is shown in the following format:

Partition -- The database partition from which activity information was collected.

Element Name -- Name of the data element that WLM_GET_ACTIVITY_DETAILS returns.

Value -- The value of the data element.

Refer to your IBM DB2 LUW 9.5 documentation for more information about the elements returned by the WLM_ GET_ACTIVITY_DETAILS table function.

Related Topics

Workload Management Analysis drilldown

Workload Management Analysis drilldown - Threshold Queues tab

The Threshold Queues tab on Workload Management Analysis drilldown shows real-time information about the each threshold queue that the Workload Manager is currently managing in your current database. These queues are means of controlling workload resources. Rather than stopping activities or connections when too much work is executing in the workload, Workload Manager can assign these entities to a queue where they await their turn to run. Queues are established when the workload reaches certain user-defined thresholds that trigger queue creation.

Threshold Queue Summary Details

The tab lists the threshold queues currently defined in your Workload Manager environment on the database. Expand a queue node to view the following summary statistics collected for the queue at each database partition:

Name	Action Set	Queue Time Latest (msec)
ID	Work Class	Queue Exit Time Latest
Partition	Workload	Threshold Curr . Concurrency
Predicate	Queue Size HWM	Threshold Max. Concurrency

Domain	Queue Time Total (msec)	Last Reset
Superclass	Queue Assign. Total	
Subclass	Queue Size Current	

Information About Specific Threshold Queues

To access in-depth information about a specific threshold, click its entry and view the following sub-graph:

Queue Size -- The number of activities or connections still existing in the selected queue at the time specified on the graph.

Queue Time -- At the time specified on the graph, the amount of time that the last activity to be removed from the queue resided in that queue.

Related Topics

Workload Management Analysis drilldown

List of predefined alarms

Spotlight on IBM® DB2® LUW (Linux, Unix, and Windows) comes with some predefined alarms for the database or database partition home page and for the associated instance home page. For a list of these of alarms, click on the appropriate home page type.

Database or database partition home page alarms

Instance home page alarms

Database or database partition home page alarms

The following is a list of predefined alarms for database or database partition home page. To the right of each alarm name is the home page component and metric for which the alarm thresholds are defined. If necessary, you can use the Metric Editor to adjust the thresholds that correspond to these alarms to reflect your DB2 database or database partition environment.

Alarm	Home Page Panel	Component
Buffer Pool Overall Hit Rate alarm	Buffer Pool	Overall Hit Rate
Buffer Pool Index Hit Rate alarm	Buffer Pool	Index Hit Rate
Package Cache Hit Rate alarm	Buffer Pool	Package Cache Hit Rate
Package Cache Overflows alarm	Buffer Pool	Package Cache Hit Rate
Catalog Cache Hit Rate alarm	Buffer Pool	Catalog Cache Hit Rate
Catalog Cache Overflows alarm	Buffer Pool	Catalog Cache Hit Rate
Lock List Utilization alarm	Buffer Pool	Lock List Utilization

Percent of Sort Heap Threshold alarm	Buffer Pool	Pct of Thresholds
Sort Overflows alarm	Buffer Pool	Overflows
DMS Container Percentage alarm	Storage	DMS graph
DMS Utilization alarm	Storage	DMS Utilization
SMS Container Percentage alarm	Storage	SMS graph
SMS Utilization alarm	Storage	SMS Utilization
Log Container Percentage alarm	Active Log	Active Log graph
Log Utilization alarm	Active Log	Active Log Utilization

Buffer Pool Index Hit Rate alarm

The Buffer Pool Index Hit Rate alarm becomes active when the buffer pool index hit ratio, expressed as a percentage, for a database is low.

When the buffer pool hit rate is high, many of the index and data pages required to satisfy a query are found in a buffer pool. When the buffer pool hit rate is low, I/O is required to load the required pages from external storage.

Certain applications have low buffer hit rates for various reasons. Applications (such as data warehousing applications) that perform frequent table scans of large tables see little buffer pool improvement even with large buffer pools.

The default thresholds are predefined as follows:

- **Threshold 1**–The zero to 69 percent range. This raises a High severity alarm stating, The buffer pool hit ratio is extremely low.

- **Threshold 2**–The 70 to 79 percent range. This raises a Medium severity alarm stating, The buffer pool hit ratio is low.

- **Threshold 3**–The 80 to 89 percent range. This raises a Low severity alarm stating, The buffer pool hit ratio is moderate.

- **Threshold 4**–The 90 to 100 percent range. The buffer pool hit ratio is high. No alarms are raised.

Note: These hit rate thresholds should take into account the normal buffer pool activity incurred by the types of applications your database supports.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Access the Buffer Pools drilldown and the I/O Activity tab on the Databases drilldown for buffer pool details and statistics.
(missing or bad snippet)

Related topics
Metric Editor
What are thresholds and severities?

Buffer Pool Overall Hit Rate alarm

The Buffer Pool Overall Hit Rate alarm becomes active when the buffer pool hit ratio, expressed as a percentage, for a database is low.

When the buffer pool hit rate is high, many of the index and data pages required to satisfy a query are found in a buffer pool. When the buffer pool hit rate is low, I/O is required to load the required pages from external storage.

Certain applications have low buffer hit rates for various reasons. Applications (such as data warehousing applications) that perform frequent table scans of large tables see little buffer pool improvement even with large buffer pools.

The default thresholds are predefined as follows:

- **Threshold 1**–The zero to 69 percent range. This raises the High severity alarm stating, The buffer pool hit ratio is extremely low.

- **Threshold 2**–The 70 to 79 percent range. This raises the Medium severity alarm stating, The buffer pool hit ratio is low.

- **Threshold 3**–The 80 to 89 percent range. This raises the Low severity alarm stating, The buffer pool hit ratio is moderate.

- **Threshold 4**–The 90 to 100 percent range. The buffer pool hit ratio is high. No alarms are raised.

Note: These hit rate thresholds should take into account the normal buffer pool activity incurred by the types of applications your database supports.

The underlying metricsfor this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Access the Buffer Pool Summary drilldown and the I/O Activity tab on the Databases drilldown for buffer pool details and statistics.

Related topics
Metric Editor
What are thresholds and severities?

Catalog Cache Hit Rate alarm

The Catalog Cache Hit Rate alarm becomes active when the catalog cache hit rate, expressed as a percentage, is low.

The catalog cache is memory that temporarily stores table descriptors–information about tables, views, and aliases in a condensed internal format–found in the database catalog. During SQL statement compilation and dynamic SQL processing, the catalog cache is searched first for descriptors of objects referenced in the SQL statement. If these descriptors are available in cache, the transaction uses descriptor information in cache and avoids having to access the database catalog.

High catalog cache hit rates indicate that the I/O required to compile SQL statements is minimal. When the hit rate is low, catalog cache is not being used efficiently.

The underlying metricsfor this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- Consider increasing the size of the catalog cache using the `CATALOGCACHE_SZ` database configuration parameter.

- Access the Databases drilldown. The Caching tab provides statistics about catalog caching on the database.

Related topics

Metric Editor

What are thresholds and severities?

Catalog Cache Overflows alarm

The Catalog Cache Overflows alarm becomes active when an insert fails because the cache is full. Catalog cache space becomes available again when table descriptor information that is not being used is purged from the cache.

Note: Once a transaction has a catalog cache overflow event, every ensuing attempt by the same transaction to insert table descriptor information into the cache will also result in an overflow. The transaction involved will continue, but its descriptor information will not be inserted.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- **Large overflows**—The cache could be of inadequate size for the workload. Try enlarging the cache to improve performance.

- **Large transactions**—If you have transactions that are compiling a large number of SQL statements that reference a lot of tables, views, and aliases in a single unit of work, try compiling fewer SQL statements in a single transaction. This may improve the performance of the cache without having to increase its size.

- **Binding packages**—If you are binding packages that contain a large number of SQL statements that look up a lot of tables, views or aliases, try splitting your packages to include fewer SQL statements. This may improve the performance of the cache without having to increase its size.

Click the Catalog Cache component to open the Databases drilldown. This drilldown provides details about the database or partition you are monitoring. The Caching tab on the drilldown is in immediate focus to help you pinpoint problems related to catalog cache usage in the database or partition.

Related topics

Metric Editor

What are thresholds and severities?

DMS Container Percentage alarm

The DMS Container Percentage alarm becomes active when the percentage of utilized Database Managed Space (DMS) approaches the total amount of DMS storage in the DB2 instance.

The thresholds are predefined as follows:

- **Threshold 1**—The zero to 69 percent range. The percentage of DMS space utilized is low. No alarms are raised.

- **Threshold 2**—The 70 to 79 percent range. This raises the low (severity 4) alarm stating, The percentage of DMS space utilized is moderate.

- **Threshold 3**—The 80 to 89 percent range. This raises the medium (severity 5) alarm stating, The percentage of DMS space utilized is high.

- **Threshold 4**—The 90 to 100 percent range. This raises the high (severity 6) alarm stating, The percentage of DMS space utilized is very high.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Access the Tablespaces drilldown to acquire detailed information about tablespaces and their containers.

Related topics

Metric Editor

What are thresholds and severities?

DMS Utilization alarm

The DMS Utilization alarm becomes active when the percentage of utilized Database Managed Space (DMS) approaches the total amount of DMS storage in the DB2 instance.

The thresholds are predefined as follows:

- **Threshold 1**—The zero to 69 percent range. The percentage of DMS space utilized is low. No alarms are raised.

- **Threshold 2**—The 70 to 79 percent range. This raises the low (severity 4) alarm stating, The percentage of DMS space utilized is moderate.

- **Threshold 3**—The 80 to 89 percent range. This raises the medium (severity 5) alarm stating, The percentage of DMS space utilized is high.

- **Threshold 4**—The 90 to 100 percent range. This raises the high (severity 6) alarm stating, The percentage of DMS space utilized is very high.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Access the Tablespaces drilldown to acquire detailed information about tablespaces and their containers.

Related topics

Log Container Percentage alarm

The Log Container Percentage alarm becomes active when the percentage of log space currently used by the database is approaching the total amount of log space available for the database.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- Access the Databases drilldown. The Log Status tab identifies the values of key log configuration parameters as well as the log transaction rate and the average amount of log space used per transaction.

- Consider adjusting the LOGFILSZ , LOGPRIMARY , and LOGSECOND database configuration parameters.

Related topics

Log Utilization alarm

The Log Utilization alarm becomes active when the percentage of log space currently used by the database is approaching the total amount of log space available for the database.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- Access the Databases drilldown. The Log Status tab identifies the values of key log configuration parameters as well as the log transaction rate and the average amount of log space used per transaction.

- Consider adjusting the LOGFILSZ , LOGPRIMARY , and LOGSECOND database configuration parameters.

Related topics

Package Cache Hit Rate alarm

The Package Cache Hit Rate alarm becomes active when the package cache hit rate, expressed as a percentage, is low.

Package cache is memory that temporarily stores package and section information required for the execution of static and dynamic SQL statements. When applications obtain static SQL information from package cache, they

eliminate I/O on system catalogs. When applications obtain dynamic SQL information from package cache, they avoid the cost of compiling the SQL statements.

Low package cache hit rate means that you are incurring more overhead to reload static SQL and recompile dynamic SQL.

The default thresholds are predefined as follows:

- **Threshold 1**—The zero to 69 percent range. This raises High severity alarm stating, The package cache hit ratio is extremely low.

- **Threshold 2**—The 70 to 79 percent range. This raises a Medium severity alarm stating, The package cache hit ratio is low.

- **Threshold 3**—The 80 to 89 percent range. This raises a Low severity alarm stating, The package cache hit ratio is moderate.

- **Threshold 4**—The 90 to 100 percent range. The package cache hit ratio is high. No alarms are raised.

The underlying metricsfor this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

You can consider increasing the size of package cache using the PCKCACHESZ database configuration parameter. If your database supports heavy transaction processing volumes, your current package cache size might not be able to accommodate these volumes.

Related topics
Metric Editor
What are thresholds and severities?

Package Cache Overflows alarm

The Package Cache Overflows alarm becomes active when package cache inserts have resulted in overflows. An overflow occurs when the limit set by the PCKCACHESZ database configuration parameter is exceeded. The package cache is forced to borrow memory from other entities in database shared memory, causing possible lock list escalation, loss of concurrency, out of memory errors, and overall performance degradation.

This alarm has been predefined with a severity level of 4 (low alarm). The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Determine whether or not you need to increase the size of the package cache to avoid overflows. Clicking on the Package Cache element opens the Databases drilldown where you can review statistics for the package cache.

Related topics
Metric Editor
What are thresholds and severities?

Lock List Utilization alarm

The Lock List Utilization alarm becomes active when the amount of lock list storage used, expressed as a percentage, approaches its maximum size, as defined by the LOCKLIST database configuration parameter.

Lock list is the memory allocated for locks in a database. Once the lock list is full, database performance degrades. Lock escalation increases, reducing concurrency on shared objects in the database. Additionally, because applications are waiting on a limited number of table locks, more deadlocks between applications can occur, causing transactions to be rolled back.

The thresholds are predefined as follows:

- **Threshold 1**—The zero to 69 percent range. The lock list utilization percentage is very low. No alarms are raised.

- **Threshold 2**—The 70 to 79 percent range. This raises a Low severity alarm stating, The lock list utilization percentage is moderate.

- **Threshold 3**—The 80 to 89 percent range. This raises a Medium severity alarm stating, The lock list utilization percentage is high.

- **Threshold 4**—The 90 to 100 percent range. This raises a High severity alarm stating, The lock list utilization percentage is very high.

The underlying metricsfor this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- Perform frequent COMMITs to release locks.

- When performing many updates, lock the entire table before updating (using the SQL LOCK TABLE statement). This tactic uses only one lock and keeps others from interfering with the updates. However, it does reduce concurrency of the data.

- Use the LOCKSIZE parameter of the ALTER TABLE statement to control how locking is performed for a specific table.

- Access the Databases drilldown. The Locking tab allows you to profile the database's lock list utilization.

- Access the Client Applications drilldown. The Locking tab identifies those applications holding locks as well as those waiting for locks.

Related topics
Metric Editor
What are thresholds and severities?

Percent of Sort Heap Threshold alarm

The Percent of Sort Heap Threshold alarm becomes active when the amount of sort heap space allocated approaches or exceeds the SHEAPTHRES database manager configuration parameter value. The SHEAPTHRES parameter value is the sort heap threshold for the entire instance.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- Determine whether the current SHEAPTHRES parameter value is appropriate for your system. Insufficient sort heap allocation causes the database manager to select less efficient sorting techniques that can severely impact system performance.

- Access the Client Applications and Databases drilldowns.

  - The Sorting tab in both drilldowns can help you profile the sorting requirements and sort exception rates at the instance, database, and application levels.

  - The Details tab can help you identify databases and applications that are consuming significant amounts of CPU time to perform sorts.

Related topics
Metric Editor
What are thresholds and severities?

SMS Container Percentage alarm

The SMS Container Percentage alarm becomes active when the percentage of utilized System Managed Space (SMS) approaches the total amount of SMS storage in the DB2 instance.
The default thresholds are predefined as follows:

- **Threshold 1**–The zero to 69 percent range. The percentage of SMS space utilized is low. No alarms are raised.

- **Threshold 2**–The 70 to 79 percent range. This raises a Low severity alarm stating, The percentage of SMS space utilized is moderate.

- **Threshold 3**–The 80 to 89 percent range. This raises a Medium severity alarm stating, The percentage of SMS space utilized is high.

- **Threshold 4**–The 90 to 100 percent range. This raises a High severity alarm stating, The percentage of SMS space utilized is very high.

The underlying metrics for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Click the SMS database files component to open the Tablespaces drilldown. This drilldown provides details about the tablespaces in which the database you are currently monitoring resides. You can select a specific tablespace to view its storage and activity statistics.

Related topics
Metric Editor
What are thresholds and severities?

SMS Utilization alarm

The SMS Utilization alarm becomes active when the percentage of utilized System Managed Space (SMS) approaches the total amount of SMS storage in the DB2 instance.

The default thresholds are predefined as follows:

- **Threshold 1**–The zero to 69 percent range. The percentage of SMS space utilized is low. No alarms are raised.

- **Threshold 2**–The 70 to 79 percent range. This raises a Low severity alarm stating, The percentage of SMS space utilized is moderate.

- **Threshold 3**–The 80 to 89 percent range. This raises a Medium severity alarm stating, The percentage of SMS space utilized is high.

- **Threshold 4**–The 90 to 100 percent range. This raises a High severity alarm stating, The percentage of SMS space utilized is very high.

The underlying metricsfor this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Click the SMS database files component to open the Tablespaces drilldown. This drilldown provides details about the tablespaces in which the database you are currently monitoring resides. You can select a specific tablespace to view its storage and activity statistics.

Related topics
Metric Editor
What are thresholds and severities?

Sort Overflows alarm

The Sort Overflows alarm becomes active when the percentage of sort at the database or partition exceeds the overflow threshold limits. Sort heap is controlled by the SHEAPTHRES database manager configuration parameter, which sets the sort heap threshold for the entire instance and by the SORTHEAP database configuration parameter, which sets the sort heap size. Sort overflows are sorts that run out of sort heap and require disk space for temporary storage.

The thresholds are predefined as follows:

- **Threshold 1**–The zero to 69 percent range. The percent of sort overflows is very low. No alarms are raised.

- **Threshold 2**–The 70 to 79 percent range. This raises a Low severity alarm stating, The percent of sort overflows is moderate.

- **Threshold 3**–The 80 to 89 percent range. This raises a Medium severity alarm stating, The percent of sort overflows is high.

- **Threshold 4**–The 90 to 100 percent range. This raises a High severity alarm stating, The percent of sort overflows is very high.

The underlying metricsfor this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

- If the percentage of overflowed sorts is high, increase the sort heap.

- If there were any post-threshold sorts, increase the SHEAPTHRES configuration parameters. To determine if there were any post threshold sorts, click the Overflows component to open the Databases drilldown. This drilldown provides details about all the active databases in the DB2 instance. You can select any database to view details about that database, including sort activity and sort heap usage statistics.

Related topics
Metric Editor
What are thresholds and severities?

Instance home page alarms

The following is a list of predefined alarms for the instance home page. To the right of each alarm name is the home page component and metric for which the alarm thresholds are defined. If necessary, you can use the Metric Editor to adjust the thresholds that correspond to these alarms to reflect your DB2 instance environment.

Alarm	Home Page panel	Component
Post Threshold Sorts alarm	Global Sort/Join	Post Threshold Sorts
Post Threshold Hash Joins alarm	Global Sort/Join	Post Threshold Hash Joins
Piped Sort Rejects alarm	Global Sort/Join	Piped Sort Rejects
Sort Overflows alarm	Global Sort/Join	Sort Overflows
Hash Loops alarm	Global Sort/Join	Hash Loops
Hash Join Small Overflows	Global Sort/Join	Hash Join Overflows
Hash Join Small Overflows alarm	Global Sort/Join	Hash Join Small Overflows
Agent Request Overflows alarm	Global Application/Agent	Agent Request Overflows
Agents Pending Locks alarm	Global Application/Agent	Agents Pending Locks
Agents Pending Token alarm	Global Application/Agent	Agents Pending Token
Missed Beats alarm	HADR	Missed Beats

Agents Pending Locks alarm

The Agents Pending Locks alarm becomes active when the number of agents waiting on a lock in comparison to the total number of registered agents for the DB2 instance reaches a threshold percentage. Lock waits occur for a number of reasons, including internal database contention.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

The application agents could be experiencing concurrency problems. Try to identify the applications that are holding locks for long periods of time and determine the cause. One possible solution is to use small units of

work (for example frequent COMMITs) to reduce concurrency issues.

Access the Client Application Analysis drilldown to identify the applications holding locks and review statistics about those locks.

Notes:

- When you change the locking configuration parameters in one partition for a database, be sure that the changes are also made to all of the other partitions for the database.

Related topics

Metric Editor

What are thresholds and severities?

Agents Pending Token alarm

The Agents Pending Token alarm becomes active when the number of agents waiting for a token (a permission to run) in comparison to the total number of registered agents for the DB2 instance reaches a threshold percentage.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Determine whether the current value for the MAXAGENTS database configuration parameter is appropriate for your system. Some environments might use the MAXAGENTS parameter value to limit the number of processing tokens, thereby limiting the amount of execution concurrency in the database manager. Although this strategy might be beneficial to the system as a whole, it can impact applications, causing problems such a high percentage of agents awaiting tokens.

Related topics

Metric Editor

What are thresholds and severities?

Agent Request Overflows alarm

The Agent Request Overflows alarm becomes active when a request for an agent was made after the limit for the MAXAGENTS database manager configuration parameter was reached. Once this limit is reached, the application begins to search for idle associated agents currently residing in the instance.

This alarm has been predefined with Low and Information severity levels. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Try reducing the workload if the MAXAGENTS limit is exceeded on a regular basis.

Related topics

Metric Editor

What are thresholds and severities?

Hash Join Small Overflows

The number of times that a hash join overflowed available sort heap space by less than 10%. If this number and the number of hash join overflows are both high, try increasing SHEAPTHRES parameters. If the small overflows number is more than 10% of the number of hash join overflows, try increasing the sort heap size.

Note: In a multiple-partition database, information is presented for each active partition on the system. An All Nodes row keeps a running total for all partitions. In a single-partition database, only the All Nodes information is presented.

Hash Join Small Overflows alarm

The predefined Hash Join Small Overflows alarm becomes active when a hash join overflows available sort heap space by less than 10%.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

If **Hash join small overflows** and **Hash join overflow** values are both high, try increasing SHEAPTHRES parameters. If the small overflows number is more than 10% of the number of **Hash join overflows**, try increasing the SORTHEAP size.

Related topics

Metric Editor

What are thresholds and severities?

Hash Loops alarm

The predefined Hash Loops alarm becomes active when a single partition of a hash join is larger than the available sort heap. This could indicate that you have a poorly performing hash join.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Try increasing the SHEAPTHRES for the instance or the SORTHEAP for the database or partition to alleviate this issue.

Related topics

Metric Editor

What are thresholds and severities?

Piped Sort Rejects alarm

The Piped Sort Rejects alarm becomes active when piped sort requests are rejected. A piped sort is rejected if the sort heap threshold is exceeded at the time the sort heap is allocated.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Increasing the sort heap threshold for the instance, or the sort heap space for the database can improve piped sort acceptance.

Related topics
Metric Editor
What are thresholds and severities?

Post Threshold Hash Joins alarm

The predefined Post threshold Hash Joins alarm becomes active when a hash join request was limited because of concurrent use of sort heap space (shared or private).

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

If this value is larger than 5% of total hash join overflows, increase the sort heap threshold (SHEAPTHRES configuration parameter).

Related topics
Metric Editor
What are thresholds and severities?

Post Threshold Sorts alarm

The Post Threshold Sorts alarm becomes active when a sort has been requested after the sort heap threshold has been reached. Sorts that start after the sort heap threshold (sheapthres configuration parameter) has been reached may not receive all the memory they need to run efficiently. By adjusting the sort heap threshold and sort heap size configuration parameters, you can improve sort operations and overall system performance.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

If post-threshold sort values are high, you can try the following possible solutions:

- Increasing `SHEAPTHRES`

- Making SQL query changes that allow applications to use fewer or smaller sorts.

Related topics
Metric Editor
What are thresholds and severities?

Sort Overflows alarm

The predefined Sort Overflows alarm becomes active when sorts have run out of sort heap and may have used disk space for temporary storage. There is an additional overhead cost when sorts overflow because the sort will need a merge phase. There is also the possibility of added I/O if the data does need to be written out to disk.

This alarm has been predefined with a Low severity level. The underlying metric for this alarm can be adjusted to suit your environment by using the Metric Editor.

What you can do

Try these performance turning methods:

- **At the database or application level**–If sort overflows are high, try increasing the sort heap size.

- **At the statement level**–SQL statements that require large sorts may benefit from further tuning to reduce the amount of sorting.

Related topics
Metric Editor
What are thresholds and severities?

Acquisition Options

Note: This topic focuses on information that may be unfamiliar to you. It does not include all step and field descriptions.

To set Acquisition Options

1. Select **View | Options | Spotlight on IBM DB2 LUW options | Acquisition**.

2. Review the following for additional information:

General	Description
Inherit settings from selected connection	Select the checkbox to inherit settings from existing connections. Default: *unchecked*
Show current running static SQL.	Select the checkbox to display static SQL currently running in **Current SQL** or**Client Application	Unit of work-SQL**. Default: *unchecked*
Application Log	**Description**
Log Size	Set the log size for each connection.
Log Location	Select the location for the log file.
Log Level	Designate the level of information detail in the log file.

DB2 Client Options

Note: This topic focuses on information that may be unfamiliar to you. It does not include all step and field descriptions.

To set a default DB2 client for Spotlight

1. Select **View | Options | Spotlight on IBM DB2 LUW options | DB2 Client**.

2. Select a DB2 client.
   Note: You must close and restart Spotlight on DB2 when you change this option.

3. Clickl **OK**.

To import a DB2 Catalog from another DB2 client to the current client

1. Select **View | Options | Spotlight on IBM DB2 LUW options | DB2 Client**.

2. Click Import DB2 Catalog .

3. If a DB2 client is installed, select a client from the list to automatically import a catalog profile from another client to the selected DB2 client.

4. Click **OK.**
   Note: If no other client is installed, **OK** is disabled.
   Review the following for additional information:

Tuning an SQL statement listed in the Top SQL drilldown

You can tune SQL statements listed in the Top SQL drilldown if you have either of the following products licensed and installed on the same computer where your Spotlight on IBM® DB2® LUW client is installed:

- SQL Optimizer for DB2

To start a tuning session

1. Open the Top SQL drilldown.

2. Select one of the tabs shown as follows:

   - Select the Current SQL tab to tune currently executing statements.

   - Select the Static SQL tab to tune static SQL statements.

   - Select the Dynamic SQL tab to tune dynamic SQL statements.

3. From the list in the selected tab, click the row for the statement you want to tune. The detail pane for this statement appears in the lower half of the SQL tab

4. In the upper left of the detail pane, click the **SQL Tuning** icon .
   SQL Optimizer for DB2 opens, showing the SQL statement in the product's main window.

5. If you need instructions to proceed with the tuning session, click F1 for context-sensitive help.

Related topics

Tuning an SQL statement listed in the Client Application Analysis drilldown

You can tune SQL statements listed in the Client Application Analysis drilldown if you have either of the following products licensed and installed on the same computer where your Spotlight on IBM® DB2® LUW client is installed:

- SQL Optimizer for DB2

To start a tuning session

1. Open the Client Application Analysis drilldown.

2. From the list of applications, click in the application whose currently executing SQL statement you want to tune.

A set of tabs, containing detail information about the application, displays in the lower half of the drilldown window.

3. Click the Unit of Work - SQL tab to examine the SQL statement currently executing for the application.

4. In the upper left of the detail pane of this tab, click the **SQL Tuning** icon . SQL Optimizer for DB2 opens, showing the SQL statement in the product's main window.

5. If you need instructions to proceed with the tuning session, click F1 for context-sensitive help.

Related topics

Client Application Analysis drilldown

Spotlight Basics

Spotlight Basics describes how Spotlight applications work. It contains these sections:

Section	Description
Spotlight Connections	Create / Modify / Delete connections to systems, and monitor those systems.
Monitor Spotlight Connections	Monitor connections, detect alarms, investigate the cause of alarms.
Alarms	Alarms are the warnings that Spotlight raises when a metric falls outside its "normal" range of values.
Charts, Grids and Home Page Components	Spotlight displays provide detailed statistics of the monitored system, relevant to diagnosing the cause of alarms. Make the best use of them.
View | Options	Customize Spotlight.
Troubleshooting	Solve problems using Spotlight.

Spotlight Connection Manager

Use the Spotlight Connection manager to create connections to the systems you intend to monitor, to monitor those connections and to modify the properties of those connections.

To open the Spotlight Connection Manager

Click **File | Connect**

Use the Spotlight Connection Manager to:

Option	Description
Add a connection	Click Connection Properties - Details Page **Tip**: You may want to select the type of connection in the **Connections** pane first.
Start monitoring a connection	Select the connection in the Spotlight Connection Manager and click **Connect**. **Note**: To stop monitoring a connection, close the Spotlight Connection Manager and click **File | Disconnect**.
Rename a connection	Right click the connection in the Spotlight Connection Manager and select **Rename**.
Edit properties for a connection	Right click the connection in the Spotlight Connection Manager and select **Properties**. Connection Properties - Details Page
Delete a connection	Right click the connection in the Spotlight Connection Manager and select **Delete**. This removes the connection from the Spotlight Connection Manager.

Tips:

- You can use the Spotlight Browser to manage connections once they have been added. Spotlight Browser

- To simplify the list of connections select the connection type to view. For example, if Spotlight on Windows is installed then click Spotlight on Windows to view only Windows connections. If the connection list is still long you may like to click **Tools | View | List**.

- You may have access to shared connections. Migrate (Move And Share) Connections

Related Topics
Spotlight Browser
Connection Properties - Details Page

Connection Properties - General Page

View general information about the spotlight connection.

To open the Connection Properties dialog, General page

1. Locate the connection in the Spotlight Connection Manager or the Spotlight Browser.

2. Right click the connection and select **Properties | General**.

About the General page

Note: You cannot edit the fields in this page directly while the connection is open.

Field	Description
Connection name	The display name for the connection.
Type	The type of connection. (DB2, Windows etc.)
Location	The location of the connection profile. Connection profiles store information that Spotlight uses to connect to the system.
Calibrated	Has the connection been calibrated? When? Calibrate Spotlight Connections
Template	The name of the template applied to the connection. Connection Templates **Note**: Click **Change template** to change the template.
Created	The date the connection was created.
Modified	The date the connection was last modified.
Accessed	The date the connection was last accessed.

Notes:

- User defined connection profiles are stored by default in C:\Documents and Settings*user*\Application Data\Quest Software\Shared\Profiles

- User defined templates are stored by default in C:\Documents and Settings*user*\Application Data\Quest Software\Shared\Templates

- Locations for connection profiles and templates may differ from the default in your case. Migrate (Move And Share) Connections

Related Topics
Connection Properties - Details Page
Connection Properties - Overhead Page

Connection Properties - Details Page

View / Enter / Modify the settings for the Spotlight connection.

To open the Connection Properties dialog, Details page

1. Locate the connection in the Spotlight Connection Manager or Spotlight Browser.

2. Right click the connection and select **Properties | Details**.

About the Details page

The fields in the Details page are dependent on the type of Spotlight connection.

- For help on a specific field, click in the field and press **F1**.

- For help on all fields, click **Help | Contents |** *Connection type* **| Connect to A ... System**.

Note: You cannot edit the fields directly while the connection is open.

Related Topics
Connection Properties - General Page
Connection Properties - Overhead Page

Connection Properties - Overhead Page

The queries that Spotlight uses to collect data from a Spotlight connection can place an additional load on the monitored system.

- Simple queries that retrieve small amounts of data have little impact on the performance of the monitored system.

- Most Spotlight queries consume a measurable but not significant amount of system resources.

- Queries that are complex or collect large amounts of data may put a significant load on the system.

The controls on the Overhead Management page allow you to adjust collection rates for different categories of Spotlight data, and so limit the load that Spotlight places on the monitored system.

Notes: Use this page to change the overall settings for the connection. To change the settings for a Spotlight component use the Component Properties editor. Component Properties - Overhead Page

To open the Connection Properties dialog, Overhead page

1. Locate the connection in the Spotlight Connection Manager or the Spotlight Browser.

2. Right click the connection and select **Properties | Overhead**.

Alternatively

1. Select the connection in the Spotlight Browser Live connections.

2. Select **Tools | Performance Settings** or click .

About the Spotlight overheads dialog

Control	Description
	Drag the slider to one of the standard overhead settings:
	• Higher impact - Spotlight collects more data but may put a higher load on the system.
	• Medium impact - Spotlight collects less data or collects it less often.
	• Lower impact - Spotlight puts a minimal load on the system while still collecting essential data.
Overhead Impact slider	Accompanying text further defines the data collected under each overhead setting. The definitions are dependent on the connection type. The criteria may include:
	• The kind of data being collected
	• The load that collecting the data will put on the system
	• Where the data in the category is to be displayed
	• The importance of the data
	• How often the data is needed
	Note: The slider is not displayed if a component in the connection already has a custom setting. Click Restore default to remove all custom settings.
Custom...	Click to select a collection rate for each available category of Spotlight data. Custom Connection Performance
Restore default	Click to remove all user-specified collection rates. Revert to Spotlight default settings. Set the rate at which data is collected
Make this the default setting for all new connections	When selected, all new Spotlight connections collect data at the rates defined in this dialog.

Related Topics

Connection Properties - General Page
Connection Properties - Details Page

Custom Connection Performance

Use the Custom Connection Performance page to choose specific collection rates for every available category of Spotlight data.

To open the Custom Connection Performance dialog

1. Open the Connection Properties - Overhead Page.

2. Click **Custom**.

About the Custom Connection Performance dialog

The **Category / Impact** table lists each category of Spotlight data and defines the effect of collecting that data on the monitored system.

1. Select a category. This enables the Refresh setting list.

2. Select from the Refresh setting list how often Spotlight will collect data in the specified category.

Refresh setting	Description
Refresh at the default rate	Spotlight data is refreshed as per the definition in Spotlight Options. Set the rate at which data is collected **Note**: In most cases (but not all) this will be at the foreground rate when the component is visible and the background rate when the component is not visible.
Always refresh at the foreground rate	Always refresh at the foreground rate, whether the data is displayed or not.
Refresh only while data is displayed on the screen	Refresh data at the foreground rate when it is visible in the Spotlight window.
Do not refresh data in this category	The data is never collected. Spotlight displays one of the following icons on components that display this category of data: Data for this control is not available. (There are no alarms associated with this control.) Data for this control is not available. (Alarms associated with the control are disabled.)

Related Topics

Connection Properties - Overhead Page

Set the rate at which data is collected

Migrate (Move And Share) Connections

Spotlight templates and connection profiles control the appearance and behavior in Spotlight of monitored Spotlight connections.

New Spotlight connections base their settings on templates.

Spotlight uses connection profiles to store property data for individual Spotlight applications and connections.

To migrate connections

Migrate your Spotlight connection files to a network where other Spotlight users can share them. By migrating profiles to a shared location, you can share connection information between multiple users. User-specific information (including user names and passwords) is still stored on a per-user basis. Users are identified by the Windows accounts under which they log in.

1. Click **File | Disconnect All** to disconnect ALL open Spotlight connections.

2. Select **Tools | Migrate Connections**.

3. Click **Next | Migrate connections | Next**.

4. Browse to select the location to migrate connections to.

Note:

- Do not migrate connections to locations that may be offline frequently or for extended periods.

- You cannot migrate connections that are currently in use or inaccessible.

- You cannot move connections to another location if you have no Spotlight connections to move.

- Templates shipped with Spotlight are stored by default in Console\Templates.

- User defined connection profiles are stored by default in C:\Documents and Settings*user*\Application Data\Quest Software\Shared\Profiles

- User defined templates are stored by default in C:\Documents and Settings*user*\Application Data\Quest Software\Shared\Templates

- The locations for updated user templates and profiles for EXISTING Spotlight installations remain as they were before.

To use shared (migrated) connections

First ensure the other user has migrated Spotlight connection files for you to use.

1. Select **File | Disconnect All** to disconnect ALL open Spotlight connections.

2. Select **Tools | Migrate Connections**.

3. Click **Next | Use shared connections | Next**

4. Browse to select the shared location.

Related Topics
Connection Templates
Connection Properties - General Page

Connection Templates

The template file for a Spotlight application contains the information that controls the standard appearance and behavior of Spotlight connections for that connection type. Spotlight is shipped with a set of templates that are suitable for most connections.

Once a Spotlight connection has been configured and calibrated, you can save its configuration as a template for use with Spotlight connections to similar systems. Changes that are made to a template propagate to all connections that use that template, unless individual changes have already been made to those connections.

The information contained in the template includes:

- The configuration of controls in the home page and drilldowns, including the metrics displayed and the data sources from which they are derived.

- Calibration values for those controls.

- The threshold values for alarms on those controls.

To save the configuration for the current Spotlight connection as a template

1. Select the Spotlight connection in the Spotlight Browser in Live connections

2. Select **File | Save as Template**.

3. Select **Save calibration** to save calibration data in the template.

Notes:

- You can change the template assigned to an existing Spotlight connection. Connection Properties - General Page

- Templates shipped with Spotlight are stored by default in Console\Templates.

- User defined templates are stored by default in C:\Documents and Settings*user*\Application Data\Quest Software\Shared\Templates

- Template files are given extension **.stx**

- Locations for template files may differ from the default in your case. Migrate (Move And Share) Connections

Related Topics
Calibrate Spotlight Connections
Metrics Dialog
Thresholds Dialog
Spotlight Home Page Components

Spotlight Browser

Use the Spotlight browser to view the details of your Spotlight connections. The browser contains a number of connection groups, each of which can be opened by clicking on its bar within the browser:

Option	Description
Welcome to Spotlight	View Spotlight's introductory page and web pages of interest to Spotlight users.
Live connections	View the Spotlight connections currently being monitored by Spotlight.
All connections	View all Spotlight connections (open or closed). The connections are displayed in tree form with connection types as the topmost branches, and individual connections as the child branches of each type.
History Browser	View historical snapshots for the current Spotlight connection. **Note**: The History Browser may be hidden until you select **View	History browser**.

Notes:

- To hide the Spotlight browser click or **View | Navigation Tree**.

- You can customize the appearance of the Spotlight Browser. Navigation tree

Related Topics

Navigation tree

Spotlight Connection Manager

Welcome to Spotlight

Welcome to Spotlight is the first page you see when you start Spotlight. New Spotlight users can use this page to discover and experiment with various Spotlight features.

Action	Description
Open the Welcome to Spotlight page	From the Spotlight Browser, select **Welcome to Spotlight | Welcome to Spotlight**.
	From the Spotlight Browser, select **Welcome to Spotlight | Spotlight on the Web** | *Web Site*.
Open web pages of interest to Spotlight users	***Toolbar buttons on Spotlight on the Web***

Click	Description
	Navigate to the previous page in the web browsing sequence.
	Navigate to the next page in the web browsing sequence.
	Stop downloading the web page.
	Refresh (Reload) the web page.
	Return to the page selected from the **Spotlight Browser**.

Related Topics

Spotlight Browser

Live connections

All connections

History Browser

Live connections

View the Spotlight connections that are currently monitored by Spotlight.

Action	Description
Select / Click the connection	The Spotlight menus and toolbar adjust to the selected connection. The last viewed page for this connection opens.
	Note: The name of the connection takes on the color of the most urgent severity now active for the connection. Spotlight Alarms
Right click | **view**	Sort the list of connections:
	- Sort by Name - Sort alphabetically.
	- Sort by Severity - Sort according to the severity of alarms currently raised for each connection.

Action	Description
Right click | **Disconnect**	Stop monitoring this connection.
Right click | **Browse History**	Open the History browser for the connection. View snapshots of the system as it operated in the past. History Browser
Right click | **View Alarm Log**	View the Alarm log for the connection. Alarm log
Right click | **View Alarms by Time**	View Alarms by Time for the connection. Alarms by Time
Right click | **Properties**	View properties for the connection. • Connection Properties - General Page • Connection Properties - Details Page • Connection Properties - Overhead Page **Note**: Most fields cannot be modified while the connection is open.

Notes:

- The color of the Spotlight browser Show / Hide button reflects the highest current severity raised by connections in Live connections.

- To stop monitoring a connection, select the connection and click **File | Disconnect**.

- You can start monitoring a connection in the Spotlight Browser. All connections

Related Topics

Spotlight Browser

Welcome to Spotlight

All connections

History Browser

All connections

View all Spotlight connections (open or closed) defined in the Spotlight Connection Manager.

Action	Description
Click on the connection type	List all connections defined for this connection type.
Click on the connection	Start monitoring the connection. This opens the connection in Live connections.
Right click | **Connection Manager**	Open the Spotlight Connection Manager. Spotlight Connection Manager
Right click | **Browse History**	Open the History browser for the connection. View snapshots of the system as it operated in the past. History Browser
Right click | **View Alarm Log**	View the Alarm log for the connection. Alarm log
Right click | **View Alarms by Time**	View Alarms by Time for the connection. Alarms by Time
Right click | **Properties**	View properties for the connection. • Connection Properties - General Page

Action	Description
	• Connection Properties - Details Page
	• Connection Properties - Overhead Page

Note: Most fields cannot be modified while the connection is open.

Related Topics

Spotlight Browser

Welcome to Spotlight

Live connections

All connections

History Browser

An important feature within Spotlight is the ability to display, collect and replay the behavior of a Spotlight connection at a specified point in time such as when an important alarm was raised.

The details shown via the History browser are identical in format to those displayed by Spotlight in its "live" state, so you can view historical data in the same way as you would view a live connection.

To view Spotlight history

1. Right click the connection in the Spotlight Browser Live connections or All connections.

2. Select **Browse History**.

Select an event in the snapshot list.

The Spotlight browser displays a tree of alarms and snapshots in chronological order, together with some basic navigation controls.

Action	Description
Click Last week, Today etc.	Spotlight maintains a list of snapshots that have been raised for each day in the specified historical period. Click a day to have Spotlight display a list of all snapshots that were recorded for that day.
Click a snapshot or alarm	View the status of the connection at the time of the snapshot or alarm.
Right click | **Live view**	Return the view of this connection to the present time. Live connections
Right click | **History Interval**	Define the interval between snapshots displayed in the History browser. From the drop-down list, choose the frequency with which you want to display a snapshot.

History browser toolbar

Control	Description
	Previous alarm. Show the snapshot corresponding to the alarm before the one currently displayed.
	Previous snapshot. Show the snapshot immediately before the one currently displayed.
	Play (plus Playback speed). Activate playback mode. This automatically steps through the snapshots in sequence. **Note**: Click the down arrow to select a playback speed. Select **Custom** to define a playback speed.

Control	Description
	Next snapshot. Show the snapshot immediately after the one currently displayed.
	Next alarm. Show the snapshot corresponding to the next alarm after the one currently displayed.
	Select date/time. Click to nominate the date and time to display. Spotlight will find the snapshot closest to this time to display.

Notes:

- While you are viewing historical data the connection identifier at the top right of the window indicates the date and time of the snapshot. On the Spotlight home page the console status bar indicates "History Browser".

- An easy way to view the history for a particular alarm is through the Alarm Log. Right-click the alarm and select **Show Selected Entry In | History Browser**. Alarm log

- To view the recent history of a Spotlight home page component, right click the component and select **Show History**. Spotlight Home Page Components

- Spotlight History is installed if you click **Help | About Spotlight | Spotlight Modules** and under the heading **Spotlight Console** see **Light-weight XML repository**. You can customize the storage of history data. Configure Spotlight - View | Options

Related Topics

Spotlight Browser

Welcome to Spotlight

Live connections

History Browser

Spotlight Home Pages

The Spotlight Home Page is the main page for the monitored system. It highlights obvious bottlenecks and problem areas. Statistics and flows are updated in real time.

To see the Spotlight Home Page

1. Select a connection from the Spotlight Browser.

2. Click Home.

About the Spotlight home page

- The connection identifier in the upper right corner of the Spotlight console identifies the system being monitored. Alongside the connection identifier there may be additional icons. The following activity icon indicates whether Spotlight is collecting data from the connected system.

- The home page can show real (live) data or a historical snapshot. When the home page is showing a historical snapshot the words **History Browser** are written in the bottom left corner of the console and the recorded date and time are shown in the top right corner. History Browser

- A panel is a visual feature that groups related components on the Spotlight home page. Panels may be connected by data flows that illustrate the rate at which the system is performing. The groupings reflect how your system works. Components change color as alarms are raised for the metrics they contain. Spotlight Home Page Components

Related Topics

Spotlight Home Page Components

Calibrate Spotlight Connections

Spotlight Drilldown Pages

A drilldown contains one or more pages that display a detailed breakdown of the monitored system's performance.
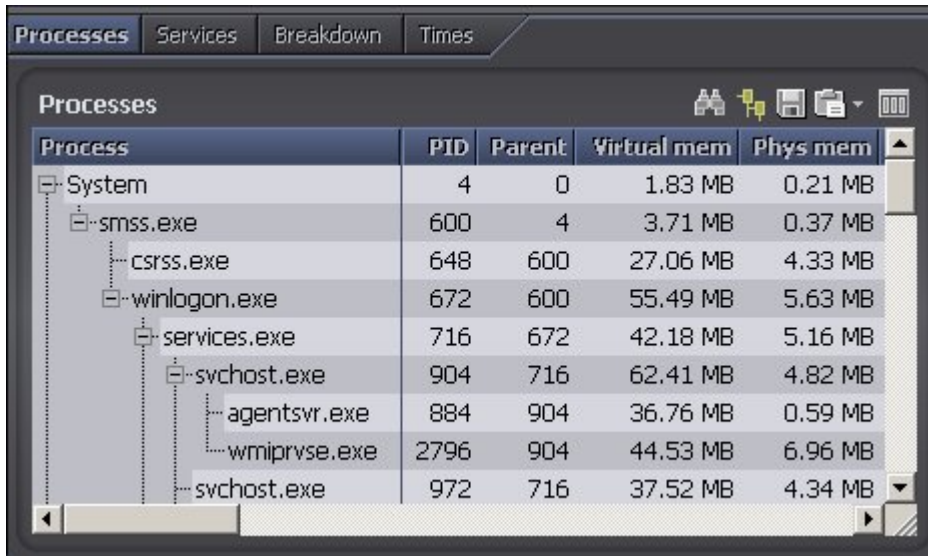
To open a drilldown

1. Select a connection from the Spotlight Browser Live connections or History Browser.

2. Click the relevant drilldown button on the Spotlight toolbar. Drilldown buttons are dependent on the Spotlight application. For help, click **Help | Contents |** *Connection type* **| Drilldown**.

About Spotlight drilldowns

Drilldowns mainly use two different types of displays – tables and charts.

About Spotlight Grids

About Spotlight Charts

Drilldowns have the following features:

- They can be configured to show all or some of the metrics associated with components.

- You can access further information about displays in drilldowns by moving the mouse over the displays, or by clicking or right-clicking on them.

- You can copy the data shown in drilldowns to other applications or save it to a file.

Related Topics
About Spotlight Grids
About Spotlight Charts

Spotlight Navigation

Return to the previous page or drilldown in the browse sequence

Click the down arrow to select the page to return to.

Note: Back is disabled at the start of the browse sequence.

Toolbar	Menu	Shortcut
Click	Click **View	Back**.	Press ALT+LEFT

View the next page or drilldown in the browse sequence

Click the down arrow to select the page to go to.

Note: Forward is disabled at the end of the browse sequence.

Toolbar	Menu	Shortcut
Click	Click **View	Forward**.	Press Alt+Right

Return to the Spotlight home page

Toolbar	Menu	Shortcut
Click	Click **View	Close Drilldown**.	N/A

Pause / Resume / Refresh

Pause monitoring the current Spotlight connection. Spotlight stops retrieving data from the connection. The Spotlight home page and drilldowns for the connection are not updated. Use the Resume command to resume monitoring the Spotlight connection. While the connection is paused, use Refresh to update the current window display.

Notes:

- Any actions that are in progress when you pause Spotlight are allowed to complete. No further actions are performed until you resume Spotlight's collection of data.

- Pausing may affect Calibration. If calibration is in progress when the pause is started, it continues to run. The calibration finishes at the time it was originally set to finish. (It is not delayed by the Pause command.)

- Pausing may affect drilldowns. Any drilldown that gathers information when it is started is affected by the Pause command.

To Pause Spotlight

Toolbar	Menu	Shortcut
Click	Click **View	Pause**.	N/A

To Resume Spotlight

Toolbar	Menu	Shortcut
Click	Click **View	Resume**.	N/A

To Refresh (update) the current Spotlight window

Use while Spotlight is paused to refresh the current window manually.

Toolbar	Menu	Shortcut
Click	Click **View	Refresh**.	Press **F5**.

Note: Alarms raised during refresh may be included in historical snapshots, depending on snapshot preferences.

Related Topics

Set the rate at which data is collected

Calibrate Spotlight Connections

Spotlight Favorites

Add frequently viewed pages and connections to the favorites menu.

To view a favorite

View a page you have added to the favorites menu.

1. Select **Favorites**.

2. Select the favorite.

Note: The favorite may be in a folder as per how you have organized your favorites. The display name of the favorite is a name of your choosing.

To add a favorite

1. Navigate to the page you want to add to favorites.

2. Select **Favorites | Add to Favorites**.

3. How do you want the page to appear in the Favorites list?

Option	Description
Name	Give a display name to the favorite.
Save the current connection as part of the favorite	Select when the same display name is used for favorite pages for different connections.
Create in	Select the folder to put the new favorite into. **Tip**: Click **New folder…** to create a new folder to save the favorite in.

To organize favorites

1. Select **Favorites | Organize Favorites**.

2. Click and drag entries in the **Favorites** list to rearrange them.

Notes:

- Select an item then click **Delete**, **Rename** or **Move to...** another location in the favorites list.
- Click **New folder...** to create a new folder to organize the favorites in.

Spotlight Slideshow

Spotlight, Slideshow cycles through the current page of Spotlight's live connections.

Select **View | Slideshow**.

Note: You can configure Spotlight to start slideshow mode when you are not actively using the application. Slideshow

Connection Status bar

The Connection status bar is a floating bar that contains a Spotlight icon for each active (live) connection.

To show / hide the Connection Status bar

Click **View | Connection Status Bar**

About the Connection Status bar

- The color of the connection icon represents the status of the connection.
- Click the connection icon to display the connection's Spotlight home page.
- Click and drag the checkered handle on the left to resize the bar.
- Right-click the bar to display its shortcut menu. The following options are available:

Option	Description
Transparent	"Fades" the bar when the mouse pointer is not placed over it.
Text labels	Displays the names of the connections in the status bar.
Automatic resize	Resizes the status bar automatically and activates a scroll arrow.
Sort by	Sorts the connection icons by Name or Severity.

Related Topics

Spotlight Connection Manager

Calibrate Spotlight Connections

Calibrate the Spotlight home page components to display data in the way best suited to the monitored system.

Note: Spotlight automatically calibrates new connections the first time that they are opened. Use this dialog to re-calibrate a connection at any other time.

Scenario: Calibration of data flows

Data flows on the Spotlight home page indicate the rate at which data is being transferred through a system. The faster the rate of traffic, the faster the data flow moves. However, the data flow is an accurate representation of system activity ONLY if Spotlight knows the normal range of values for your system.

If the normal range of values for a data flow is from 0 to 100, a value of 8 is low, and the data flow moves slowly.

If the range of values for your system is from 0 to 10, a value of 8 is high and the data flow moves much faster.

The Spotlight calibration tool calculates the normal range of metric values for your system. Over a set period of time, Spotlight measures the changes in value of important metrics within the system. It uses this data to set the upper and lower display limits for the corresponding components. You can accept the results or adjust them if necessary.

To re-calibrate a connection

1. Select / Open the connection from the Spotlight Browser Live connections.

2. Select **File | Calibration**.

3. Set the Calibration period of time.

4. Click **Start**.

Notes:

- To stop calibration select **File | Calibration | Stop**.

- You can set Spotlight to save calibration data when saving configuration data in a Spotlight template. Connection Templates

Related Topics
Spotlight Home Page Components

Spotlight Alarms

Alarms are the warnings that Spotlight raises when a metric falls outside its "normal" range of values. A new alarm is raised whenever the severity for a metric changes. When the severity returns to normal, the alarm is canceled.

Alarms are recorded in the Spotlight alarm log pages:

- Spotlight Today - Alarms according to priority.

- Alarms by Time - Alarms according to duration and severity.

- Alarm log - Alarms in a sorted table.

Many Spotlight elements use color to indicate severity.

High Medium Low Normal Information Disabled

You can configure Spotlight to notify you of changes to severity through color or intensity, audible warnings, sending email or network notifications or running a program.

Spotlight elements whose color represents severity	Color indicates	What to do
The open/close button on the Spotlight browser	The highest severity currently raised - across all Live	Locate the connection in the Spotlight Browser (Live connections). Open the connection's Home Page, Alarm Log or Alarms by Time.

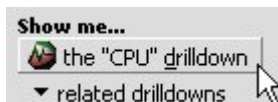| Spotlight elements whose color represents severity	Color indicates	What to do
The system tray icon The connection name in the Spotlight browser (Live connections) and connection status bar	(monitored) connections. The highest severity currently raised against the connection.	Locate the connection in the Spotlight Browser (Live connections). Open the connection's Home Page, Alarm Log or Alarms by Time.
The component on the Spotlight home page	The highest severity currently raised against the component.	Click on the component to open help regarding the alarm. The text in the help window gives some details about the alarm, why it was raised, and ways in which the problem indicated can be resolved. Use the help to open the drilldown page that is directly related to the alarm. Drilldown pages contain charts and tables that you can use to diagnose the cause of the alarm. You may like to view one or more related drilldown pages that may be of some value in diagnosing the problem. Use the information shown in the help window, together with the details shown in the associated drilldown page(s), to discover why Spotlight has raised the alarm. If the alarm is transient, you at least will see where in the system it occurred. If the alarm persists, Spotlight indicates what measures need to be taken to resolve the problem. **Note**: These actions are customizable. Balloon help
The entry in the alarm log page	The severity of the alarm.	Alarms are recorded in the Spotlight alarm log pages: • Spotlight Today - Alarms according to priority. • Alarms by Time - Alarms according to duration and severity.

Spotlight elements whose color represents severity	Color indicates	What to do
		• Alarm log - Alarms in a sorted table.

Historical alarms

For an alarm raised in the past, view a snapshot of the system as a whole at the time the alarm was raised.

1. Open the Alarm Log, Alarms by Time or Spotlight Today.

2. Right-click the alarm and select **Show Selected Entry In | History Browser**.

3. History Browser

Customize the way that Spotlight handles alarms:

Option	Description
Set new thresholds	You can edit the metric and measures by which an alarm is raised in the Metrics Editor. About Alarms, Metrics, Thresholds And Severities
Set action rules	Govern Spotlight behavior when an alarm is raised. Alarm actions
Set filters	Govern which alarms are displayed. Alarm Filters
Ignore specific alarms	Ignore (Snooze) Alarms

Spotlight Today

Spotlight Today shows recent alarms for all Spotlight connections. Spotlight Today can be grouped or filtered to show alarms by connection or severity. Spotlight Today displays: the number of current Spotlight connections, the number of connections for which an alarm has been raised, a Spotlight Today entry for each alarm raised.

To see Spotlight Today

Toolbar	Menu	Shortcut
Click	Click **View	Spotlight Today**.	N/A

About the Spotlight Today grid

Column	Description
Severity	The degree of urgency of the alarm.
Alarm	The name of the alarm. **Note:** This column is unavailable on **Group By	Alarm**.
Connection	The name of the Spotlight connection that raised the alarm. **Note:** This column is unavailable on **Group By	Connection**.
Raised	The time the alarm was raised.
Message	More information on the alarm.
Connection type	The type of Spotlight connection that raised the alarm (for example, Spotlight on Oracle).
Last updated	The last time the alarm was updated.
Occurrences	How often the alarm has occurred in the monitoring period.

Actions on Spotlight Today

Click	Description
	Find Text
	Alarm Filters - Filter by severity and connection. Show snoozed alarms.
	Save To File
	Copy To Clipboard
	Show, Hide & Order Columns - Some of the columns in the grid may be hidden by default. Use the Column Organizer to select the columns you want to see.
	Alarm actions - Define actions to perform when an alarm meets pre-defined conditions.
Right click the header row and select **Arrange by** | *Column name*	Sort the Grid - Sort the grid according to the contents of a column
Right click the header row and select **Freeze First column**	Grid Properties - Options Page - Keep the first column visible while you scroll through the columns of the table.
Right click the content and select **Group By**	Group By | Connection - Arrange the grid in a tree like structure by connection.
	Group By | Alarm - Arrange the grid in a tree like structure by alarm.
	Group By | Do Not Group - List the alarms.
Right click the grid content and select **Collapse** or **Expand**.	Collapse / Expand the tree structure. Applicable on **Group By | Connection** and **Group By | Alarm**.
Right click an alarm and select **Open Connection**	Open the Spotlight Home Page for the selected connection.
Right click an alarm and select **Properties**	Open the Connection Properties dialog for the selected connection. Connection Properties - Details Page

Related Topics

Alarms by Time

Alarm log

Alarms by Time

Alarms by Time displays the alarms that have been logged for your Spotlight connections. It shows (in connection order) the start time, duration, and severity of the logged alarms.

To see alarms by time

Toolbar	Menu	Shortcut
Click	Click **View | Alarms by Time**.	N/A
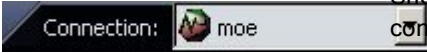
About alarms by time

Feature	Description
Time scale	This spans the top of the chart from left to right, and indicates when alarms (represented as colored horizontal bars) were raised, changed, or canceled.
List of connections and alarms	The vertical axis of the chart contains a tree structure. The roots of the tree represent your Spotlight connections; the branches represent the alarms that have been raised for each connection. The order of the alarm branches in each connection changes to reflect the alarms that have occurred most recently.
Severity of	The color of the horizontal alarm bars change to reflect the severity of the relevant alarm. In the

Feature	Description
alarms	illustration the Threads alarm was raised with High severity, and thereafter downgraded in a series of steps to Information status, and then canceled.

Actions on alarms by time

Click	Description
Connection: moe	Show alarms by time for the selected connection(s). Select a connection, all connections of a given type, a customized selection of connections or all connections.
	Find Text
	Alarm Filters - Filter by severity and age. Show snoozed alarms.
10 minutes	Zoom in and out of the displayed time. Click the down arrow to choose a suitable time scale.
	Move to the Previous / Next time the alarm was raised, upgraded, downgraded, or canceled. **Note**: Select an alarm to enable Previous / Next.
	Click to return to the bookmarked time. **Note**: To bookmark a moment in time, click in the body of the grid. The vertical line indicates the bookmarked time.
	Alarm actions - Define actions to perform when an alarm meets pre-defined conditions.
Right click the grid content and select **Collapse** or **Expand**.	Collapse / Expand the tree structure. The name of the connection is at the top of the tree. Logged alarms are listed for each connection.
Right click the Alarm header and select **Sort Ascending** or **Sort Descending**.	Sort connections alphabetically. For any expanded roots of the tree, the alarms are similarly sorted. **Tip**: To re-sort alarms by time, collapse then expand that connection in the tree.

Related Topics

Spotlight Today

Alarm log

Alarm log

The Alarm Log page displays a table of logged alarms for your Spotlight connections.

To see the alarm log
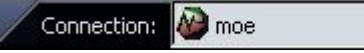
Toolbar	Menu	Shortcut
Click	Click **View	Alarm Log**.	CTRL+L

About the alarm log grid

Column	Description
Time	The time at which the alarm was raised or canceled.

Column	Description
Connection	The name of the Spotlight connection that raised the alarm.
Connection Type	The type of Spotlight connection that raised the alarm (for example, Spotlight on Oracle).
Action	The alarm action that was performed. **Note**: Snoozed alarms are indicated by 🔔. Ignore (Snooze) Alarms
Details	A brief description of the cause of the alarm.
Severity	The degree of urgency of the alarm.
Alarm	The name of the alarm.
Help	Click ❓ (if available) to view more information about the alarm and help to resolve it. Spotlight Alarms

Actions on the alarm log

Click	Description
Connection: 🔴 moe	Show the alarm log for the selected connection(s). Select a connection, all connections of a given type, a customized selection of connections or all connections.
🔍	Find Text
▼	Alarm Filters - Filter by severity and age. Show snoozed alarms.
💾	Save To File
📋	Copy To Clipboard
▥	Show, Hide & Order Columns - Some of the columns in the grid may be hidden by default. Use the Column Organizer to select the columns you want to see.
🎬	Alarm actions - Define actions to perform when an alarm meets pre-defined conditions.
Right click the header row and select **Arrange by | *Column name***	Sort the Grid - Sort the grid according to the contents of a column
Right click the header row and select **Freeze First column**	Grid Properties - Options Page - Keep the first column visible while you scroll through the columns of the table.
Right click an alarm and select **Show Selected Entry In | History Browser**	View a historical snapshot of the selected alarm. Spotlight opens the relevant snapshot / the home page for the connection / in History mode. History Browser
Right click an alarm and select **Show Selected Entry In | Alarms by Time**	Alarms by Time - Open the selected alarm in Alarms by Time.

Related Topics

Spotlight Today

Alarms by Time

Alarm Filters

Filter the contents of the Alarm Log pages (Spotlight Today, Alarms by Time and the Alarm Log).

The performance of Spotlight generally improves when you choose a smaller period of time to view.

To filter alarms

1. Open one of the Alarm Log pages.

   Click one of: 🔲🔲🔲.

2. Click 🔲.

3. Select from the following:

Option	Description
Show All	Do not filter.
Filter by severity	Filter by the value of the alarm severity.
Filter by age	Filter by the age of the alarm. **Notes**: • Not applicable to Spotlight Today. • Select a predefined time period or the **Custom age filter** to set a customized time period.
Filter by Connection	Filter by the connection that raised the alarm. **Note**: Applicable to Spotlight Today.
Show / Hide Snoozed Alarms	Ignore (Snooze) Alarms
Modify Custom Filters	Specify filter rules that govern when alarms are displayed. Custom Filter Rules Dialog
Restore default filters	Reset modified filters to their default values.

Related Topics
Spotlight Today
Alarms by Time
Alarm log

Ignore (Snooze) Alarms

You may choose to snooze an alarm when:

- A temporary situation causes an expected alarm whose existence is unimportant.

- There is a known problem (a network problem, for example) that is beyond your control; and a continuing reminder of the problem is neither necessary nor desirable.

About Snoozed Alarms

- Alarms that are snoozed still contribute to the Alarm Log.

- Alarms that are snoozed still trigger historical snapshots, and are shown in the History Browser.

- Alarms that are snoozed do not fire Alarm Log rules or take any associated actions.

- An entry is written to the Alarm Log when a snoozed alarm is "unsnoozed".

Snooze

Alarms are snoozed from the Spotlight home page.

Requirement	Procedure
To snooze an alarm	1. Right click the home page component with the raised alarm and select **Snooze alarm**. 2. Select the period of time: from five minutes to the duration of the Spotlight session. The snooze icon is added to the home page component and alarm log entry. **Note**: During the snoozed period, alarms on ALL metrics for this home page component are ignored.
To restore snoozed alarms	Right-click the snoozed home page component and select **Snooze alarm**. The snooze icon is removed from the home page component and alarm log entry.
To show only snoozed alarms	Use the Alarm Log Filter dialog. Custom Filter Rules - Alarm Log Filter Dialog This filter will recognize only components whose alarm status has changed during the snooze period. Other snoozed alarms will not be recognized. For example, if an alarm on the snoozed component was raised before the snooze period, and if that alarm status does not change during the snooze period, the alarm will not appear in the filtered Alarm Log.

Related Topics

Spotlight Home Page Components

Alarm log

Custom Filter Rules Dialog

Specify filter rules that govern when alarms are displayed in the Alarm Log pages.

To open the Custom Filter Rules dialog

1. Open one of the Alarm Log pages.

   Click one of: .

2. Click | **Modify Custom Filters**.

About the dialog

Option	Description
Show alarm entries when the following conditions are met	Rules are listed in the order they are applied. De-select a rule to retain it but not apply it. Highlight a rule to view, modify, delete, rename or move it.
Rule description	View the highlighted rule.
New	Click to create a rule. Custom Filter Rules - Alarm Log Filter Dialog
Modify	Highlight a rule. Click to modify the highlighted rule. Custom Filter Rules - Alarm Log Filter Dialog
Delete	Highlight a rule. Click to delete the highlighted rule.
Rename	Highlight a rule. Click to rename the highlighted rule.
Move Down	Highlight a rule Click to move it lower down the list. It is applied later.

Option	Description
	Rules are applied in the order they are listed.
Move Up	Highlight a rule. Click to move it higher up the list. It is applied sooner. Rules are applied in the order they are listed.

Related Topics

Alarm Filters

Spotlight Today

Alarms by Time

Alarm log

Custom Filter Rules - Alarm Log Filter Dialog

Create and Modify the filter rules that govern when alarms are displayed in the Alarm Log pages.

To open the dialog

1. Open the **Custom Filter Rules** dialog. Custom Filter Rules Dialog

2. Click **New** or **Modify**.

To fill in the dialog

1. Select one or more conditions. Selected conditions are transcribed to the **Rule description** window.

2. Refine the rule description by clicking the underlined links as appropriate. Underlined links in red must be configured to acceptable values for the rule to be accepted.

Additional information to the Rule description

- Where the action is…
- Where the alarm has (not) been snoozed
- Where the alarm severity is...
- Where the connection is...
- Where the connection type is...
- Where the control is...
- Where the details contain…
- Where the highest alarm for the connection is...
- Where the value is...

Related Topics

Custom Filter Rules Dialog

Alarm Filters

Spotlight Today

Alarms by Time

Alarm log

Alarm Actions - Alarm Log Filter Dialog

Create and Modify rules that govern action taken on an alarm severity.

To open the dialog

1. Open the **Alarm Actions** dialog. Alarm actions

2. Click **New** or **Modify**.

To fill in the dialog

1. Select one or more conditions which when met will trigger one or more actions. The conditions are transcribed to the **Rule description** window.

2. Select the actions to be performed when the conditions are met. The actions are transcribed to the **Rule description** window.

3. Refine the rule description by clicking the underlined links as appropriate. Underlined links in red must be configured to acceptable values for the rule to be accepted.

Additional information to the Rule description - Conditions

- Where the action is…
- Where the alarm has (not) been active in the last n minutes
- Where the alarm has (not) been snoozed
- Where the alarm severity is...
- Where the connection is...
- Where the connection type is...
- Where the control is...
- Where the details contain…
- Where the highest alarm for the connection is...
- Where the value is...

Additional information to the Rule description - Actions

- Play a sound
- Run a program
- Send e-mail to…
- Send network message to…
- Stop processing further rules - Exit the filter rule when this condition is reached. Ignore subsequent commands.

Related Topics

Alarm actions

Configure Spotlight - View | Options

Where the action is…

Define appropriate actions.

How to open this dialog
From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the action is…**

3. Click to define the action in the **Rule description** pane.

Where the action is any of the following:
Select all appropriate actions.

Value	Description
Alarm raised	When the severity of a metric rises above normal.
Alarm upgraded	When the severity of a metric increases.
Alarm downgraded	When the severity of a metric decreases.
Alarm canceled	When the severity of a metric returns to normal.
Alarm snoozed / unsnoozed	If you choose to ignore (snooze) a triggered alarm for a specified period of time, the Spotlight Alarm Log records both when the alarm is snoozed, and when it is unsnoozed. Ignore (Snooze) Alarms
Information message logged	When Spotlight logs an information message.
Warning message logged	When Spotlight logs a warning message. Warning messages are more urgent than information messages.
Error message logged	When Spotlight logs an error message. Error messages are more urgent than information messages or warning messages.
Log opened	When Spotlight opens the Alarm Log.
Log closed	When Spotlight closes the Alarm Log.

Related Topics
Custom Filter Rules - Alarm Log Filter Dialog
Alarm Actions - Alarm Log Filter Dialog

Where the alarm has (not) been active in the last n minutes

Use to ignore transient alarms.

Some Spotlight alarms may last only a few seconds or minutes, and are not significant for the system under investigation. When these transient alarms occur frequently, they can even be counter-productive: they fill the Alarm Log with useless data; and they can obscure other alarms that are more important.

How to open this dialog

From the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the alarm has (not) been active in the last n minutes...**

3. Click to define the action in the **Rule description** pane.

Where the alarm has (not) been active in the last n minutes:

In the **minutes** box, enter or select the amount of time that Spotlight will use as a threshold value.

Related Topics

Alarm Actions - Alarm Log Filter Dialog

Where the alarm has (not) been snoozed

Request Spotlight display only ignored alarms or those that are not being ignored.

How to open this dialog

From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the alarm has (not) been snoozed**

3. Click to define the action in the **Rule description** pane.

Where the alarm has (not) been snoozed:

Select **Has been snoozed** or **Has not been snoozed**.

Ignore (Snooze) Alarms

Related Topics

Custom Filter Rules - Alarm Log Filter Dialog
Alarm Actions - Alarm Log Filter Dialog
Ignore (Snooze) Alarms

Where the alarm severity is...

Define the appropriate severity to use in an Alarm Log rule. Severity is the degree of urgency of an alarm.

How to open this dialog

From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the alarm severity is...**

3. Click to define the action in the **Rule description** pane.

Where the alarm severity is any of the following:

1. From the list of operators, select the operator that you want to use.

2. From the drop-down list, choose the severity you want to use.

Note: If you want the Spotlight connection's alarm with the highest severity to trigger the Alarm Log rule then use Where the highest alarm for the connection is...

Related Topics
Custom Filter Rules - Alarm Log Filter Dialog
Alarm Actions - Alarm Log Filter Dialog

Where the connection is…

Select the connections to (not) apply the rule to.

How to open this dialog
From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the action is…**

3. Click to define the action in the **Rule description** pane.

Where the connection is any of the following:
Select **Is one of the following** or **Is not one of the following**.
From the tree list, select all appropriate Spotlight connections. (The top level in the tree is the connection type.)

Related Topics
Custom Filter Rules - Alarm Log Filter Dialog
Alarm Actions - Alarm Log Filter Dialog

Where the connection type is…

Select the connection types to (not) apply the rule to.

How to open this dialog
From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the connection type is…**

3. Click to define the action in the **Rule description** pane.

Where the connection type is any of the following:
Select **Is one of the following** or **Is not one of the following**.
Select all appropriate Spotlight connection types.

Related Topics

Custom Filter Rules - Alarm Log Filter Dialog
Alarm Actions - Alarm Log Filter Dialog

Where the control is...

Select the Spotlight control (component) to (not) apply the rule to.

How to open this dialog
From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the control is...**

3. Click to define the action in the **Rule description** pane.

Where the control is any of the following:

1. Select **Is one of the following** or **Is not one of the following**.

2. From the Plug-in list, select the Spotlight plug-in application (Spotlight on Windows, for example) that contains the control.

3. From the Drilldown list, select the Spotlight page that contains the control.

4. From the Control list, select the specified control.

Related Topics
Custom Filter Rules - Alarm Log Filter Dialog
Alarm Actions - Alarm Log Filter Dialog

Where the details contain...

Require specific values to (not) apply the rule to.

How to open this dialog
From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the details contain...**

3. Click to define the action in the **Rule description** pane.

Where the details contain any of the following:

1. Select from

    - contain any of the following phrases

    - contain all of the following phrases

    - do not contain any of the following phrases

2. As appropriate:

- Click **Add** to create a phrase.

- Highlight a phrase and click **Modify** to change it.

- Highlight a phrase and click **Delete** to delete it.

Related Topics

Custom Filter Rules - Alarm Log Filter Dialog

Alarm Actions - Alarm Log Filter Dialog

Where the highest alarm for the connection is…

An active Spotlight connection may raise more than one alarm during a session, and each alarm raised may have a different level of severity. Use this condition when you want ONLY the alarm with the highest severity to apply to the Alarm Log rule.

How to open this dialog

From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the highest alarm for the connection is…**

3. Click to define the action in the **Rule description** pane.

Where the highest alarm for the connection is:

1. From the Where the severity is options, select the operator that you want to use.

2. From the drop-down list, choose the severity you want to use.

Note: When you want the severity of a specified alarm (not necessarily the highest alarm) to trigger an Alarm Log rule, use Where the alarm severity is...

Related Topics

Custom Filter Rules - Alarm Log Filter Dialog

Alarm Actions - Alarm Log Filter Dialog

Where the value is…

Define as appropriate to the rule.

Note: This condition is used in combination with other conditions. An example of this may be a rule with two conditions, as in: "Where the control is... and Where the value is..."

How to open this dialog

From the Custom Filter Rules - Alarm Log Filter Dialog or the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Where the value is…**

3. Click to define the action in the **Rule description** pane.

Where the value is any of the following:

1. Select the appropriate operator.

2. Select the appropriate value

Related Topics

Custom Filter Rules - Alarm Log Filter Dialog

Alarm Actions - Alarm Log Filter Dialog

Play a sound

Define the sound to play when the alarm condition is met.

How to open this dialog

From the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Play a sound**

3. Click to define the action in the **Rule description** pane.

Play the sound:

1. Click to browse to the sound file on the computer.

   Tip: Click **Preview** to play the sound to ensure you have selected the right one. The success (or otherwise) of the action depends on your system setup. Spotlight does not monitor the success of the action. We recommend that you preview it.

2. Select the desired Play sound option.

Option	Description
Once	Play once.
Until the severity of the control ... returns to normal	Play until the effected Spotlight control no longer registers an alarm.
Until the connection severity returns to normal	Play until the Spotlight connection no longer registers an alarm.
Until the following condition is met.	Play until the condition is met.

Tip: When an alarm has been triggered and the relevant sound plays continuously, you can mute the sound by clicking . This button is not shown on the toolbar unless a **Play sound** action is enabled.

Related Topics

Alarm Actions - Alarm Log Filter Dialog

Run a program

Define the program (executable file) to run when the alarm condition is met.

Usually that program is one that you have written to perform a particular task. The task may be a simple action, or may be one that depends on a supplied Spotlight parameter – a value or range of values, or the name of a component or connection.

How to open this dialog

From the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Run a program**

3. Click to define the action in the **Rule description** pane.

Run program:

1. Click to browse to the sound file on the computer.

   Tip: Click **Preview** to play the sound to ensure you have selected the right one.

2. Select the desired Play sound option.

Option	Description
Program	Click to browse to the program file on the computer.
Parameters	Enter any parameters the program may need to run successfully, or click **{}** to choose a parameter from the supplied list.
Test	Click to test the performance of the selected program. **Note**: The success (or otherwise) of the action depends on your system setup. Spotlight does not monitor the success of the action. We recommend that you test it.

Caution: If the program that runs in response to an alarm is still executing when the alarm fires again, subsequent commands to run the program are ignored until that execution finishes.

Related Topics

Alarm Actions - Alarm Log Filter Dialog

Send e-mail to…

Send an email when the alarm condition is met.

How to open this dialog

From the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Send e-mail to…**

3. Click to define the action in the **Rule description** pane.

Send e-mail to:

Option	Description
To	Enter the recipient email address.
Subject	Supply a subject for the email. **Tip**: Click **{}** to select a parameter from the supplied list. The selected parameter is displayed within a set of braces {}. You can choose multiple parameters if required.

Option	Description
	Enter the message to send.
Message	**Tip**: Click **{}** to select a parameter from the supplied list. The selected parameter is displayed within a set of braces {}. You can choose multiple parameters if required.
	Click to test the performance of your email system.
Test	**Note**: The success (or otherwise) of the action depends on your system setup. Spotlight does not monitor the success of the action. We recommend that you test it.

Note: For the email message to reach its intended recipient, an appropriate email client must be installed correctly on both source and target machines.

Related Topics

Alarm Actions - Alarm Log Filter Dialog

Send network message to…

Send a message to a machine on the network using the Windows messaging service.

How to open this dialog

From the Alarm Actions - Alarm Log Filter Dialog

1. Click **New** or **Modify.**

2. Select **Send network message to…**

3. Click to define the action in the **Rule description** pane.

Send network message to to:

Option	Description
Address	Enter the name of the machine where you want to send the message.
Message	Enter the message to send.
	Tip: Click **{}** to select a parameter from the supplied list. The selected parameter is displayed within a set of braces {}. You can choose multiple parameters if required.
	Click to test the performance of the network messaging system.
Test	**Note**: The success (or otherwise) of the action depends on your system setup. Spotlight does not monitor the success of the action. We recommend that you test it.

Note: For the message to reach its intended recipient, the Windows Messaging service must be active on both the source and target machines on the network.

Related Topics

Alarm Actions - Alarm Log Filter Dialog

About Alarms, Metrics, Thresholds And Severities

Terms	Definitions	Customizations and Actions
alarm	Alarms are the warnings that Spotlight raises when a metric falls outside its "normal" range of values.	A new alarm is raised whenever the severity for a metric changes. When the severity returns to normal, the alarm is canceled. You can see this behavior clearly in the Alarm log. Alarm log

Terms	Definitions	Customizations and Actions
metric	A metric is a piece of information that Spotlight collects about the performance of a system. The information may be a numeric value (a number or percentage), a string of text, or some other piece of data. Every time the Spotlight window is refreshed, Spotlight retrieves the latest value of the metric, which can then be displayed on the home page.	Use the Metrics Dialog to: • Select the metrics to collect for the connection / component. • Configure Alarm smoothing, to smooth out the anomalies and spikes that may appear momentarily in metric values. The purpose is to prevent Spotlight from reporting alarms for spikes that do not persist.
threshold	For each numeric metric, you can set thresholds – ranges of values – that indicate levels of severity for that metric.	Use the Thresholds Dialog to: • Select the metrics to collect for the connection / component. • Define the severities to apply to the metric. • Define the start and end values for each severity.
severity	A severity can be used to specify whether the information returned in the metric represents normal or abnormal behavior for the system. Within Spotlight, there can be at most the following types of severity: • Disabled • Normal • Information • Low • Medium • High "Normal" indicates that the system is performing within acceptable limits. If a metric returns a value with any other severity, Spotlight raises an alarm that indicates the system is behaving outside acceptable limits. A "Disabled" status means the system is not responding, and no information is being returned.	When an alarm is raised Spotlight performs one or more of the following actions: • Changes the color or intensity of relevant components. • Gives audible warnings. • Runs a program. • Sends email or network notifications. You can control how that notification is issued. Alarm actions

Related Topics

Spotlight Alarms

Metrics Dialog

Thresholds Dialog

Metrics Dialog

View the alarms defined for a Spotlight connection or component of the Spotlight home page.

Note: These are not the alarms raised for a connection / component. To see the alarms raised for a connection open the Alarm log.

To open the Metrics dialog

View all metrics for the connection	View all metrics for the Spotlight home page component
Open the Spotlight home page for the connection.	Right click the component on the Spotlight home page and select **Metrics…**
Select **View |Show All Alarms…**	

About the Metrics dialog

Option	Description
List of metrics	All metrics associated with the component / connection are listed. Data is collected for selected metrics. Clear the selection to stop collecting data for the metric. Highlight the metric to display its details in the **Metric** page. You can sort the list. • Click the header to sort the metrics in ▲ ascending or ▼ descending name order. • Click the severity icon in the header to sort the metrics according to severity 🔴🟡🟠🔴.
Metric Name	This is the metric highlighted in the list of metrics.
Averaging	Set Averaging for the highlighted metric. • Use default - Every value used to calculate a metric's alarm status is obtained by averaging the metric over a period of time specified in Spotlights Options. • Don't average this metric - Use the raw metric data retrieved from the current system – anomalies and spikes included. • Use custom averaging - Use a moving average, as in the Use default... case, but not the default setting. Set the time period by using the Average over... slider. • Average over... - Use the mouse or arrow keys to move the slider and so choose the time period used to generate the moving average. The longer the time period, the fewer the peaks and troughs. Spikes and other anomalies are smoothed. **Tip**: Do not to use too long a time to calculate the moving average. This may prevent the reporting of valid alarms. The maximum time period that Spotlight uses for averaging is five minutes. Averaging - is a technique that Spotlight uses to smooth out the anomalies and spikes that may appear momentarily in the metric values for a component. The purpose of averaging is to prevent Spotlight from reporting alarms for such spikes if they do not persist. **Notes**: • Averaging is used ONLY to prevent the generation of spurious alarms. It does NOT affect the charts or tables that display metric values; these continue to display metrics as they are reported. • This is not the place to set averaging for all metrics for the connection. Alarm smoothing

Related Topics

Thresholds Dialog

Define the acceptable range of values for a Spotlight metric. Metrics that exceed the Normal threshold raise an alarm.

To open the Thresholds Dialog

View all metrics for the connection	View all metrics for the Spotlight home page component
Open the Spotlight home page for the connection.	Right click the component on the Spotlight home page and select **Metrics…	Thresholds**
Select **View |Show All Alarms…	Thresholds**	

About the Thresholds Dialog

Option	Description
List of metrics	All metrics associated with the component / connection are listed.
	Data is collected for selected metrics. Clear the selection to stop collecting data for the metric.
	Highlight the metric to display its details in the **Thresholds** page.
	You can sort the list.
	• Click the header to sort the metrics in ▲ ascending or ▼ descending name order.
	• Click the severity icon in the header to sort the metrics according to severity ⦿⦿⦿⦿.
Name	This is the name of the metric highlighted in the list of metrics.
Severity table	Define severities for the highlighted metric.
	• Select the severities you wish to define.
	Define Start values for each selected severity either in the Severity Table or on the Threshold Bar. In the Severity Table: Click in the Start cell. Enter the lower value for the threshold.
	Note: Ensure the start values for each range are consistent. A failure to do so may cause problems when alarms are raised.
	• Deselect the severities you do not wish to define.
Threshold bar	Severity Start and End values can be defined in the Severity table or the Threshold bar. On the Threshold bar:
	1. Move the mouse pointer over the threshold between two severities until it appears as a split-bar pointer.
	2. Drag the mouse to move the threshold to a new position. Release.
	3. Changes are updated automatically in the Severity table.

Related Topics

Spotlight Home Page Components

Metrics Dialog
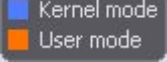
About Spotlight Charts

A chart is a component on a Spotlight page that shows data as a graphical image.
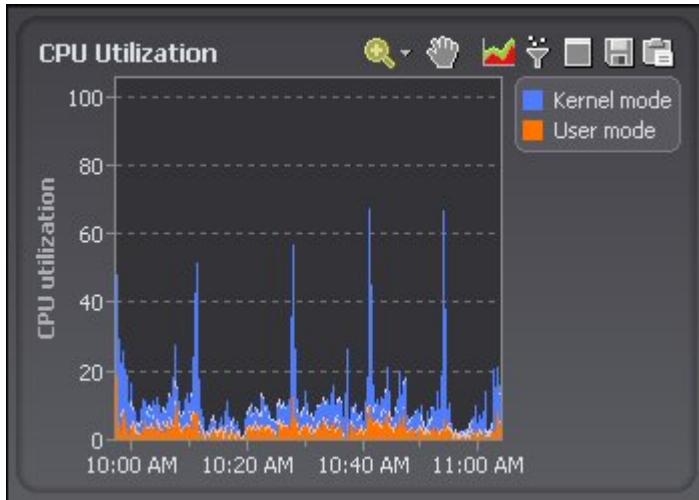
Notes:

- If there is a down arrow alongside the chart title then the chart supports more than one category of data. Click the title or arrow to open a drop-down list of all available categories. Click an item in this list to display information related to that category.

- Hover the mouse pointer over a data point in the chart to display information about that point.

- Hover the mouse pointer over an item (line or area) in the chart to display the name of the item.

- When the chart is zoomed so that not all of a chart can be seen at the one time, the chart also displays scrollbars that allow you to choose the visible portion of the chart.

- Individual toolbar buttons may not be available in all chart views.

Click	Description
	Zoom in on Charts
	Pan Over Zoomed Charts
	Set the Chart View
	Filter on the Data Series
	Maximize and Restore Charts
	Save To File
	Copy To Clipboard
Kernel mode User mode	The *legend* lists all the data series associated with the chart. Click an item in the legend to highlight its series in the chart. (Click a second time to return the series to its normal appearance.) Move the mouse pointer over an item in the legend to view the current value for that series within the chart. **Note:** Arrows at the top and bottom of a legend (if displayed) indicate that there are more

Click	Description
	items in the legend list. Click an arrow to scroll through the list.
Right click the table body and select **What's This?**	View help specific to the chart.
Right click the table body and select **Print**	Print
Right click the table body and select **Restore Default Settings**	Restore Default Settings
Right click the table body and select **View as Grid**.	View the data in a Spotlight grid. About Spotlight Grids
Right click the table body and select **Properties**.	Adjust the chart properties: • Chart Properties - Chart Page • Chart Properties - Filtering Page • Component Properties - Overhead Page
⚠	If the diagnostic server is unable to collect the data required, Spotlight displays a small icon on the affected component. Click the icon to see more information about the problem.
🔄 ⏸	Pause / Resume / Refresh

Related Topics

Colors and flashing of alarms / Colors used in charts

About Spotlight Grids

Zoom in on Charts
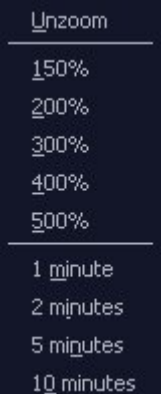
The Zoom (applied to the horizontal and vertical axes of the chart) determines the scale.

You can zoom in on a section of a chart using one of the following methods:

Action	Description

Use the magnifying glass
🔍 ▾
	1. Click the magnifying glass.
	
	2. Move the mouse pointer over the chart. (The pointer changes to a magnifying glass.)
	
	3. One of the following:
	
	   • Click on the chart. This centers the view on the chosen point and doubles the current zoom factor.
	
	   • Click-and-drag the mouse to enclose the area that you want to magnify. The chart zooms to the size of the enclosed rectangle when the mouse button is released.

Notes:

• To return the chart to its default size, right-click on the chart and select **Unzoom**.

• To switch temporarily from **Zoom** mode to **Pan** mode, hold down the **Ctrl** key.

• To zoom out, hold down the **Shift** key while clicking on the chart. This centers the view on the chosen point and halves the current zoom factor.

Action	Description
Use the drop-down menu 	1. Click the down arrow next to the magnifying glass to open the drop-down menu. 2. Click a **percentage** option to zoom the chart horizontally AND vertically by that percentage -or- Click a **time** option (where applicable) to zoom the horizontal axis of the chart to the specified time period, while resetting the vertical axis to the default value. **Note:** Time options are not available in chart views without a time axis (bar charts, for example).

Notes:

- You cannot zoom a Pie chart.
- When the chart is zoomed so that not all of a chart can be seen at the one time, the chart also displays scrollbars that allow you to choose the visible portion of the chart. Pan Over Zoomed Charts

Related Topics
Maximize and Restore Charts
About Spotlight Charts

Pan Over Zoomed Charts

When you zoom in on a chart you can use the Pan tool in the chart toolbar to view other sections of the expanded chart.

To pan over a zoomed chart

1. Click .

2. Move the mouse cursor over the chart. It will appear as a 'hand' icon.

3. Click-and-drag the mouse cursor across the chart to view other sections of the chart.

Tip: To switch temporarily from **Pan** mode to **Zoom** mode, hold down the **Ctrl** key.

Related Topics
Zoom in on Charts
Maximize and Restore Charts
About Spotlight Charts

Set the Chart View

To change the style of chart

1. Click .

2. Select the style of chart that is appropriate to your needs.

Related Topics
About Spotlight Charts
Chart Properties - Chart Page
Colors and flashing of alarms / Colors used in charts

Filter on the Data Series

A Spotlight chart may contain multiple graphs representing data series from multiple data sources. You can filter the chart so that it shows only the data series that you want to see.

To filter on the data series

1. Click .

2. Select as appropriate.

Option	Description
Show all items	Do not filter the contents of the chart.
Show top 5 items	Filter out all but the top 5 data series that appear in the chart.
Show top 10 items	Filter out all but the top 10 data series that appear in the chart.
Custom...	Create a custom filter from the available parameters. Chart Properties - Filtering Page

Related Topics
About Spotlight Charts
Chart Properties - Filtering Page

Maximize and Restore Charts

Where a drilldown displays several charts you can choose to **maximize** one chart so it occupies the whole drilldown. Once maximized, you can **restore** the chart to its normal size.
Note: All charts are automatically restored to their normal size when you view a different drilldown.

Toolbar	Menu	Shortcut
Click	Right click the chart and select **Maximize**.	N/A
Click	Right click the chart and select **Restore**.	N/A

Related Topics
Zoom in on Charts

Chart Properties - Chart Page

You can edit the visual properties of the charts displayed in Spotlight.

To set chart properties

1. Right-click the chart and select **Properties | Chart**.

2. Configure as appropriate:

Option	Description
Name	Select the chart to configure. Applicable when the chart name is selectable.
Chart style	Select a style for the chart.

Notes:

- You can use a **Summary** chart to view the statistics on how the charted data varies over time, including details on average, maximum, and minimum values. The percentile range (10% to 90%) shows the spread that contains 80% of all values in the chart.

- You can use a **Histogram** to sort data into ranges of values. The horizontal axis shows the distribution of data across those ranges. The vertical axis shows the number of items in each range. (A **current data** histogram displays the most recent data collected; a **trend** histogram displays an average of the data collected over the available history of the chart.)

- Any charts under the **Other views** heading (as in the Sample item here) are specific to individual Spotlight applications.

- Some charts that show task or status data may be drawn as **Gantt charts**. These are custom-coded charts and cannot be displayed in other styles.

Option	Description
Arrange by	Applicable when the Chart style is set to **Bar** or **Pie**.

- Select to arrange the data in **Name** or **Value** order.

Option	Description
Include legend in chart	Select to show the legend associated with the chart. Not all charts have a legend.
Series are stacked	Applicable when the Chart style is set to **Area**.

- Select to plot the cumulative values of the different series.

- Clear to superimpose the different series on the same set of axes.

Option	Description
Manual scaling Minimum value Maximum value	Select **Manual Scaling** to set the scaling factor for the chart in the **Minimum value** and **Maximum value** fields. Clear **Manual Scaling** to accept Spotlight's default scaling of the chart.

Related Topics

About Spotlight Charts

Set the Chart View

Colors and flashing of alarms / Colors used in charts

Chart Properties - Filtering Page

A Spotlight chart may contain multiple graphs representing separate data series from multiple data sources. You can filter the chart so that it shows only the data series that you want to see.

Spotlight supplies a set of filters that you can use, but if none of the standard filters are suitable for your chart, you can design a custom filter.

To set a custom filter on a chart

1. Right-click the chart and select **Properties | Filtering**.

2. Configure as appropriate:

Option	Description
Name	Select the chart to configure. Applicable when the chart name is selectable.
Show all items	Select to display all current data series.
	For example, if a system has the multiple disks **C:**, **D:**, **E:**, and **F:**, the relevant disk usage chart displays the disk usage statistic for all four disks.
Show only the top items	Select to display the most significant data series.
	Select the number of data series to show from the associated drop down list.
Show the following items	Select to choose the data series to show.
	Select the data series to show from the associated list. Note also the associated **Select all** and **Clear all** buttons.

Related Topics
[Filter on the Data Series](#)
[About Spotlight Charts](#)
[Component Properties - Overhead Page](#)

About Spotlight Grids

A table or grid shows values in tabular form. The grid component is used in Spotlight alarms and drilldowns.

Actions common to Spotlight grids

Notes:

- Not all items are available with all grids.

- If not all the grid can be seen all at the one time then scroll bars allow you to choose the visible portion.

Click	Description
	Find Text
	Alarms: Alarm Filters.
	Tree View
	Select Multiple Rows
	Select Multiple Rows
	Save To File
	Copy To Clipboard
	Show, Hide & Order Columns
	Sort the Grid
Right click the table body and select **What's This?**	View help specific to the gird.
Right click the table body and select **Print**	Print
Right click the table body and select **Restore Default Settings**	Restore Default Settings
Right click the table body and select **View as Chart**.	View the data in a Spotlight chart. About Spotlight Charts
Right click the table body and select **Properties**.	Freeze the first column or word-wrap text. Grid Properties - Options Page

Click	Description
	If the diagnostic server is unable to collect the data required, Spotlight displays a small icon on the affected component. Click the icon to see more information about the problem.
	Pause / Resume / Refresh

Related Topics
About Spotlight Charts

Find Text

To find a specified string of text in a Spotlight grid

1. Open the **Find** dialog.

Toolbar	Menu	Shortcut
Click	Right click the grid data and select **Find**.	N/A

2. Use the controls:

Control	Description
Text to find	Enter or paste in the text you want to find.
Match case	Select to match the case of the text exactly. When selected a search for the string AbC will not find the string abc.
Match entire cell contents	Select to match the contents of the table cell exactly. When selected a search for the string Spot will not find the string Spotlight.
Up	Search upwards from the current location then stop.
Down	Search downwards from the current location then stop.
Search all columns	Search all columns of the table.
Search selected column only	Search only the highlighted column.
Find next	Find the next occurrence of the text in the grid.

Note: The Find function may be of limited use in dynamic columns (*Elapsed time* is one example) whose values are continually updated.

Related Topics
About Spotlight Grids
Sort the Grid
Show, Hide & Order Columns

Sort the Grid

To sort the grid according to the contents of a column

Click the column header. An arrow appears to the right of the header .
Notes:

- The arrow points down when the column is sorted in descending order, up when the column is sorted in ascending order.

- Click the arrow to alternate ascending / descending order.

- Click a different header to select a different column to sort by.

- To sort by a column that is hidden, right-click the header row and select **Arrange by |** *column name* .

To sort the grid according to the contents of multiple columns

Keep the SHIFT key depressed when sorting by the second or later column.

For example, to sort the grid in the Alarm Log page by severity and then, within that sort, by action:

1. Click the **Severity** header.

2. While keeping the SHIFT key depressed, click the **Action** header.

Related Topics

Tree View

About Spotlight Grids

Tree View

The first column in a Spotlight grid may show data as a simple text string or (if available) a hierarchical tree structure.

Simple Text String	Hierarchical Tree Structure (Tree view)

Process
OUTLOOK.EXE
explorer.exe
MAPISP32.EXE
WZQKPICK.EXE
acrotray.exe
sqlmangr.exe
Msmsgs.exe
agentsvr.exe

Process ▲
⊞ explorer.exe
RoboHTML.exe
⊟ System
⊟ smss.exe
csrss.exe
⊟ winlogon.exe
lsass.exe
⊞ services.exe

Note: When you sort a table in tree view, the table is sorted first according to the highest level in the tree. The individual branches and sub-branches of the tree are sorted separately in the same order. Sort the Grid

Tree View in Drilldowns	Tree View in Alarm Pages

Click to toggle Tree View on or off.

Drilldown grids with this button on their toolbar can be viewed as a simple text string or a hierarchical tree structure.

Alarm pages Spotlight Today and Alarms by Time are displayed in a tree structure when **Group By Connection** or **Group By Alarm** is applied .

To manage the Tree View display, right click on the grid content and select **Collapse** or **Expand**.

Related Topics

About Spotlight Grids

Show, Hide & Order Columns

Alarms by Time

Show, Hide & Order Columns

To hide a column
Right-click the column header and select **Hide Column**.

To reorder a column
Click-and-drag the column header to its new location.

To open the Column Organizer
Use the column organizer to show or hide columns and choose the order in which to display columns. Right - click the grid header row and select **Organize columns**

Control	Description
Available columns	Hidden columns.
Selected columns	Visible columns in the order they appear in the grid.
Add	Adds the columns selected in **Available columns** to the **Selected columns** list.
Add All	Adds all columns in the **Available columns** list to the **Selected columns** list.
Remove	Moves the columns selected in the **Selected columns** list to the **Available columns** list.
Remove All	Moves all columns in the **Selected columns** list to the **Available columns** list.
Move Up	Moves the selected columns in the **Available columns** list further up the list.
Move Down	Moves the selected columns in the **Available columns** list further down the list.

Tip: To choose multiple columns in the **Available columns** and **Selected columns** lists, use the shortcut key combinations **CTRL+click** or **SHIFT+click**.

Related Topics

About Spotlight Grids

Select Multiple Rows

On grids with the **Select All** and **Clear All** buttons you can select multiple rows in the grid.

Click	To do this…
▤	Select all rows
▢	Clear all rows
CTRL + click	Select several individual rows (or clear rows already selected)
SHIFT + click	Select a range of rows

Related Topics

About Spotlight Grids

Grid Properties - Options Page

To freeze the first column

When selected the first column will remain visible while you scroll through the columns of the table:

Right-click the grid header and select **Freeze First column**.

To word-wrap text

By default, Spotlight truncates text that is too long to fit in the space allocated to the column. To allow text to wrap onto multiple lines:

1. Right-click the data content of the grid and select **Properties | Options**.

2. Select **Allow text to word-wrap onto multiple lines if it is wider than the column**

Related Topics

About Spotlight Grids

Component Properties - Overhead Page

Spotlight Home Page Components

The components in the Spotlight home page are visual objects that correspond to important features of the system Spotlight is monitoring. Components can be of many different types, including buttons, gauges, data flows, queues, spinners, and containers.

Note: They change color to alert you to performance bottlenecks. Spotlight Alarms

Component	Description	Component	Description
	Button - Sometimes called a process icon. A Spotlight component that contains a single value that represents the state or existence of a process.		Data flow - Panels in the Spotlight home page may be connected to other panels by data flows that show the rate at which data is moving through a system. As the rate increases, so does the speed of the flow.
	Gage - A Spotlight component that shows a measurement as a continuous range of possible values.		Queue - Spotlight component that shows a measurement as a range of discrete values. The globe at the end of the queue is an alarm indicator.
	Container - A cylindrical component that fills up as the value it measures increases. Normally used to indicate file size or disk usage.		Spinner - Spotlight component that shows a measurement as a rotating wheel. The higher the value, the faster the speed of rotation.

Action	Description
Hover the mouse over the component	Opens information on the status of the component. **Total CPU Usage** Total Windows processor utilization is 8.07%.
Click on the component.	Shows **What's this** help or performs other customized action. Balloon help
Right click on the component | **What's this**	Opens information about the component in balloon help. If an alarm is raised on the component, then opens help for resolving the alarm. Spotlight Alarms
Right click on the component | **Show History**	View the recent history of the component. **Tips**: • Change the size of the popup chart by clicking and dragging a corner of the popup. • Maximize and restore the size of the chart via the Maximize chart buttons at the top right of the popup chart. • Zoom in on a section of the chart by clicking in the chart, and zoom out by clicking the Unzoom chart button at the top right of the popup chart. • Right click and drag over the chart and time axis to zoom in on a time frame. • Right click the chart to copy to clipboard. • You can customize how history data is stored. see Spotlight History for more information. • You can customize the behavior and display of the history charts. Chart configuration
Right click on the component | **Show Details**	View a drilldown page for the component. Spotlight drilldown pages contain a detailed breakdown of the statistics used to diagnose the component.
Right click on the component | **Snooze alarm**	Ignore (Snooze) Alarms Components with snoozed alarms are represented by 🔔.
Right click on the component | **Copy to Clipboard**	Copy To Clipboard
Right click on the component | **Restore Default Settings**	You can change the appearance and behavior of Spotlight components. Restore Default Settings
Right click on the component | **Metrics**	Metrics Dialog Thresholds Dialog
Right click on the component | **Properties**	Home Page Properties - Options Page Component Properties - Overhead Page
Click ⚠.	This icon indicates the diagnostic server is unable to collect the data required

Action	Description
	for the affected component. Click the icon for more information.
	Data for this control is not available. (There are no alarms associated with this control.) Component Properties - Overhead Page
	Data for this control is not available. (Alarms associated with the control are disabled.) Component Properties - Overhead Page

Related Topics

Spotlight Home Pages

Calibrate Spotlight Connections

Home Page Properties - Options Page

To set the maximum display value for a Spotlight home page component

1. Right-click the component and select **Properties | Options**.

2. Configure as appropriate:

Option	Description
Name	The display name of the component.
Maximum value	Some components on the Spotlight home page display metric data as positions on a bar or graph. You can say how much of this bar or graph to display by entering a value in the Maximum Value property.
	Click … and enter a new maximum value.
	Note: For components that do not display metrics as a position in a range, the **Maximum value** is set to zero by default.

Related Topics

Spotlight Home Page Components

Component Properties - Overhead Page

Spotlight Home Pages

Restore Default Settings

For the Spotlight connection

This deletes all customized properties for the Spotlight connection, including any Spotlight overhead settings.

1. Select the Spotlight connection in the Spotlight browser Live connections

2. Select **File | Restore Connection Defaults**.

3. Click **Restore**.

For a Spotlight component

This deletes all customized properties applied to a chart or grid or Spotlight home page component. For a grid, this removes all changes to column headings and word-wrapping. For a chart, this removes changes made to chart style, scaling, legends and series.

1. Right-click on the data in the chart or grid or on the Spotlight home page component.

2. Select **Restore Default Settings**.

3. Click **Restore**.

Related Topics

Connection Properties - Overhead Page

Spotlight Home Page Components

About Spotlight Grids

About Spotlight Charts

Copy To Clipboard

To copy the chart or grid to the clipboard

1. Copy:

Toolbar	Menu	Shortcut
Click	Right click the chart or grid data and select **Copy to Clipboard**.	N/A

2. Paste from the clipboard to the desired destination – as an image within an image editor, or as text within a text editor.

Tips:

- In Spotlight grids you can choose to copy the whole table, the selected row, column or cell.

- Some applications allow you to use **Paste Special** to choose the format to use to paste the chart into a document.

- When you copy and paste the details of a **zoomed** chart to an **image** file, the image contains only the visible portion of the chart. When you paste the details to a **text** file, the pasted text includes the details of the entire chart, not only of the visible section.

Related Topics

Save To File

About Spotlight Charts

About Spotlight Grids

Save To File

To save the chart or grid

1. Open the **Save As** dialog.

Toolbar	Menu	Shortcut
Click 💾	Right click the chart or grid data and select **Save As…**.	N/A

2. Use the controls in the Save As... window to specify where (and in which format) to save the chart or grid.

Control	Description
File name	Type the name of the file that to save the chart or grid to.
Save as type	Select the file type.

                                   • For charts, the choices include **bitmap** image, Windows **metafile** image, and **comma-delimited** text file.

                                   • For grids, the choices include **formatted** text file and **comma-delimited** text file.

Text encoding	Choose the text encoding to use when saving the file.

                                   • ANSI, the default choice, supports only the character set used by Western European languages. Unsupported characters are converted to question marks.

                                   • UNICODE (UTF-16) is the 16-bit version of the Universal Character Set Transformation Format. This format is maintained by the Unicode Consortium (http://www.unicode.org/), and provides the basis for processing, storage and interchange of text data in any language.

                                   • UNICODE (UTF-8) is the 8-bit version of the Universal Character Set Transformation Format.

Related Topics

Copy To Clipboard

About Spotlight Charts

About Spotlight Grids

Print

To print a Spotlight home page or the current drilldown

Select **File | Print**.

To set print options

1. Select **File | Page Setup**.

2. Set how Spotlight windows are to be printed. The results of the choices you make are shown in the Preview pane of the window.

3. If required, click **Printer…** to select the printer to use when printing Spotlight pages.

To view a print preview

1. Select **File | Print Preview**.

2. Use the toolbar at the top of the Print Preview window to adjust display settings.

Button…	Description…
Print	Send this page to the printer.
Page setup	Choose how you want to print the contents of the page.
Show one page	
Show multiple pages	Choose how to view the preview.
Zoom	Choose how much you want to magnify the print preview to get a clear look at page details. Alternatively, move the mouse pointer over the print image, and click to magnify the image.
Close	Close the Print Preview window.

Related Topics

Spotlight Home Pages

About Spotlight Charts

About Spotlight Grids

Component Properties - Overhead Page

Spotlight components are visual objects on the Spotlight home page and in Spotlight drilldowns that correspond to important features of the system to which Spotlight is connected. You can set the options that affect the behavior and display of Spotlight components.

Note: You can set overhead properties for the Spotlight connection as a whole. Connection Properties - Overhead Page

To set the overheads for a Spotlight component

1. Right click the component and select **Properties | Overhead**

2. Select from the Refresh setting list how often Spotlight will collect data for the category of which this Spotlight component is a member.

Refresh setting	Description
Refresh at the default rate	Spotlight data is refreshed as per the definition in Spotlight Options. Set the rate at which data is collected **Note**: In most cases (but not all) this will be at the foreground rate when the component is visible and the background rate when the component is not visible.
Always refresh at the foreground rate	Always refresh at the foreground rate, whether the data is displayed or not.
Refresh only while data is	Refresh data at the foreground rate when it is visible in the Spotlight window.

Refresh setting	Description
displayed on the screen	
Do not refresh data in this category	The data is never collected. Spotlight displays one of the following icons on components that display this category of data: Data for this control is not available. (There are no alarms associated with this control.) Data for this control is not available. (Alarms associated with the control are disabled.)

Related Topics

Connection Properties - Overhead Page

Set the rate at which data is collected

Spotlight Home Page Components

About Spotlight Grids

About Spotlight Charts

Configure Spotlight - View | Options

Customize the appearance and behavior of Spotlight.

To open Spotlight options

Click **View | Options**

The categories you see in Spotlight depend on which Spotlight applications you have installed. This help page covers the standard categories. Individual Spotlight applications may contribute additional categories and options.

Category	Description
	How Spotlight reports unexpected events for your connections.
Alarms and notifications	- Alarm actions - Alarm smoothing - Balloon help - E-mail server - Taskbar
	The look and feel of Spotlight, and details of what is shown.
Appearance	- Chart configuration - Color scheme - Colors and flashing of alarms / Colors used in charts - Slideshow
Data collection	When to collect data from a target system, and the best way to display it. Set the rate at which data is collected

Category	Description
Security	Password protection and user access. • Password storage
Startup	How Spotlight behaves when it starts. • Navigation tree • Initial view
Storage	The locations for, and space assigned to, Spotlight history. • See Spotlight History for more information. **Note**: Spotlight History is installed if you click **Help | About Spotlight | Spotlight Modules** and under the heading **Spotlight Console** see **Light-weight XML repository**. If your Spotlight application does NOT include the lightweight XML repository, it may employ a different method to store historical data and you should View the online help for your Spotlight application for more information.
Troubleshooting	Settings that may be needed to address Spotlight performance issues. • Display performance

Notes:

- To locate an option in the Options dialog, type a keyword into the Search field.

- Additional customization options are available in the Spotlight Connection Manager, Metrics Dialog, Chart Properties - Chart Page, Grid Properties - Options Page and Home Page Properties - Options Page.

Alarm actions

Define rules that determine the action Spotlight takes when alarms are triggered.

Note: An action rule does not take effect until an associated alarm is triggered.

To set action rules for alarms

1. Click **View | Options**.

2. Click **Alarms and notifications | Alarm actions**.

3. Configure as appropriate

Option	Description
Perform the following actions when the conditions are met.	Rules are listed in the order they are applied. Highlight a rule to view, modify, delete, rename or move it. Only selected rules are applied. De-select a rule to retain it but not apply it.
Rule description	A description of the rule selected: condition(s) triggering action(s) in the order listed.
New...	Click to create a rule. Alarm Actions - Alarm Log Filter Dialog
Modify...	Highlight a rule. Click to modify the highlighted rule. Alarm Actions - Alarm Log Filter Dialog
Delete	Highlight a rule. Click to delete the highlighted rule.
Rename...	Highlight a rule. Click to rename the highlighted rule.
Move Down	Highlight a rule Click to move the highlighted rule lower down the list. It is applied later. Rules are applied in the order they are listed.
Move Up	Highlight a rule. Click to move the highlighted rule higher up the list. It is applied sooner. Rules are applied in the order they are listed.

Related Topics
Spotlight Alarms
About Alarms, Metrics, Thresholds And Severities
Configure Spotlight - View | Options

Alarm smoothing

Smoothing (also known as metric averaging) is used to prevent Spotlight from generating spurious alarms. It smooths out the peaks and troughs in Spotlight components by averaging the information that Spotlight collects over a specified period of time. It does not affect the charts or tables that display metric values; these continue to display metrics as they are reported.

Use this dialog to set default smoothing for all Spotlight metrics for the connection. To set smoothing for a specific Spotlight metric use the Metric editor. This overrides the default smoothing value for the metric. Metrics Dialog

To set default smoothing for Spotlight metrics for the connection

1. Select the Spotlight connection in the Spotlight Browser Live connections.

2. Select **View | Options**.

3. Select **Alarms and Notifications | Alarm smoothing**.

4. Use the **Average over…** slider to choose the time period over which you want to average metric values.

Note: The change is effective from the next refresh. Pause / Resume / Refresh

Related Topics
Metrics Dialog
Set the rate at which data is collected

Balloon help

By default Spotlight opens a context-sensitive balloon help window when you click a Spotlight home page component. Context-sensitive help may be component help or alarm help.

To configure balloon help

1. Click **View | Options**.

2. Click **Alarms and Notifications | Balloon help**.

3. Select from the following:

Option	Description
Always show balloon help	View component help or alarm help when you click on a component. Component help is shown when no alarm is raised for the component; alarm help is shown when an alarm is raised.
Only show help if an alarm is being raised	If NO alarm is raised for the component, view the related drilldown. If an alarm is raised, view the alarm help for the component.
Always jump directly to the related drilldown	Always view the related drilldown whether an alarm is raised or not.

Related Topics

E-mail server

You can set the message service that Spotlight uses when sending e-mail.

To set the Spotlight mail server

1. Select **View | Options**.

2. Select **Alarms and Notifications | E-mail server**.

3. Select the method that you want Spotlight to use when sending email

Service	Description
Messaging Application Program Interface (MAPI)	A standard Windows interface that you can use to send e-mail.
Simple Mail Transfer Protocol (SMTP)	An independent TCP/IP protocol used to send e-mail. Click **Change SMTP settings** to set the required SMTP mail options.

Related Topics

Taskbar

You will not always have the Spotlight console open in front of you while Spotlight is diagnosing the performance of your connections. In this situation, you can set up Spotlight to continuously oversee the status of all your connections, and to alert you whenever important alarms require your attention.

To set notification options

1. Select **View | Options**.

2. Select **Alarms and Notifications | Taskbar**.

3. Change the available options to specify Spotlight's notification behavior.

Option	Description
Show the Spotlight icon in the task bar notification area.	Select this option to display a Spotlight icon in the System Tray area of your Windows task bar. The icon has the following properties: • Its color reflects the highest-severity alarm raised by a current connection. • Double-clicking the icon restores the Spotlight console (if minimized). **Note**: You can also employ a floating **Connection status bar** to display the status of all current Spotlight connections. Connection Status bar
Pop up an alert where the severity is...	Select this option to enable the Spotlight icon to display an appropriate text message when the status of the current Spotlight connection reaches or exceeds the severity chosen in the associated drop-down list. Use the drop-down list to choose the threshold that will trigger the message. **Note**: The Spotlight icon in the Windows system tray displays the highest-severity alarm currently triggered for the corresponding connection.
Do not show alerts if Spotlight is the active window.	Select this option if you are currently using Spotlight, and do not want to be notified of new alarms that may already be visible.

Related Topics

Chart configuration

Spotlight makes extensive use of charts (historical and otherwise) in its drilldown windows. You can set a number of display options for these charts.

To set the appearance of charts

1. Select **View | Options**.

2. Select **Appearance | Chart configuration**.

3. Configure as appropriate.

Option	Description
Chart line style	Click the box, and from the list choose a line style that can be used to display charts. The sample chart above the control demonstrates how charts will appear under the chosen line style.
Show data tool tips when hovering over charts	Select this option to display tool tips when you hover the mouse pointer over an item in a Spotlight chart.
Show data from previous sessions	Select this option to display data from previous connections to this system (if available). This affects only the display of data from previous sessions. It does not affect the collection of data.
Amount of history to display in charts	This affects the time scale in all Spotlight charts, including those displayed when you select **Show History** from the shortcut menu invoked when you right click a Spotlight home page component. Spotlight Home Page Components

Related Topics

About Spotlight Charts

Configure Spotlight - View | Options

Color scheme

Change Spotlight's appearance by selecting a saved color scheme.

To select a saved Spotlight color scheme

1. Click **View | Options**

2. Click **Appearance | Color scheme**

3. Select the color scheme to use.

Note: You can create your own color scheme. Colors and flashing of alarms / Colors used in charts

Related Topics

Configure Spotlight - View | Options

Colors and flashing of alarms / Colors used in charts

You can change the appearance of the Spotlight console via Spotlight Options by using a simple point-and-click method. You can also create and save your own color scheme for Spotlight.

To set the appearance of Spotlight

1. Click **View | Options**

2. Select **Appearance | Colors and flashing of alarms** or **Colors used in charts**.

3. Select a Spotlight GUI component:

Component	Description
Main Screen	Change the appearance of the Spotlight home page.
Drilldowns	Change the appearance of Spotlight drilldowns.
Windows	Change the appearance of Spotlight editors and message windows.
Severities	Change the appearance of Spotlight alarms.
Chart colors	Change the colors used in Spotlight charts.

4. Click an item in the sample image or select it from the **Item** list.

5. Edit the item's appearance from the available (enabled) controls in the Appearance group.
    Review the following for more information.

Appearance	Description
Item	Display the type of features whose appearance you can change.
Foreground	Click to display a color chart, and then click a color on the chart to use it as the foreground color for the current feature (or click **More colors** to display a wider range of colors to choose from).
Background	Click to display a color chart, and then click a color on the chart to use it as the *background* color of the feature (or click **More colors** to display a wider range of colors to choose from).
Gradient fill	Select to use a gradient fill color on the feature; clear to remove an existing color fill.
Text shadow	Select to use a text shadow on the feature; clear to remove an existing shadow.
Image	When you choose a feature that has a bitmap image associated with it, choose an alternative image from the list. Select *Custom image* to use a different bitmap in place of the default image.
Flash speed	A display option for severities. Select an option from the list.
Flash intensity	A display option for severities. Select an option from the list.

6. Optional: Click **Save as** to save the settings as a new color scheme.
    Tip: Saved schemes can be selected from **View | Options | Appearance| Color scheme**.

Related Topics
Color scheme
About Alarms, Metrics, Thresholds And Severities
Configure Spotlight - View | Options

Slideshow

Configure Spotlight to start slideshow mode when you are not actively using the application.

1. Select **View | Options**.

2. Select **Appearance | Slideshow**.

3. Configure as appropriate.

Option	Description
Start Slideshow mode if Spotlight remains idle	Select this option to put Spotlight into Slideshow mode for unattended activation.
	Specify the delay period before an unattended Spotlight begins its slideshow.
Only start if Spotlight is the active window	Start Slideshow mode only if the Spotlight window is the one you are currently working in.
Connection switching interval	Use this option to specify how long Spotlight displays an individual connection before moving on to the next.
Show the following connections during slideshows	Specify how Spotlight will cycle through the available connections
	If you choose **Show connections with the following severity**, select an option from the list.
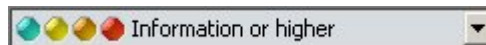
Related Topics

Spotlight Slideshow

Configure Spotlight - View | Options

Set the rate at which data is collected

Data is collected from monitored connections at foreground rate or background rate. Use this page to quantify foreground rate and background rate.

To set the rate at which data is collected

1. Select **View | Options**.

2. Select **Data Collection | Data refresh rates**.

3. Set values as appropriate.

Option	Description
Collect data at the following rate when it is being displayed	Enter a value. This value quantifies the foreground rate.
	While you are viewing a page in a Spotlight window, Spotlight collects the data that is displayed at the foreground rate.
Collect data at the following rate when it is being collected in the background	Enter a value. This value quantifies the background rate.
	Data that is reported in a hidden window is collected at the background rate.

Note: To improve the performance overhead for individual Spotlight connections, you can vary the data collection rates for a connection as a whole, or for individual components within that connection.

- Connection Properties - Overhead Page

- Component Properties - Overhead Page

Related Topics

Configure Spotlight - View | Options

Alarm smoothing

Metrics Dialog

Password storage

You can save passwords to Spotlight connections to allow other Spotlight users to diagnose target systems without needing direct access to the systems themselves.

To save passwords

1. Select **View | Options**.

2. Select **Security | Password storage**.

3. Select **Store passwords for new Spotlight connections**.

Note: To delete any passwords that have been saved previously, click **Clear all passwords**.

Related Topics

Configure Spotlight - View | Options

Navigation tree

To customize the display of the Spotlight browser

1. Select **View | Options**.

2. Select **Start-up | Navigation tree**.

3. Configure as appropriate.

Option	Description
Display the following pages in the navigation tree	Select the pages to be available in the Spotlight browser.
Open the following page in the navigation tree when Spotlight starts	Select the page to be open when Spotlight starts.

Related Topics

Spotlight Browser

Configure Spotlight - View | Options

Initial view

The startup page is the page displayed when you first start a Spotlight application.

To set the start up page

1. Select **View | Options**.

2. Select **Start-up | Initial**.

3. Select as appropriate.

4. At the dialog that opens, click **Change the start-up view**.

Related Topics
Configure Spotlight - View | Options

Display performance

Caution: This category contains Spotlight options that are rarely needed and that are designed for use when addressing Spotlight performance issues.

You can choose to balance display quality within the Spotlight home page and drilldowns against the speed at which the application deliver information.

To set display quality and performance

1. Select **View | Options**.

2. Click **Troubleshooting | Display performance**.

3. Drag the Display Performance slider to a position that marks an acceptable balance between display quality and display speed.

Related Topics
Configure Spotlight - View | Options

Troubleshooting

This section identifies general problems that you may encounter when using Spotlight applications, and details how to address those problems. For problems caused in specific Spotlight applications, see the help for that application.

- Product Authorization Errors
- I Cannot Open a Migrated Connection
- High Spotlight Load On Database Server
- Main Window Hard To Read
- No Change to Data Flows
- Disabled Controls
- Delay in Reporting Alarms
- Alarm Help Is Displayed Instead of The Drilldown
- Build the Support Bundle

Product Authorization Errors

When first installed, Spotlight can be used with full functionality for a specified trial period. When the trial period expires, contact your Questrepresentative to obtain an authorization key to continue.

Notes:

- Make sure that you enter the authorization key exactly as stated on your Product Authorization sheet or as provided by your Questdistributor. If you enter the authorization key incorrectly, the following message is displayed:

  Invalid authorization key

- Once the trial period has expired, if an authorization key is not supplied, only the Spotlight home page will be visible. If you click any of the drilldown buttons, the following message is displayed:

  You do not have a license for "Spotlight..." that enables access to this functionality. Please contact Questfor information about obtaining the correct license.

- If the trial date has not passed, check that the date on your system is set correctly.

How to enter an authorization key

1. Click **Help | About Spotlight**.

2. Click **Installed products**. The Installed Products page lists all the Spotlight products that are currently installed on your system.

3. Locate your Spotlight product in the list of installed products.

4. Click the associated **View/change product license** link.

5. Click **Change this license**.

6. Enter the new **Authorization key** and **Site message**.

Tip: Individual Spotlight applications may handle invalid license numbers in different ways (for example, if the license has expired on a trial copy of the software). For more information, refer to the corresponding help in that application.

I Cannot Open a Migrated Connection

Problem

Spotlight is unable to open a migrated connection.

Solution

When you open a Spotlight connection, you use connection details that are stored in files somewhere on your local computer or (in the case of migrated connections) somewhere on your network.

When Spotlight cannot access those connection details (for example, when the relevant network machine is unavailable), it displays a Spotlight message box that you can use to change the location where connection details are stored. The choices you have are:

Option	Description
Yes	Click to open the Spotlight Migration wizard to change the location where connection details are stored.
No	Click to abandon the attempt to open the desired Spotlight connection.
Retry	Sometimes the network location that you have been using is only temporarily unavailable. Click to reuse the existing connection details to open a Spotlight connection.

Note: You should not migrate connections to locations that may be offline frequently or for extended periods. You cannot migrate connections that are currently in use or inaccessible.

Related Topics

Troubleshooting

Migrate (Move And Share) Connections

High Spotlight Load On Database Server

The queries Spotlight uses to collect data can place an additional load on the system under diagnosis.

Most Spotlight queries consume a measurable but not significant amount of system resources. However, queries that are complex or collect large amounts of data may put a significant load on the system.

If you find that the load that Spotlight places on your server is too high, you can adjust collection rates for different categories of Spotlight data, and so limit the load that Spotlight places on the system under analysis.

To choose a standard Spotlight setting or create your own Custom settings

1. Close and re-open Spotlight.

2. Select **File | Connect | Spotlight**.

3. Right click the Spotlight connection experiencing the high load and select **Properties**.

4. Select **Overhead**.

5. Adjust the Spotlight load for the connection:

    - Use the Overhead Management slider to adjust the overall Spotlight load.

    - Alternatively, click **Custom…** to adjust specific load categories.

6. Open the Spotlight connection.

Note: The procedure for adjusting the Spotlight load is covered in more detail in the online help.

Main Window Hard To Read

The screen resolution is too low.

Increase the number of colors your computer is using. Spotlight requires a monitor that supports 65536 colors or higher (or Hi Color/16-bit) so that it can properly shade its graphs.

1. Open the Windows Control panel and double-click **Display**.

2. Click the **Settings** tab.

3. Increase the number of colors in the palette. If this option is not available you should upgrade your hardware.

No Change to Data Flows

Either the refresh rate is too slow or Spotlight is paused. Do the following:

- Update the screen. Click **View | Refresh**.

- Check the foreground refresh rate setting (try a rate of approximately 30 seconds).
  To view / change the rate click **View | Options** and select **Data collection | Data refresh rates**

- Resume Spotlight diagnostics. Click **View | Resume**. (If **Resume** is not available, Spotlight is not paused.)

Disabled Controls

Problem

I have a number of disabled controls on a Spotlight page. Every disabled control has one of these icons associated with it:
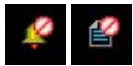
Solution

To solve the problem, view the Overhead Management page in the Component Properties editor for one of the disabled controls. You should find that the refresh setting for that component has been set to Do not refresh this control. Component Properties - Overhead Page

Background to The Problem

When Spotlight collects data from a system under investigation, it places a load on that system's resources.
Usually, the Spotlight load does not have a significant effect on the performance of the system, but when it does you can use Spotlight's overhead management tools to reduce or eliminate the load. Two factors influence overhead management:

- The categories of data collected.

- The refresh rate assigned to each category.

The categories of data collected are defined by the current Spotlight application, and are set by criteria that may include:

- The kind of data being collected.

- The load that collecting the data will put on the system.

- Where the data in the category is to be displayed.

- The importance of the data.
- How often the data is needed.

By default, every category is assigned one of the following refresh rates:

- Always refresh at the foreground rate.
- Refresh at the default rate.
- Refresh only while data is displayed on the screen.
- Do not refresh data in this category.

When you change the refresh rate for the data in a Spotlight control (a component, table or chart), you also change the refresh rate for ALL data in the same category. When you disable the refresh rate for the data in a control, you disable the refresh rate for ALL data in the same category.

A control whose data collection is disabled displays one of the indicated icons:

Icon	Description
	Data for this control is not available. (Alarms associated with the control are disabled.)
	Data for this control is not available. (There are no alarms associated with this control.)

Related Topics
Troubleshooting
Component Properties - Overhead Page

Delay in Reporting Alarms

Problem

The current displayed value of a Spotlight metric should be raising an alarm, but that alarm has not yet been reported.

Solution

There are two parts to this problem:

- **Refresh rates**
  Spotlight refreshes the metric data it collects at a rate specified in Spotlight Options. In most cases (but not all) this is the foreground rate when the control is visible, or the background rate when the control is not visible. These refresh rates (which you can set) are usually set to different values, with foreground data being collected more often. Set the rate at which data is collected

- **Metric smoothing**
  Spotlight uses metric smoothing (averaging) to prevent the generation of spurious alarms. It does this by averaging the data that it collects over a specified period of time. Alarm smoothing Metrics Dialog

The combination of these two features means that an alarm is raised faster for metrics in the page you are viewing than it is for those in a page that is hidden.

Related Topics

Alarm Help Is Displayed Instead of The Drilldown

Spotlight can be configured to show the help or jump directly to the related drilldown on clicking a home-page control.

To show the drilldown

1. Click **View | Options**.

2. Select **Alarms and notifications | Balloon help**.

3. Select **Always jump directly to the related drilldown**.

Build the Support Bundle

In emailing a request for assistance to Quest (support@quest.com) be sure to include a support bundle.

To create the support bundle

1. Click **Help | Support Bundle**.

2. Select the Spotlight application(s) whose information you want to collect.

3. Click **Collect**.
   A file is created containing a snapshot of your Spotlight installation. The file is called **SpotlightSupport.zip** The location of this file is C:\Users\<user>\AppData\Roaming\Quest Software\Spotlight\Default\Support

About Spotlight History

An important feature within Spotlight is the ability to display, collect and replay the behavior of a Spotlight connection at a specified point in time such as when an important alarm was raised.

Action	Description
Select a connection in Spotlight Browser Live connections and click **View | History Browser**.	View historical snapshots for the Spotlight connection. **Note**: While you are viewing historical data the connection identifier at the top right of the window indicates the date and time of the snapshot. On the Spotlight home page the console status bar indicates "History Browser".
Right click on a Spotlight home page component and select **Show History**	View the recent history of the Spotlight home page component. **Tips**: • Change the size of the popup chart by clicking and dragging a corner of the popup. • Maximize and restore the size of the chart via the Maximize chart buttons at the top right of the popup chart. • Zoom in on a section of the chart by clicking in the chart, and zoom out by clicking the Unzoom chart button at the top right of the popup chart. • Right click and drag over the chart and time axis to zoom in on a time frame. • Right click the chart to copy to clipboard.
From the Alarm log, right click an alarm and select **Show Selected Entry In | History Browser**	The Alarm Log provides an easy way to view the history for a particular alarm.
Click **View | Options | Storage**	Configure the locations for, and space assigned to, Spotlight history. • Clear history • History limits • History location

Related Topics

History Browser
Spotlight Home Page Components
Alarm Log

Clear history

To clear information from Spotlight history

1. Select **File | Disconnect All** to close all open Spotlight connections.

2. Select **View | Options**.

3. Select **Storage | Clear history**.

4. Select the repositories to delete.

5. Click **Delete**.

Related Topics
About Spotlight History

History limits

To set disk space for Spotlight history

1. Select **File | Disconnect All** to close all open Spotlight connections.

2. Select **View | Options**.

3. Select **Storage | History limits**.

4. Configure as appropriate.

Option	Description
Limit disk space usage	Select to limit the amount of space allocated for files in the history repository. Select how much space to allocate.
Limit repository period	Select to limit the length of time files are stored in the history repository. Select the period of time. **Note**: When clear, Spotlight stores the snapshot files indefinitely.

Related Topics
About Spotlight History

History location

To set where to save Spotlight history

1. Select **File | Disconnect All** to close all open Spotlight connections.

2. Select **View | Options**.

3. Select **Storage | History location**.

4. Configure as appropriate.

Option	Description
Select the path where history will be stored	Enter where you want to store historical information, or click on the folder icon to browse for its location.
Compress history files	The Spotlight history files stored in the repository may contain very large amounts of information. Select this option to compress history files when storing them.
Space available on	Shows how much space is available to store your snapshot files in the specified location.

Related Topics

About Spotlight History

Welcome on Windows

Spotlight is a powerful diagnostic and problem-resolution tool for Windows operating systems. Its unique user interface provides you with an intuitive, visual representation of the activity on the host machine.

Connect to Windows Systems

1. From the Spotlight Browser select **All Connections | Spotlight on Windows**

2. Select the Windows system. If the Windows system is not listed then follow the instructions to add a new connection.
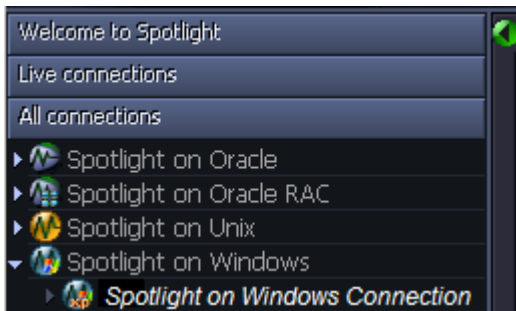
Add A New Connection

You will need	Description
Administrator access	Ensure your login to the Windows machine has Administrator access privileges. To add a new connection to the Windows machine you are currently using, ensure you are logged in to the machine with Administrator access privileges. If necessary, logout and login again.
Privileges	Spotlight retrieves its data from Windows performance counters and the windows registry of the monitored system. An administrator login to the Windows machine has such access.
Remote connectivity	To add a new connection to a Windows machine other than the one you are currently using, ensure the machine is accessible to the Spotlight client. Spotlight uses NetBIOS traffic to retrieve **perfmon** and registry information, so any firewalls between the Spotlight on Windows client and the machine being monitored must allow this traffic to pass through. Troubleshooting Spotlight on Windows

1. Click **File | Connect**

2. Select **Spotlight on Windows** on the Connections menu.

3. Double-click **Add new connection**.

4. Fill in the **Details** page of the Properties window as follows:

Field	Description
Connection name	Enter the preferred display name for the Windows machine. **Note**: If left blank, the **Connection name** field resets to the value of the **Address** field.
Local Machine	Select to diagnose the machine you are now using.
Address	Enter the hostname, IP address or URL to the Windows machine. **Note:** The value changes to "." and cannot be edited when **Local Machine** is selected.

Logon Details
Note: Not applicable when **Local Machine** is selected.

Field	Description
Domain	The name of the domain the user belongs to (to login to the Windows machine). **Note**: If you intend to leave **User** blank then **Domain** should also be left blank.
User	Enter the user name to login to the Windows machine. **Note:** Ensure this user has administrator rights.
Password	Enter the password to login to the Windows machine. **Note**: Applicable when **Save password details (for this connection)** is selected.

Select **Save password details (for this connection)** to save the password.
Note: Not applicable when **Local Machine** is selected.

5. Close the dialog.

Click	Description
Connect	Save changes and open the connection in Spotlight.
OK	Save changes. Do not open the connection in Spotlight.
Cancel	Do not save changes.

What Is Spotlight?

Spotlight is a powerful diagnostic and monitoring tool for Unix operating systems. Its unique user interface provides you with an intuitive, visual representation of the activity on your host machine. Graphical flows illustrate the rate at which data is moving between system components. Icons display the value of key statistics and measurements (metrics).

The power of Spotlight lies in its ability to provide visual and audible warnings if the performance metrics exceed acceptable thresholds. The components and dataflows change color to show you the source of the problem.

A range of visual graphs and tabular grids provide you with detailed information about your Unix hosts. This information can be viewed on the screen or printed.

You can set Spotlight to warn you when a threshold is reached. You may set a number of thresholds so that warning messages are displayed well before the traffic levels into or out of a host become critical. Spotlight uses a number of different techniques to warn you when a Unix host is exceeding a threshold.

When Spotlight detects a condition that it considers is a potential problem, it not only informs you about it, but also advises you what you could look at to diagnose the problem further, and suggests corrective actions.

Features of Spotlight on Windows

- Spotlight presents a visual representation of process flows within a Windows operating system, allowing you to observe actual system activity in real time.

- It visually identifies system bottlenecks and provides extensive drilldown capabilities.

- It displays the details of problem areas, including CPU resources, paging activity, and memory use, for rapid problem resolution.

- It uses visual and audible warnings to alert you when performance metrics exceed acceptable thresholds.

- Its drilldowns provide detailed information about a specific component allowing you to pinpoint the source of problems.

- Spotlight learns the normal range of values for your system.

- Spotlight assesses the normal rate of process flows via a calibration process and sets the display speed of the visual indicators accordingly.

- It can simultaneously observe multiple systems.

- It is easy to install.

Windows Architecture Overview

This diagram summarizes the architecture of Windows. This architecture is the basis for the design of the Spotlight home page. Spotlight Home Page
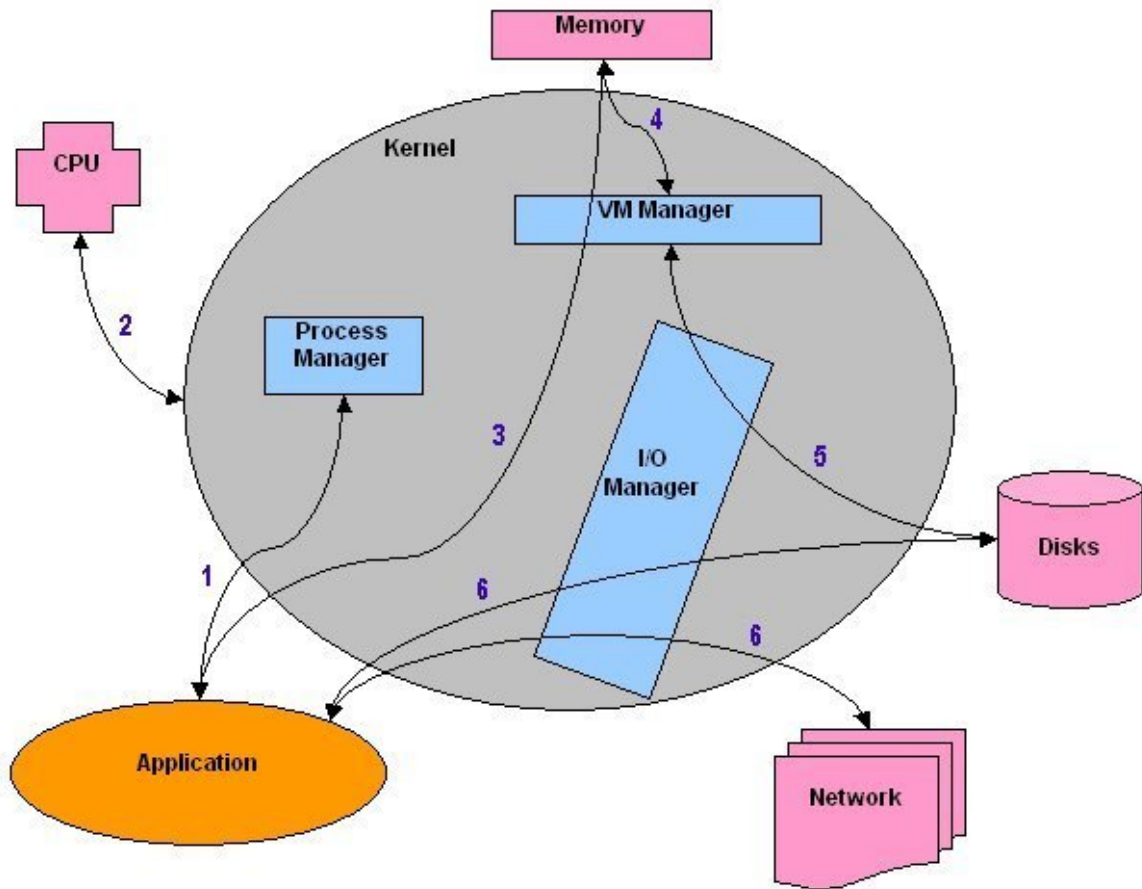
For more information on the Windows architecture, see *A Review of the Windows Architecture* in the online help.

Spotlight Home Page

The Spotlight home page shows the flow of information and commands between various sub-components and the size and status of internal resources such as processes, disk files and memory structures.

Related operating system statistics are grouped together on panels that are connected by a series of graphical flows and icons. Spotlight updates these flows in real time so you can see how quickly data is moving through the system. The icons change color as Spotlight alarms are raised, upgraded, downgraded and canceled.

The panels on the Spotlight home page are described in more detail next. For a full discussion of the panels, see the Spotlight online help.

Tip: Hover the mouse pointer over a panel component for more information.

System Panel

The **System** panel indicates:

- The type and version of operating system.
- The name of the service pack installed on the system (if any).

Network Panel
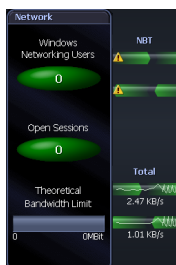
The **Network** panel shows the total number of connected users, and the rate at which packets are being sent from and received by the system:

- The number of clients connected to this system.
- The current number of Windows Networking sessions open on this machine.
- The level of network traffic graphed against a "theoretical" maximum bandwidth for the network card specified.

The flows between the **Network** panel and the **CPU** panel represent the rate of data transfer between the Windows host and the connected network.

Event Log Panel

The **Event Logs** panel provides a link to the Event Log drilldown.

- It alerts you to items that have reached alarm status as per your configuration of Event Log alarms.
- The button icon shows the number of Event Log alarms that have been raised but not yet cleared. You can acknowledge and clear alarms via the Event Log drilldown.

Note: This button is enabled on **Enable Event Logs** selected. See the online help for more information.

CPU Panel
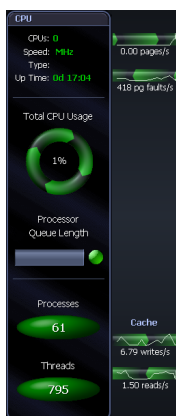
The **CPU** panel shows processor and load information for the Windows system:

- The number of processors installed in the system.
- The speed of the processors installed in the system.
- The type of processors installed in the system: Intel Pentium for example.
- The length of time since Windows was last booted (started).
- The amount of CPU being used on the system.
- The average number of threads (program execution units) that are waiting to run on each processor.
- The number of applications and services that are currently active on the system.
- The number of threads currently active in all processes.

The flows between the **CPU** panel and the **Memory** panel represent paging information between CPU and memory on the Windows system. The flows include the number of memory pages written out and read in to memory per second.

Memory Panel

The **Memory** panel shows detailed information on the Windows system.:

- The total amount of physical memory in RAM.

- The amount of physical memory available to applications.

- The percentage of physical memory that is available to applications.

- The amount of physical memory (RAM) Windows is using.

- The maximum amount of virtual memory available without expanding the paging file.

- The amount of virtual memory available.

- The percentage of virtual memory available to applications.

- The amount of Virtual Memory available without expanding the paging file

- How much Virtual Memory is in use.

- The percentage of soft pages versus hard pages found in RAM.

- The percentage of disk requests that are found in cache memory.

The flows between the **Memory** panel and the **Paging Files** panel represent the rate at which virtual memory pages are written to disk from memory, and swapped from disk into memory.

The flows between the **Memory** panel and the **Disks** panel represent the rate at which memory is written to disk and read from disk.

Paging Files Panel
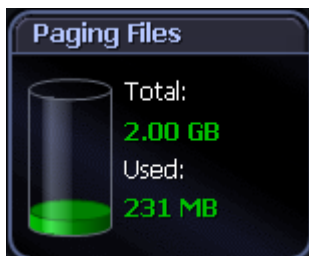
The **Paging Files** panel shows information on paging space for all paging files active on the system.

- The total size (MB) of all paging files on all disks in the system.

- The current amount of space (MB) being used by all paging files on all disks in the system.

Disks Panel
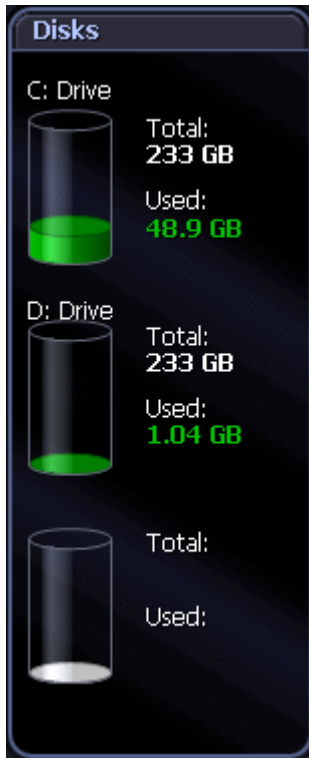
The **Disks** panel shows information for each logical disk on the system.:

- The drive letter.
- The size of the disk (MB).
- The space used on the disk (MB).

Files Panel
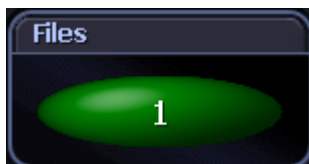
The **Files** panel shows the number of files whose file size is currently being tracked by Spotlight on Windows.

Note: This panel is hidden while no files are being tracked. To add / remove files, right click on the Files panel and select **Files Options**. See the online help for more information.

Spotlight Alarms

Spotlight alerts you to problems with your system by issuing an alarm. You can configure Spotlight in the level of severity that constitutes an alarm, to disable an alarm, and the actions Spotlight takes on raising the alarm.

When an alarm is raised Spotlight performs one or more of the following actions:

- Changes the color or intensity of relevant components.
- Gives audible warnings.
- Sends email notifications.

Actions you can take on an alarm being raised include:

- View details of the statistics that caused the alarm in a Spotlight drilldown page. Spotlight Drilldowns

- View the Spotlight online help.

- View details of the alarm in the Spotlight **Alarm Log** drilldown.

- Snooze the alarm.

- Save the alarm details.

- Filter the list of alarms.

Following are the alarms specific to Spotlight. For possible solutions to the problems indicated by these alarms or for information on how to diagnose problems further, see Spotlight alarms in the online help.

This alarm…	Alerts you…
Cache alarms	When the level of cache activity exceeds a threshold.
Disk alarms	As the logical disk is filling up, and also when disk activity exceeds a threshold.
Event Log alarm	When an Event Log entry is found which matches an active Event Log severity rule.
Memory alarms	When the available memory (both physical and virtual) drops below specified thresholds.
Paging alarms	When there is excessive paging activity on the target system.
Process alarms	When processes sustain a high CPU load.
Windows CPU alarm	When the server's CPU usage exceeds a specified percentage.

Cache Reads Alarm

This alarm is activated when the level of cache reads exceed a threshold.

High cache reads do not usually indicate a problem.

When this alarm is current you should look at the **Processes** tab in the **Processes** drilldown to see which applications are busy. Look at **Page Faults/Sec** and **% CPU**. These columns will show which applications are causing the high level of cache reads.

Related Topics
Processes Page

Cache Writes Alarm

This alarm is activated when the level of Cache Writes exceeds a threshold.

High cache writes usually mean one of the following:

- There is a lot of updating occurring in memory pages.

- An application is requesting to write information directly to disk.

When this alarm is current you should look at the **Processes** tab on the **Processes** drilldown to see which applications are busy. Look at page faults delta and CPU time. These columns will show which applications are causing the high level of cache writes.

Related Topics

Current Users Nearing Limit Alarm

When the current user count nears the number of user licenses available on the machine, it may be time to upgrade the license count on that particular system.

Diagnose Single Application Alarm

This alarm is raised when a process is terminated that was selected for diagnosis in the Windows Applications Metrics options in the Spotlight on Windows Options window.

Related Links

Disk Alarm

This shows that the logical disk is filling up or is full.

To rectify this:

- Inform users of the disk volume to clean up space by deleting or moving files no longer required.

- Go through and delete files that are no longer required from the volume.

- Move files to a volume with spare space.

- Purchase more disk capacity for the system, and move data onto it.

It is recommended that a backup of the system is taken before any files are moved or deleted.

Disk Reads Alarm

The **Disk Reads** alarm is activated when the level of traffic being transmitted from this disk exceeds a threshold. Sustained high disk activity can cause an overall degradation of system responsiveness.

When this alarm is current you should look at:

- The **Page File Transfers** chart on the **Paging Activity** page of the **Memory** drilldown. High paging in suggests that the high number of disk reads could be due to system paging. Adding more memory may alleviate the problem.

- The **NBT** page of the **Network** drilldown and see if any systems are transferring a large amount of data.

- The **Disk Queue** chart on the **Physical Disks Activity** page of the **Disks** drilldown. If the disk queue length is high on a single physical disk you can look at:

  - Spreading frequently accessed data over multiple physical disks.

  - Spreading page files over multiple physical disks.

  - Implementing a faster disk subsystem.

  - Implementing RAID striping technology.

If the high number of reads is causing an overall degradation of the system, consider upgrading the disk subsystem (controller and/or disks).

Related Topics
Summary Page
NBT Page
Physical Disk Activity Page

Disk Writes Alarm

The **Disk Writes** alarm is activated when the level of traffic being transmitted to this disk exceeds a threshold. Sustained high disk activity can cause an overall degradation of system responsiveness.

When this alarm is current you should look at:

- The **Page File Transfers** chart on the **Paging Activity** page of the **Memory** drilldown. High paging out suggests that the high number of disk writes could be due to system paging. Adding more memory may alleviate the problem.

- The **NBT** page of the **Network** drilldown and see if any systems are transferring a large amount of data.

- The **Disk Queue** chart on the **Physical Disks Activity** tab of the **Disks** drilldown. If the disk queue length is high on a single physical disk you can look at:

  - Spreading frequently accessed data over multiple physical disks.

  - Spreading page files over multiple physical disks.

  - Implementing a faster disk subsystem.

  - Implementing RAID striping technology.

If the high number of writes is causing an overall degradation of the system, consider upgrading the disk subsystem (controller and/or disks).

Related Topics
Summary Page
NBT Page
Physical Disk Activity Page

Event Log Alarm

This alarm becomes active when an Event Log entry is found which matches an active Event Log rule.

You can enable event logs and set event log rules. Windows Event Log Metrics

Related Topics

Logical Disk Statistics Alarm

This alarm is activated when disk-based performance counters are not enabled or not updated.

Windows 2000

Not all disk-based performance counters are enabled by default on Windows 2000 systems. If this is the case, Spotlight on Windows cannot retrieve the data in these counters, and the disk activity flows (to the left of the **Disks** panel) are disabled for this machine.

To rectify this situation:

1. Go to the machine's command prompt and type:

   diskperf -y

2. For this change to take effect, you then need to reboot the system under diagnosis.

Dynamic disks and Windows 2000

Windows 2000 systems now have additional support for dynamic disks, but the disk performance counters under Windows 2000 have not been updated to report this information correctly. Because of this, Spotlight on Windows cannot retrieve the data in these counters, and the disk activity flows (to the left of the **Disks** panel) are disabled for this machine.

To rectify this situation:

- Remove all dynamic disks from the machine, or
- Upgrade the machine to the Windows XP operating system.

Related Topics

Low Memory Alarm

This **Low Memory** alarm is activated when the available memory drops below a threshold. This can result in:

- Applications being unable to secure additional memory.
- New programs being unable to open successfully.
- General performance degradation due to excessive paging.

When this alarm is current you should:

- Look at the **Processes** page in the **Processes** drilldown. Check the **Mem Usage (MB)** column for processes using excessive amounts of memory. You can:

  - Close those applications using excessive memory.

  - Change the maximum memory usage setting for those applications using excessive memory. Refer to the relevant documentation to find out if this setting is supported by the application.

- Purchase more physical memory for your system.

- Determine if any programs have a memory leak. To do this:

  1. Restart your machine to ensure fresh memory usage by all applications.

  2. Use the **Microsoft Performance Monitor** tool and set the update interval to 10 minutes.

  3. Add the **Working Set Peak** counter for all applications from the **Process** object.

  4. Leave the application running on the system overnight.

  5. If there is a memory leak that can be viewed in a day, it will show up with the **Working Set Peak** graph for one of the applications rising smoothly. If the results are inconclusive after a day, some system applications that have a memory leak may exhibit the behavior over a week or more. Adjust the update interval to accommodate the extended period (update once per hour or more).

Related Topics
Processes Page

Missing Performance Counters Alarm

The Missing Performance Counters alarm becomes active when performance counter data is unavailable. This may occur when:

- The machine is unusually busy, or

- The performance counter has been disabled.

To verify the performance counters are unavailable, run Microsoft Performance Monitor and look for the Memory and Processor performance objects in the Add Counters dialog for the machine being diagnosed. If these objects are not listed then they need to be enabled. To do this, download and install Exctrlst.exe from the Microsoft download site.

Run the utility and do the following:

1. Edit the Machine Name to show the name of the machine being diagnosed.

2. Click the **Refresh** button.

3. Select **Service** as the sort order.

4. Select **PerfOS** in the counter list.

5. Check the **Performance Counters Enabled** checkbox if unchecked.

6. Repeat steps 4-5 for the PerfDisk, PerfNet and PerfProc counters.

7. Restart the machine being diagnosed to enable the counters.

Network Failure Alarm

This alarm becomes active when you lose the Spotlight connection to a machine under diagnosis. This may occur when:

- The network is down (and all open connections are lost), OR
- The remote machine is rebooted, OR
- One of the services required by Spotlight is down.

Look for the following on the machine that has raised the alarm:

- Is the machine running?
- Does it respond to commands?
- Do error messages appear on the screen or in the event logs?
- Can it "ping" other machines successfully?
- Can it access shared folders on the network?
- Are the Remote Procedure Call, Remote Registry and Windows Management Instrumentation services running?

If no problems are found, try to reconnect to the machine via Spotlight.

Page Faults Alarm

This alarm is activated when the number of pages being swapped exceeds a threshold. Sustained high paging rates can adversely effect the performance of a system.

When this alarm is current, you should:

- Look at the **Processes** page on the **Processes** drilldown to see which process is causing the paging. Look at the **Page Faults/sec** column. This will help you determine the cause of paging.
- Look at the **Paging** chart on the **Summary** page of the **Memory** drilldown. This shows how long the high paging has been occurring. Short periods of high paging are acceptable, but if the paging rate is high for a sustained period, there may be a problem with the system.
- Consider adding more physical memory to the system.
- Stop unnecessary services and processes on the system.

Related Topics
Processes Page
Summary Page

Page File Location Alarm

This alarm is activated when paging is low, but the disk activity on a paging file disk is high.
You may need to investigate moving paging files to another disk.

Page Hits Alarm

This alarm is activated when the proportion of page requests that can be sourced from memory rather than disk drops below a threshold.

When this alarm is current you should:

- Look at the **Paging** chart in the **Summary** page of the **Memory** drilldown to see if a high level of paging has been occurring for an extended period. Spikes in requirements are expected.

- Look at the **Page Faults/sec** column on the **Processes** page in the **Processes** drilldown.

  This will tell you which applications are making the most page requests. If this is an ongoing problem, look at spreading the load between multiple systems.

- Look at increasing the physical memory of the system.

Related Topics

Summary Page

Processes Page

Page Outs Alarm

This alarm is activated when the number of pages being written to disk exceeds a threshold. Sustained high paging rates can adversely effect the performance of a system.

Only those pages that have been altered are written to disk. Pages that have not been changed are dropped.

When this alarm is current, you should:

- Look at the **Processes** page on the **Processes** drilldown to see which process is causing the paging. Look at the **Page Faults/sec** column. This will help you determine the cause of paging.

  A high level of page faults and page outs indicates memory thrashing. If this is the case, the system may require more memory.

- Look at the **Paging** chart on the **Summary** page of the **Memory** drilldown. This shows how long the high paging has been occurring. Short periods of high paging are acceptable, but if the paging rate is high for a sustained period, there may be a problem with the system.

- Consider adding more physical memory to the system.

- Stop unnecessary services and processes on the system.

Related Topics

Processes Page

Summary Page

Page Reads Alarm

This alarm is activated when the number of pages being read in exceeds a threshold. Sustained high paging rates can adversely effect the performance of a system.

When this alarm is current, you should:

- View the **Processes** page on the **Processes** drilldown to see which process is causing the paging. Look at the **Faults/s** column to determine the cause of paging.

- View the **Page File Transfers** chart on the **Paging Activity** page of the **Memory** drilldown. This shows how long the high paging in has been occurring. Short periods of high paging are acceptable, but if the paging rate is high for a sustained period, there may be a problem with the system.

- Consider adding more physical memory to the system.

- Stop unnecessary services and processes on the system.

Related Topics
Processes Page
Summary Page

Page Writes Alarm

This alarm is activated when the number of pages being written to disk exceeds a threshold. Sustained high paging rates can adversely effect the performance of a system.

Only those pages that have been altered are written to disk. Pages that have not been changed are dropped.

When this alarm is current, you should:

- Look at the **Processes** page on the **Processes** drilldown to see which process is causing the paging. Look at the **Page Faults/Second** column to determine the cause of paging. A high level of page faults and page outs indicates memory thrashing. If this is the case, the system may require more memory.

- View the **Paging** chart on the **Summary** page of the **Memory** drilldown. This shows how long the high paging has been occurring. Short periods of high paging are acceptable, but if the paging rate is high for a sustained period, there may be a problem with the system.

- Consider adding more physical memory to the system.

- Stop unnecessary services and processes on the system.

Related Topics
Processes Page
Summary Page

Paging File Alarm

This alarm is activated when the utilization of the paging file exceeds a threshold.

When this alarm is active you should:

- Look at the **Processes** page of the **Processes** drilldown. Look at the **Virtual MB** column to see which applications are using the most virtual memory.

- Some applications (such as Microsoft Exchange or Microsoft SQL Server) can have their memory utilization limited.

- Close any superfluous processes.

- Look at increasing the size of the page file.

- Look at increasing the amount of RAM in the machine.

Related Topics
[Processes Page](#)

Paging File Disk Location Alarm

This alarm becomes active when there is more than one paging file on a single physical disk.

This can cause performance degradation – especially on IDE disks. IDE disks allow only a single disk operation to be active on the bus at any time.

To rectify this:

1. Open the Windows Control Panel.

2. Open the System control panel.

3. Click the Advanced tab.

4. Choose Performance Settings, and change the paging file allocations.

Percentage Bandwidth Alarm

The total network bandwidth capacity of the specified network card is nearing the limit where it is saturating the network link. If this is happening regularly, look at:

- The **NBT** page of the **Network** drilldown to see if any users or other systems are copying an inordinate level of data between systems.

- Moving network applications or shared files to another machine to balance the load.

- Upgrading the network subsystem to a faster technology.

- Adding an additional network card to the machine and configuring the system to utilize it.

Tip: If there are multiple network cards on the target system, use the Windows Network Card Display options in the Spotlight on Windows Options window to choose the one whose data you want to display.

Related Topics
[NBT Page](#)
[Windows Network Card Display](#)

Physical Memory Usage Alarm

This alarm becomes active when a process is suffering a potential memory leak.

A memory leak occurs when a program continues requesting memory over a period without releasing any memory.

For some programs (such as Spotlight) this behaviour is expected as the program stores information over an extended period for dynamic near-history analysis.

Process CPU Utilization Alarm

This alarm becomes active when a process is sustaining a high CPU load.

When this alarm is active, you may need to:

- Investigate what the application is doing: a high CPU load may be an expected activity.

- Terminate the application if it is not responding. Some applications do not terminate through Windows Task Manager; if this does not work, use the KILL.EXE command from the Windows 2000 resource kit, or the TASKKILL.EXE command on Windows XP or Windows 2003 systems.

- Consider upgrading the system's processor.

Processor Queue Length Alarm

This alarm becomes active when the length of the processor thread queue reaches a threshold for sustained activity.

A long sustained queue length indicates a processor bottleneck - leading to overall system degradation.

To rectify this, you should look at:

- The **Processes** page of the **Processes** drilldown to see which Windows process is consuming the CPU.

- Possibly moving some processing activities to another system.

- Add an additional CPU to the system – the waiting threads will be load balanced between all available processors.

- Upgrade to a faster CPU. This will allow the CPU to process threads faster, but may not reap the same rewards as adding an additional CPU.

Related Topics

[Processes Page](#)

Read Hit Alarm

This alarm is activated when the proportion of file read requests that can be sourced from memory cache rather than disk drops below a threshold.

When this alarm is current you should:

- Look at the **NBT** page of the **Network** drilldown.

  If there is a large number of read requests coming from the network clients, the cache can be inundated and temporarily of little use.

- Look at increasing the physical memory of the system.

Related Topics
NBT Page

System Up Time Alarm

This alarm is activated when the system up time exceeds the value specified in the related threshold.

Total CPU Usage Alarm

The Windows CPU alarm is activated when the average CPU utilization of the system exceeds a threshold. This value is taken over a specific number of refresh intervals (by default, four).

Sustained high CPU utilization can adversely effect the performance of the system.

When this alarm is current, you should:

- Look at the **Processes** tab on the **Processes** drilldown to see which Windows process is consuming the CPU.

- Consider upgrading to a faster CPU or adding processors to your system.

- Look at the **Paging Activity** page on the **Memory** drilldown to see if there is a high paging rate. High paging rate can cause inflated CPU utilization. If this is the case, adding more memory to the system may overcome the problem.

Related Topics
Processes Page
Paging Activity Page

Virtual Address Space Alarm

This alarm becomes active when a process approaches the two gigabyte virtual address space limit imposed by Windows. Processes attempting to exceed this limit may fail catastrophically. Any process that approaches this limit should be closed to free the address space and then restarted if required.

Regardless of the amount of physical memory in your system, Windows uses a virtual address space of 4 GB, with 2 GB allocated to user-mode processes (for example, applications) and 2 GB allocated to kernel-mode processes (for example, the operating system and kernel-mode drivers).

NOTE: Some versions of Windows Server allow users to change this ratio to 3 GB for user-mode and 1 GB for kernel-mode processes via the /3GB switch in boot.ini. Versions that support this switch are:

- Windows 2000 Advanced Server

- Windows 2000 Datacenter Server

- Windows Server 2003 Enterprise Edition

- Windows Server 2003 Datacenter Edition.

Windows Server 2003 x64 Editions do not support the /3GB switch. On these systems:

- 32 bit processes compiled with the /LARGEADDRESSAWARE switch can increase their virtual address space limit to 4GB

- 64 bit processes can access up to 8TB of virtual address space.

Thus this alarm is disabled for 64 bit systems.

Virtual Memory Alarm

This alarm is activated when free virtual memory drops below a threshold.

When this alarm is current you should:

- Look at the **Processes** page on the **Processes** drilldown. Look at the VM Size (MB) column to see which applications are using the most virtual memory.

  Some applications (such as Microsoft Exchange or Microsoft SQL Server) can have their memory utilization limited.

- Close any superfluous processes.

- Look at increasing the size of the page file.

- Look at increasing the amount of RAM in the machine.

Related Topics

Processes Page

Windows CPU Alarm

The Windows CPU alarm is activated when the average CPU utilization of the system (taken over a specific number of refresh intervals; the default is four) exceeds a threshold. Sustained high CPU utilization can adversely effect the performance of the system.

When this alarm is current, you should:

- Look at the **Processes** page on the **Processes** drilldown to see which Windows process is consuming the CPU.

- Consider upgrading to a faster CPU or adding processors to your system.

- Look at the **Summary** page on the **Memory** drilldown to see if there is a high paging rate. High paging rate can cause inflated CPU utilization.

  If this is the case, adding more memory to the system may overcome the problem.

Related Topics
Processes Page
Summary Page

WMI Statistics Alarm

For Spotlight on Windows to retrieve WMI information, Spotlight must be able to make a WMI connection between the Spotlight machine and the remote machine.

This connection is made with the user credentials specified for this profile. If this user does not have administrator rights on the remote machine then the connection will fail with an access denied error. To rectify this, either make the specified user an administrator or change the credentials to those of an administrator.

You can continue to diagnose a machine without a WMI connection however some information, such as the **Disk Summary** drilldown, will not be available for this machine.

Note: For troubleshooting information, see WMI connection problems.

Related Topics
Disk Summary Page
WMI Connection Problems

Spotlight Drilldowns

When you have isolated a problem, you can display a drilldown page, whose charts and tables provide a detailed breakdown of the underlying statistics.

Following are Spotlight drilldowns.

Drilldown	Click to open	Description
Processes		Detailed information about all processes and services currently running on the machine you are diagnosing.
CPUs		Details of recent processor activity as measured by Windows.
Memory		Details about recent physical and virtual memory usage.
Disks		Information about the logical and physical disks on the system you are diagnosing.
Network		Detailed information about the network activity to and from the system being diagnosed.
Activity Summary		Summaries of recent Windows activity including, paging and packet rates, CPU usage across all processors, the number of threads waiting to be run, and the number of I/O requests that were queued for each logical disk.
Event Log		Recent event log items that have occurred on the target machine.
Single		Detailed information about specified individual processes (or a group of processes) on a

Drilldown	Click to open	Description
Application		Windows system. Choose the processes you want Spotlight to analyze by **View	Options	Data collection	Windows applications metrics**.
Alarm Log		Information on alarms, including the name of the component that issued the alarm, the date and time at which the alarm was logged, and the severity of the alarm.

Processes Page

The **Processes** page contains a table (grid) that lists all Windows processes that are currently running on the system. (A process is one instance of an application program that is currently executing on the Windows machine.)

To open the Processes page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Processes | Processes**.

Note: Not all of these columns will be visible by default. To view hidden columns, right-click a column heading and choose **Organize Columns**... from the shortcut menu. See Show, Hide & Order Columns for more information.

Column	Description
Process	The image name of the application. This can be used as a parameter in system programs, such as the TASKKILL.EXE command on Windows XP or Windows 2003 systems. **Notes**: • Click on a process to display sub pages. Process Details Page • Right click on the process and select **End Process** (if enabled) to end any processes currently running on the monitored machine (local or remote). • Right click on the process and select **Diagnose Process** (if enabled) to select the processes to monitor performance of in the Single Application drilldown. Windows Files Metrics • Some of the processes may be running other, secondary, processes. Right click the grid and select **Show as Tree** to display the secondary processes as branches of the initial process.
PID	Process ID. This is unique to each process running on the system. This can be used as a parameter in system programs, such as the TASKKILL.EXE command on Windows XP or Windows 2003 systems.
% CPU	The percentage of CPU time that the process is currently consuming. This is an instantaneous result.

Column	Description
Memory Usage (MB)	The current size of the working set of the process.
Virtual Memory Size (MB)	The current memory allocated to the process that cannot be shared with other processes.
Elapsed time	How long it has been since the process was started.
Handles	The overall number of resources that the process currently has open. A handle is a value used to uniquely identify a resource so that a process can access it.
Threads	The number of active threads in the process. A thread is a process execution unit.
Page Faults/sec	An instantaneous view of how many page faults are occurring for the process.
IO/sec	Shows the number of IO accesses (such as hard disk reads and writes and memory reads and writes) being performed by the process.
% Kernel	The percentage of CPU time that the process is currently consuming in privileged mode. (Privileged mode is designed for operating system components and allows direct access to hardware and all memory.)
% User	The percentage of CPU time that the process is currently consuming in user mode. (User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems.
Affinity	For systems with multiple CPUs, the process affinity, which can be set via the Windows Task Manager, specifies the CPUs that are permitted to run the current process. Use the Configure affinity option to tell Spotlight on Windows about the CPUs where the specified process is permitted to run. This enables Spotlight to report the CPU usage for that process with accuracy. Right click on the process and select **Configure Affinity** (if enabled). See Process Affinity Dialog for more information. **Note**: Spotlight itself CANNOT set process affinity. The Processes Page And Process Affinity
IO data bytes/sec	The rate at which the process is reading and writing bytes in all its I/O operations.
IO other bytes/sec	The rate at which the process is issuing bytes to I/O operations that do not involve data (control operations, for example).
IO other operations/sec	The rate at which the process is issuing I/O operations that do not involve data (control operations, for example).
IO read bytes/sec	The rate at which the process is reading bytes from I/O operations.

Column	Description
IO Reads/sec	Shows the number of IO reads (such as hard disk reads and memory reads) being performed by the process.
IO write bytes/sec	The rate at which the process is writing bytes to I/O operations.
IO Writes/sec	Shows the number of IO writes (such as hard disk writes and memory writes) being performed by the process.
Page file bytes	The current number of bytes that this process has used in the paging file(s).
Page file bytes peak	The maximum number of bytes that this process has used in the paging file(s).
Parent ID	The ID of the process that created the current process.
Peak Address Space (MB)	The peak size of the total address space of the process since it was started.
Peak Memory Usage (MB)	The peak size of the working set of the process since it was started.
Pool nonpaged bytes	The number of bytes of memory currently used by the server that cannot be paged out.
Pool paged bytes	The number of bytes of memory currently used by the server that can be paged out.
Priority	The priority of the process. Program priorities range from 1 to 31, and are dependant upon what the process is currently executing. Processes started in "Real Time" mode run with a priority of 16 to 31, whereas processes with "High", "Normal" or "Low" settings run in a priority range of 1 to 15.
Services	The services (if any) associated with the current process permitted to run the current process.
User	The owner of the process on a Windows Server machine (for example, Windows 2000 Advanced Server or Windows Server 2003).
Virtual Address Space (MB)	The current size of the total address space of the process. **Note**: A process is limited to 2GB of address space no matter how much free RAM may be available.

Process Details Page

The **Process Details** subpage contains a detailed list of properties for a selected process

To open the Process Details subpage

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Processes | Processes**.

3. Select a process to view its details

You can view the information that follows:

Item	Description
Process name	The image name of the application. This can be used as a parameter in system programs, such as the TASKKILL.EXE command on Windows XP or Windows 2003 systems.
Process ID	This is unique to each process running on the system, and can be used as a parameter in system programs, such as the TASKKILL.EXE command on Windows XP or Windows 2003 systems.
Memory Usage (MB)	The current size of the working set of the process.
Virtual Memory Size (MB)	The current memory allocated to this process that cannot be shared with other processes.
Virtual Address Space (MB)	The current size of the total address space of the process. **Note**: A process is limited to 2GB of address space no matter how much free RAM may be available.
Processor %	The percentage of CPU time that the process is currently consuming. This is an instantaneous result.
CPU User Mode %	The percentage of CPU time that the process is currently consuming in user mode. (User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems.)
CPU Kernel Mode %	The percentage of CPU time that the process is currently consuming in privileged mode. (Privileged mode is designed for operating system components and allows direct access to hardware and all memory.)
Elapsed Time	How long it has been since the process was started.
User	The owner of the process on a Windows Server machine (for example, Windows 2000 Advanced Server or Windows Server 2003).
Process Affinity	For systems with multiple CPUs, this shows the CPUs that are permitted to run the current process. See Process Affinity Dialog for more information.
Services	The services (if any) associated with the current process.

Item	Description
Handles	The overall number of resources that the process currently has open. A handle is a value used to uniquely identify a resource so that a process can access it.
Threads	The number of active threads in the process. A thread is a process execution unit.
Priority	The priority of the process. Process priorities range from 1 to 31, and are dependant upon what the process is currently executing. Processes started in "Real Time" mode run with a priority of 16 to 31, whereas processes with "High", "Normal" or "Low" settings run in a priority range of 1 to 15.
Page Faults/s	An instantaneous view of how many page faults are occurring for the process.
Reads/s	Shows the number of IO reads (such as hard disk reads and memory reads) being performed by the process.
Writes/s	Shows the number of IO writes (such as hard disk writes and memory writes) being performed by the process.
IO/s	Shows the number of IO accesses (such as hard disk reads and writes and memory reads and writes) being performed by the process.

Process History Page

The **Process History** subpage of the Processes drilldown opens when you click a process in the table (grid) in the Processes page, and then click the Process History tab.

This page shows the recent activity of the selected process in a series of related charts. By default, the charts in the Process History page now record data from the time that Spotlight on Windows was started.

To open the Process Details page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Processes | Processes**.

3. Select a process to view its details.

4. Click **Process History**

Charts on the Process History subpage
Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Processor	Shows the percentage of Privileged and User Time.
	High Privileged time means that the program is predominantly busy accessing resources through operating system requests.
	High User time means that the program is predominantly CPU-bound with the program code itself.
Read/Writes	Reads / Second – Shows the rate of IO reads (such as hard disk reads and memory reads) being performed by the process.
	Writes / Second – Shows the rate of IO writes (such as hard disk writes and memory writes) being performed by the process.
Memory	Shows the current memory allocated to the selected process.
	Memory Usage – Shows the current size of the working set of the selected process. The working set is the set of memory pages touched recently by the threads in the process.
	Virtual Memory Size – Shows the current size of the memory allocated to this process that cannot be shared with other processes.
Page Faults	Shows the number of page faults being generated by the program. A consistently high value may indicate a lack of memory. However, this metric includes both soft and hard page faults; as such, there may be no associated problems.

Note: The time frame displayed in the charts will depend on the historical settings you have chosen. See History Browser for more information.

Process Threads Page

The Process Threads sub-page of the Processes drilldown opens when you click a process in the table (grid) in the Processes page, and then click the Process Threads tab.

This page shows current information about the threads that execute the selected process, including:

To open the Process Details subpage

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Processes | Processes**.

3. Select a process to view its details.

4. Click **Process Threads**.

You can view the information that follows:

Item	Description
Instance	The instance name of the thread.
Thread ID	The unique identifier of the thread.
% CPU	The percentage of CPU elapsed time that the thread has used in order to execute

Item	Description
	instructions.
% User	The percentage of elapsed time that the thread has spent executing code in user mode.
% Kernel	The percentage of elapsed time that the thread has spent executing code in privileged mode.
Elapsed time	The total elapsed time the thread has been running.
Switches/sec	The rate of switches from one thread to another.
Thread state	The current state of the thread (Ready, Running or Waiting).
Thread Wait Reason	This is applicable only when the thread is in the Waiting state, and shows the reason why the thread is waiting.

Services Page

The **Services** page shows details of Windows services on the current system.

To open the Services page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Processes | Services**.

The **Services** table contains the following information.

Notes:

- Not all of these columns will be visible by default. To view hidden columns, right-click a column heading and choose **Organize Columns**... from the shortcut menu.

- To hide or show the tree structure of the grid, right-click the grid and choose **Hide Tree** or **Show Tree** from the shortcut menu.

Column	Description
Service	This shows the name of the service or driver as reported to the system. The "tree bar" to the left of the names is a dependency tree, showing all services or drivers that depend on the particular service.
Display name	The "friendly" name of the service or driver.
Start up	This shows how the service acts on Windows start: - Automatic – Starts every time the system starts, after the Boot and System devices start. - Manual – Requires manual startup or another service or device to request its startup.

Column	Description
	• Disabled – Does not start and cannot be started manually.
	• Boot – Starts every time the system starts, before any other devices start.
	• Demand – Starts when the device is detected or needed for a specific event.
	• System – Starts every time the system starts, after the Boot devices start.
Service type	Shows what type of program this is.
Current state	Shows what the current status is of the service or driver. The status can be Running, Not Running or Paused.
Controls accepted	Identifies what can be done with a service or driver. This information is only available for currently running or paused services.
Running PID	The ID of the process associated with the current service (if any).
Process path name	The location of the service executable file.

System Drivers Page

The **System Drivers** page shows details of the Windows kernel drivers and file system drivers on the current system.

To open the Services page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Processes | System Drivers**.

The **System Drivers** table contains the following information for each device.

Note: Not all of these columns will be visible by default. To view hidden columns, right-click a column heading and choose **Organize Columns…** from the shortcut menu.

Column	Description
Driver	The name of the driver as reported to the system. The "tree bar" to the left of the names is a dependency tree, showing all drivers that depend on the particular driver.
Description	The "friendly" name of the driver.
Type	Shows if the device is a Kernel or File System driver.
State	Shows if the driver is Running or Stopped.

Column	Description
Start Mode	This shows how the driver acts when Windows starts.

- Automatic – Starts every time the system starts, after the Boot and System devices start.
- Manual – Requires manual startup or another service or device to request its startup.
- Disabled – Does not start and cannot be manually started.
- Boot – Starts every time the system starts, before any other devices start.
- Demand – Starts when the device is detected or needed for a specific event.
- System – Starts every time the system starts, after the Boot devices start.

Column	Description
Accepts Stop	Identifies whether a driver can be stopped. This information is available only for currently running devices.
File	Shows the file location of the device. This cannot be retrieved for all devices.

CPUs Page

The **CPUs** page shows performance details of all the CPUs on the current Windows system. The **CPUs** page contains several charts.

To open the CPUs page

1. Select the **Spotlight** connection in the **Spotlight Browser**.
2. Click **CPUs**.

Charts on the CPUs page
Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.
- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.
- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
CPU Utilization	Shows the CPU utilization for all the processors on the system. Each processor is represented by a single line on the chart, and by an item in the legend at the top right of the chart. **Note:** To view CPU utilization analyzed according to processing mode, see the **Processing** page in the **CPUs** drilldown.
Server Work Queues	Shows the number of threads (program execution units) that are waiting to run on each processor. A sustained queue length greater than four might indicate processor congestion.

Chart	Description
Interrupts	The Interrupts chart shows the number of interrupts per second per CPU on the Windows system. Each CPU is represented by a single graph on the chart, and by an item in the legend at the top right of the chart. The Interrupts chart is designed to show if any CPUs on the system are bearing a disproportionate level of the I/O load.
Kilobytes Transferred	The Kilobytes Transferred chart shows the rate of kilobytes transferred per CPU on the Windows system. Each CPU is represented by a single graph on the chart, and by an item in the legend at the top right of the chart.

Processing Page

The **Processing** page on the **CPUs** drilldown shows the overall performance details of CPUs on the current Windows system. To view performance details for individual CPUs in the system, see the CPUs Page and Multiprocessor Page. The **Processing** page contains several charts.

To open the CPUs page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **CPUs | Processing**.

Charts on the Processing page

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Total CPU Utilization	Shows the combined CPU utilization for all processors on the system. Two data series are displayed: - User – The percentage of CPU time that is being consumed in User mode. (User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems.) - Kernel – The percentage of CPU time that is being consumed in privileged mode. (Privileged mode is designed for operating system components and allows direct access to hardware and all memory.) **Note:** To view CPU utilization for the individual CPUs in the system, see the CPUs Page.

Chart	Description
Total Processor Queue Length	Shows the total number of threads (program execution units) that are waiting to be run on ALL processors. A sustained processor queue length greater than three times the number of processors can indicate processor congestion. **Note:** To view Server Work Queues for the individual CPUs in the system, see the CPUs Page.
Total Process & Thread Count	This chart displays the total number of Windows processes and threads that exist. A process is one instance of an application program or system service that is currently executing on the server. Each process will have one or more threads which are the basic entity that can be scheduled. Sophisticated application processes such as SQL Server or Exchange can have dozens of threads running concurrently.
Context Switching	The Context Switching chart shows the rate at which Windows has been processing context switches. A Context Switch occurs when a processor switches from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service.

Multiprocessor Page

The **Multiprocessor** page in the **CPUs** drilldown displays in separate charts the total processor utilization for all CPUs on the Windows system.

To open the Multiprocessor page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **CPUs | Multiprocessor**.

Summary Page

The **Memory | Summary** page summarizes recent memory usage.

To open the Summary page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Memory | Summary**.

Charts on the Summary page
Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Physical Memory	The **Physical Memory** chart shows how much physical memory (RAM) Windows is using. Physical memory usage will normally remain close to the total amount of physical memory installed on the system unless the amount of physical memory you have exceeds the amount of virtual memory that Windows is using. Windows normally keeps some physical memory available (free) for immediate reuse.
Virtual Memory	The total amount of memory in use by the program. This includes Physical Memory and space in the paging file. A steady increase in virtual memory usage can indicate that a process on the system has a memory leak.
Paging	Shows the rate at which pages are being swapped in and out of memory. The chart displays two data series: • Page ins – The Page Ins value includes hard pages (paging requests that have to go to the paging file on disk) NOT soft pages (requests for memory pages that are not in the program's working set, but still in memory). • Page outs – The Page Outs value provides the number of write requests to the paging file on disk. A sustained high rate of paging can cause problems with overall system degradation due to disk thrashing and CPU load.
Physical Memory Usage	The Physical Memory Usage chart shows how Windows is using physical memory. It displays the following data series:

Series	Description
Free	Available physical memory not currently in use.
Processes	Memory being used by Windows processes. This will normally be the largest area by a significant degree.
Kernel	Memory being used by the Windows kernel.
File Cache	Memory that Windows is using to cache disk files in order to speed reads and writes.

Paging Activity Page

The **Memory | Paging Activity** page shows details of Windows paging activity and page files.

To open the Paging Activity page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Memory | Paging Activity**.

Charts and Grid on the Paging Activity page

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Charts and Grid	Description
Page File Transfers Chart	The Page File Transfers chart shows recent Windows paging activity. It shows the number of pages read (in) and written (out) per second to and from the Page Files. Paging occurs when the Windows Virtual Memory Manager moves data or code between physical memory and disk. Sustained high paging rates can degrade system performance significantly.
Page Faults Chart	The Page Faults chart shows the rate at which Windows is processing page faults. It compares the two types of page fault (soft and hard) and makes it easy to see the ratio between them. A page fault occurs when a process references a page that is not in that processes working set (the set of pages visible to that process in physical memory). When this happens, the process has to wait while the Windows Virtual Memory Manager retrieves the page from virtual memory. - A soft page fault occurs when Windows finds the required page somewhere in physical memory. - A hard page fault occurs when the page is not in physical memory and Windows has to read it from the page files. This is by far the more expensive of the two as it involves disk I/O. Hard page faults are the cause of paging and can degrade performance significantly.
Page Files Grid	The Page Files grid shows the page files in use by Windows. Page files are disk files that the Windows Virtual Memory Manager uses to back physical memory. Code and data is moved between physical memory and the page files as required, giving processes on the system the illusion that there is much more physical memory available than there really is. The process of moving data and code between memory and disk is called paging. The Page Files grid shows the following: **Column** — **Description** Location — Page File location (file name).

Charts and Grid	Description	
	Column	**Description**
	Size	The page file size (MB).
	Used	The current Used space (MB) in the page file.
	% Used	The current Used space as a percentage of the page file size.
	Peak %	The peak percentage of space used in the specified page file.

Cache Page

The **Memory | Cache** page shows detailed information about system cache.

To open the Cache page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Memory | Cache**.

Charts on the Cache page

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Cache Hit Rate	Shows a recent summary of the percentage of file requests that are satisfied by the file cache, and that do not require a disk read.
Cache Size	Shows the recent history of memory allocated to the file cache. Memory allocated to the file cache is dynamic. This is controlled by the Disk Cache Manager, and will alter the level of memory based upon how much physical memory is being used by other applications and is available in the system.
Physical Cache I/O	Shows the rate of Reads and Writes being made to the file cache.
Cache Faults	This chart displays the number of faults that occur when a requested page is not found in the file

system cache, and must be retrieved from elsewhere.

The chart displays the number of faults, NOT the number of pages faulted in each operation.

Logical Disk Activity Page

The **Disks | Logical Disk Activity** page shows information about the logical disks on this system. It includes I/O information, as well as disk space usage.

To open the Logical Disk Activity page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Disks | Logical Disk Activity**.

To view charts on the Logical Disk Activity page

Right-click the **Logical Disk Activity** page and select **View as Chart**.

Notes:

- Each disk is represented by a single line on the chart, and by an item in the legend at the right of the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Disk Reads	Shows a recent summary of the rate at which read requests have been sent to each logical disk.
Disk Writes	Shows a recent summary of the rate at which write requests have been sent to each logical disk.
Disk I/O	Shows a recent summary of the rate at which read and write requests have been sent to each logical disk. This chart is the sum of the **Disk Reads** and **Disk Writes**.
Disk Queue	The Disk Queue chart shows the number of I/O requests that were queued for each logical disk. Disk queue length indicates how heavily loaded a disk subsystem is. High queue lengths mean the disks are struggling to process the I/O load being put on them.
Transfer Time	The Transfer Time chart shows how long it is taking for data to be transferred between disk and memory, and includes both **Disk Reads** and **Disk Writes**. If disk transfers are taking consistently longer than 50 ms, a disk bottleneck may be developing.
Disk Load	Shows how busy overall the individual logical disks are.

To view the Logical Disk Activity page as a table

Right-click the **Logical Disk Activity** page and select **View as Grid**.

In table (grid) mode, the page displays the information contained in the charts, plus additional information in the following columns:

Column	Description
KBs Read/s	The rate (in kilobytes/s) at which data is read from the specified logical disk.
KBs Written/s	The rate (in kilobytes/s) at which data is written to the specified logical disk.

Physical Disk Activity Page

The **Disks | Physical Disk Activity** page shows information about the logical disks on this system. It includes I/O information, as well as disk space usage.

To open the Logical Disk Activity page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Disks | Physical Disk Activity**.

To view charts on the Physical Disk Activity page

Right-click the **Physical Disk Activity** page and select **View as Chart**.

Notes:

- Each disk is represented by a single line on the chart, and by an item in the legend at the right of the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Disk Reads	Shows a recent summary of the rate at which read requests have been sent to each physical disk.
Disk Writes	Shows a recent summary of the rate at which write requests have been sent to each physical disk.
Disk I/O	Shows a recent summary of the rate at which read and write requests have been sent to each physical disk.
Disk Queue	The Disk Queue chart shows the number of I/O requests that were queued for each physical disk. Disk queue length indicates how heavily loaded a disk subsystem is. High queue lengths mean the disks are struggling to process the I/O load being put on them.
Transfer Time	The Transfer Time chart shows how long it is taking for data to be transferred between disk and

Chart	Description
	memory, and includes both **Disk Reads** and **Disk Writes**.
	If disk transfers are taking consistently longer than 50 ms, a disk bottleneck may be developing.
Disk Load	Shows how busy overall the individual physical disks are.

To view the Physical Disk Activity page as a table

Right-click the **Physical Disk Activity** page and select **View as Grid**.

In table (grid) mode, the page displays the information contained in the charts, plus additional information in the following columns:

Column	Description
KBs Read per sec	The rate (in kilobytes/s) at which data is read from the specified physical disk.
KBs Written per sec	The rate (in kilobytes/s) at which data is written to the specified physical disk.

Logical Disk Space Usage Page

The **Disks | Logical Disk Space Usage** displays a graphical view of used and available space on all logical drives.

To open the Logical Disk Space Usage page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Disks | Logical Disk Space Usage**.

Note: Use the **Disk Space By…** control to switch between **Disk Space (GB)** and **Disk Space (%)**.

Disk Summary Page

The **Disk Summary** page provides summary information about the physical disks on the current system.

To open the Disk Summary page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Disks | Disk Summary**.

For each physical disk detected, disk summary information is displayed in the columns of a single table (grid).

Column	Description
Disk	Each physical disk in the system is listed in the grid in the order it has been discovered. The first disk is numbered 0 (zero).

Column	Description
Description	A description of each disk discovered. This displays RAID details where applicable.
Type	The type of hard disk discovered (for example, IDE).
Total GB	The total size of the disk in gigabytes.
Device type	More detailed information about the type and configuration of the disk.
Hardware vendor	The vendor code associate with each hard disk.
Partitions	The number of partitions discovered on each disk.

To open the Logical Disks sub page

To provide equivalent details about the logical disks associated with each physical disk:

Click on the entry for the physical disk on the **Disk Summary** page.

At any one time, the Logical Disks grid displays information only on the logical disks associated with the physical disk currently highlighted in the Disks grid.

Column	Description
Drive	The name of the logical disks associated with the physical disk specified in the Disk Summary page.
Label	The label assigned to the logical disk (if any).
File System	The file system used on the logical disk (for example, NTFS).
Total MB	The total amount of space on the logical disk, measured in megabytes.
Free MB	The total amount of free space on the logical disk, measured in megabytes.
Compressed	True or False. Whether the information on the logical disk is maintained in compressed format.

File Sizes Page

The **File Sizes** page displays a grid that contains information on the size of files that you are tracking from the **Windows Files Metrics** dialog.

To open the File Sizes page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Disks | File Sizes**.

The File Sizes grid contains the following information:

Chart	Description
File Name	The full pathname of the selected file.
Current	The current size of the file, measured in megabytes.

Chart	Description
Size (MB)	
Max Size (MB)	The maximum size for this file as set in the **Windows Files Metrics** dialog. Setting the maximum file size does NOT prevent the specified file from growing beyond this size. However, if the size of the file exceeds that maximum size, Spotlight raises a critical alarm.
% Max Size	The current size of the file, measured as a percentage of the maximum size.
Severity	The degree of urgency associated with the current size. The severity controls what action Spotlight takes. (For example, if the size of the file exceeds the maximum size, the severity becomes critical, and Spotlight raises a critical alarm.)
File exists	Whether or not the file exists at the specified location

Tip: You can add / remove / configure the tracking of file size. Right-click in the File Sizes grid and select **Files Options**. Windows Files Metrics

Network Page

The **Network** page shows recent network activity to and from this system.

To open the Network page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | Network**.

Charts on the Network page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Total packets in/out	Shows the incoming and outgoing packet rates as an aggregate across all network cards in the system. Packets Sent and Packets Received are each represented by a single line on the chart, and by an item in the legend at the right of the chart.
Packets by network card	Shows the total rate of packets (incoming and outgoing), broken down by network card. Each network card on the current system is represented by a single line on the chart, and by an item in the legend at the right of the chart.

Chart	Description
Total kilobytes in/out	Shows the level of traffic being received and sent by the system in kilobytes per second. Kilobytes Sent and Kilobytes Received are each represented by a single line on the chart, and by an item in the legend at the right of the chart.
Errors and Retries	Shows the number of errors and retries on the network subsystem. A consistent number of retries or errors usually means one of the following: - A network segment attached to the system is over-utilized. - There are problems with a network card, cabling or other networking device on a network segment. Inbound Errors and Outbound Errors are each represented by a single line on the chart, and by an item in the legend at the right of the chart.

TCPIP Page

The **TCPIP** page shows recent TCP/IP activity to and from the current system.

TCP/IP (Transmission Control Protocol/Internet Protocol) is one means by which computers communicate across a network. It is the basic protocol used for communication across the Internet.

Communication programs that use the TCP/IP protocol are built in two layers:

- The **TCP layer** converts a data stream into a series of packets for transmission across a network. When these packets are received at their destination, another TCP layer reassembles the data stream.
- The **IP layer** ensures that data packets reach the right destination.

To open the TCPIP page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | TCPIP**.

Charts on the TCPIP page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.
- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
TCP Segments	When TCP transfers a stream data, it breaks up the stream into small segments. The TCP Segments chart displays three different graphs representing the rate at which TCP segments are transmitted or received in the current Windows system.

Chart	Description

Graph	Description
Sent	The rate at which TCP segments are transmitted by the current system.
Received	The rate at which TCP segments are received by the current system.
Retransmitted	The rate of transmission of TCP segments that have already been sent, but need to be retransmitted because of packet time-out or some other transmission error.

Sent, Received and Retransmitted are each represented by a single line on the chart, and by an item in the legend at the right of the chart.

Chart	Description
IP Fragments	When data packets are transmitted across a network via TCP/IP, they may be further fragmented en route, and need to be reassembled. The IP Fragments chart displays the rates at which the fragmented packets are received and rebuilt.

The graphs on the chart display one of the following features of IP fragmentation:

Graph	Description
Received	The rate at which IP fragments are received successfully.
Created	The rate at which IP datagram fragments are generated as a result of fragmentation.
Re-assembled	The rate at which IP fragments are rebuilt into their original data packets. Packets are fragmented when they travel through a router that needs to send packets that are smaller than the packets received.
Datagrams	The rate at which IP fragments are created on the current system. This applies only to routed packets.
Failures	The rate at which the current system receives data packets that are too large to be transmitted, and that cannot be fragmented. The cause may be the presence of a "do not fragment" flag in the IP packet header.
Reassembly failures	The rate at which errors are reported when IP fragments are reassembled into data packets. This may be due to an error in one or more fragments.

Each data series in the IP Fragments chart is represented by a single line, and by an item in the legend at the right of the chart.

Chart	Description
UDP Datagrams	UDP (User Datagram Protocol) is an alternative communications protocol to the Transmission Control Protocol (TCP).

Unlike TCP, UDP does not provide a sequencing service, so when UDP datagrams (packets) arrive at a destination, they are reassembled into a complete message via an end-point application that is active on the port(s) specified in the datagram. The UDP Datagrams chart displays three different graphs representing the rate at which UDP

Chart	Description

datagrams are transmitted or received in the current Windows system.

Graph	Description
Sent	The rate at which UDP datagrams are transmitted by the current system.
Received	The rate at which UDP datagrams are received by the current system.
No port	During transmission, UDP provides port numbers to distinguish different user requests. The No Port metric displays the number of packets received per second that do not have an end-point application active on the specified port.

Sent, Received and No Port are each represented by a single line on the chart, and by an item in the legend at the right of the chart.

Chart	Description
IP Datagrams / IP Datagram Errors	IP datagrams (packets) are the message units that are transmitted across a network (and across the Internet) via the Internet Protocol.

Use the IP Datagrams control to choose to view IP Datagrams or IP Datagram Errors.

The IP Datagrams chart:

The rate at which IP packets are sent, received, and routed to other destinations.

Graph	Description
Sent	The rate at which IP datagrams (packets) are being sent by the current system.
Received	The rate at which IP datagrams are being received by the current system.
Forwarded	The rate at which the current system is routing IP datagrams to another destination.
Received delivered	The rate at which the system successfully receives and accepts IP packets.

The IP Datagram Errors chart:

The number of errors that occur when IP packets are sent, received, and routed to other destinations.

Graph	Description
Outbound discarded	The number of outbound packets discarded because of an issue unrelated to the packets themselves (for example, if the send buffer is full).
Outbound	The number of outbound packets discarded because the system cannot

Chart	Description

Graph	Description
no route	route the packets to the destination IP address.
Received address errors	The number of times that the system has received packets that do not have a valid return address.
Received discarded	The rate at which received packets are discarded because of an issue unrelated to the packets themselves (for example, if the receive buffer is full).
Received header errors	The number of times that the system has received packets that have errors in the IP header area (for example, a packet checksum error).
Received unknown protocol	The number of times that the system has received packets that are correctly addressed, but that use a protocol unsupported by the IP handler on the system.

Notes:

- Each data series in the IP Datagrams chart is represented by a single line, and by an item in the legend at the right of the chart.
- The information represents the number of errors detected since the chart was last refreshed.

NBT Page

The **NBT** page shows the machines that have a connection to the diagnosed system, and the corresponding level of traffic (in kilobytes per second) being generated between the connections.

Note: This is only for NBT (NetBios over TCP/IP) connections – that is, only Microsoft Networking connection over TCPIP.

To open the NBT page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | NBT**.

SPX Page

The **SPX** page shows recent SPX activity to and from the current system.

SPX (Sequenced Packet Exchange) is a communication protocol used in Novell Netware networks. It is similar to TCP, and provides connection services between nodes on a network. The combination of SPX with IPX (Internetwork Packet Exchange) provides connection services similar to TCP/IP.

To open the SPX page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | SPX**.

Charts on the SPX page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Bytes In/Out	The rate at which bytes are handled by the IPX/SPX protocol. • Sent – Number of bytes sent using the IPX/SPX protocol • Received – Number of bytes received by the IPX/SPX protocol
Frames In/Out	The rate at which SPX frames (equivalent to IP datagrams) are handled by the IPX/SPX protocol. • Sent – Number of SPX frames sent per second • Received – Number of SPX frames received per second
Errors (static)	This chart shows all of the SPX static errors counted since the machine was last restarted. The graphs on the chart are: • Adapter – The number of errors in receiving SPX packets due to adapter problems. • Link – The number of errors in receiving SPX packets due to physical network (link) issues such as ethernet problems. • No Listen – The number of errors in receiving SPX packets due to a remote computer not being found. • Resource local – The number of errors in receiving SPX packets due to a lack of computer resources on the local computer. • Resource remote – The number of errors in receiving SPX packets due to due to a lack of computer resources on the remote computer.
Errors (dynamic)	This chart shows the rate at which SPX dynamic errors occur. The graphs on the chart are: • Re-sent – The number of times per second that frames need to be resent for any reason. • Rejected – The number of data frames per second that are rejected from this computer.

Sessions Page

The sessions displayed in this page are the active connections made to the current Windows system by users on other systems.

To open the SPX page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | Sessions**.

The Sessions grid on this page displays data for each connection in the following columns:

Column	Description
Computer	The name of the remote machine connected to the current system.
Elapsed time	How long the connection has been maintained.
Idle time	The part of **Elapsed time** when the connection has been inactive.
Domain	The domain of the user session currently logged in to the remote machine. Domain information is displayed only for Terminal Services and Remote Desktop sessions.
User	The user currently logged in to the remote machine.
Type	The type of user session, for example "Terminal Services" or "Remote Desktop". For normal sessions, the Type column is blank. Remote Desktop sessions are shown only for Windows Server machines (for example, Windows 2000 Advanced Server or Windows Server 2003).
State	The current state of the session. Some examples are "active", "disconnected", and "connecting".
Logon Time	The time that the user last logged in to the remote desktop session. This column shows data only when the session type is Remote Desktop.
Disconnected Time	The time that the user last disconnected from the remote desktop session. This column shows data only when the session type is **Remote Desktop** and the state is **Disconnected**. If the user logs out of the remote desktop session, the session is closed.

Shares Page

The **Shares** page on the Network drilldown displays the resources on the monitored Windows system that can be shared with users on other systems. Each row on the Shares grid represents a shared resource on the current system. The columns in the grid contain relevant data for shared resources.

To open the Shares page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | Shares**.

Note: For each shared resource, you can view more details about the connections currently established to the resource. Double click on the resource's entry in the grid. See the Open Connections Page and the Permissions Page.

Column	Description
Share	The name of the shared resource.
Local Path	The location of the shared resource.
Comment	A short, meaningful description of the resource.
Type	The type of device that provides the shared resource (for example, Disk or Printer)
Max Users	The maximum number of users who can simultaneously share the resource.
Current Users	The number of users who are now sharing the resource.

Open Connections Page

The Open Connections sub-page provides details about the connections currently established to the shared resources shown on the Shares page of the Network drilldown.

To open the Open Connections subpage

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | Shares**.

3. Double click on the resource in the grid.

4. Click **Open Connections**.

For each shared resource, information on its open connections is displayed in the columns of a single table (grid).

Column	Description
Share	The name of the shared resource.
Computer	The remote machine connected to the shared resource.
User	The user logged in to the remote machine.
Type	The type of device that provides the shared resource (for example, Disk or Printer)
Open Files	The number of files currently open on the shared resource.
Elapsed Time	The length of time that the user on the remote machine has been connected to the shared resource.

Permissions Page

The **Permissions** page lists the trustees assigned to the shared resource currently chosen in the Shares page of the Network drilldown, together with the rights granted or denied to the trustees.

Note: Microsoft defines a trustee as:

"The user account or group account to which an access control entry (ACE) applies. Each ACE in an access control list (ACL) has one security identifier (SID) that identifies a trustee. User accounts include accounts that users or programs such as Windows services use to log on to the local computer. Group accounts cannot be used to log on to a computer, but are useful in ACEs to allow or deny a set of access rights to one or more user accounts."

To open the Permissions page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | Shares**.

3. Double click on the resource in the grid.

4. Click **Permissions**.

For each shared resource, information on its assigned permissions is displayed in the columns of a single table (grid).

Note: Not all of these columns will be visible by default. To view hidden columns, right-click a column heading and choose **Organize Columns**... from the shortcut menu. See Show, Hide & Order Columns for more information.

Column	Description
User	The name of the trustee identified by the User SID.

User Type	The trustee can be one of the following types of user:
	
	• UNKNOWN – The trustee type is unknown, but not necessarily invalid.
	
	• USER – A Windows user account.
	
	• GROUP – A Windows group account.
	
	• DOMAIN – A domain.
	
	• ALIAS – A unique SID with a specified User alias.
	Example, BUILTIN\Administrators can have different local names on different resources:
	
	   • serverA\Administrators on server A
	   • serverB\Administrators on server B
	
	• KNOWN GROUP – A well known group (for example, the **Authenticated Users** or **Everyone** groups, which are available on all Windows systems).

Grant	The access types (READ ONLY, CHANGE, FULL CONTROL) that have been granted to the current trustee.
Deny	The access types that have been specifically DENIED to the current trustee. **Deny** rights take precedence over **Grant** rights.
Full Name	The full name of the trustee. Hide this column if you do not need this information, as Spotlight obtains it by querying the domain controller for the remote machine; this may cause a significant delay when refreshing data. See Show, Hide & Order Columns for more information.

Column	Description
Share	The name of the shared resource.
User SID	The security identifier for the current trustee of the shared resource. Every account on a network is issued a unique SID when the account is first created.

Open Ports Page

The **Open Ports** page is not available if the Windows SNMP protocol is not installed on:

- The Windows machine under diagnosis, AND
- The machine running the Spotlight client.

Installing The SNMP Protocol

The open communication ports on a Windows machine are not necessarily in use. They may be the inadvertent consequence of a now-stopped service or some other application, or they may have been opened by a virus.

The **Open Ports** page on the Network drilldown summarizes the status of all the communication ports on the machine under diagnosis. Each row in the table (grid) represents an open port, and the columns in the grid display the properties associated with it.

To open the Open Ports page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | Open Ports**.

Note: In the Open Ports page, the term **local** refers to the Windows machine under diagnosis. **Remote** refers to any machines connected to a port on the diagnosed (local) machine.

Column	Description
Protocol	The communication protocol associated with the open port (for example, **TCP** or **UDP**).
Local Address	The address of the diagnosed machine.
Local Port	The registered name of the port on the diagnosed machine, if available. Otherwise it is the number of the port on the diagnosed machine.
Remote Address	The IP address of a remote host connected to the specified **local port**.
Remote Port	The port number by which the specified remote host is connected to the **local port**.
Status	The current status of the network connection between the two machines (for example, **Listening**, **Established**, or **Time_wait**.)

cLAN Gigabit Page

The **cLAN Gigabit** page of the Spotlight on Windows Network drilldown shows recent activity to and from the current system via cLAN Gigabit cards.

Note: The cLAN Gigabit network interface card (NIC) enables high-speed communications across a network. This page in the Network drilldown is enabled only if the current network contains one or more of these cards.

To open the cLAN Gigabit page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Network | cLAN Gigabit**.

Charts on the cLAN Gigabit page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Total MB in/out	Sent – The rate at which megabytes are sent over cLAN Gigabit network cards from this machine. Received – The rate at which megabytes are received by this machine from cLAN Gigabit network cards.
Frames per second	Sent – The rate at which frames are sent over cLAN Gigabit network cards from this machine. Received – The rate at which frames are received by this machine from cLAN Gigabit network cards.
Interrupts per second	Async sent – The rate at which interrupt requests are sent by the cLAN network cards (such as those to CPUs). Async received – The rate at which interrupt requests are received by the cLAN network cards.
Errors	**Note:** This chart does NOT show the **rate** of errors detected, but the number of errors detected **since the chart was last refreshed**. VPI/VCI Rx errors – The number of receive errors in VCI (Virtual Channel Identifier) and VPI (Virtual Path Identifier) packets. Rx NIB congestion – The number of receive errors originating from the NIB (Network Interface Board). Rx Cell Drop – The number of received packets dropped for any reason. Received HEC Errors – The number of received packets with HEC (Header Error Control) errors. Received 10b characters – The number of invalid 10 bit characters received.

Activity Summary Page

The **Activity** page displays summaries of recent Windows activity.

To open the Activity Summary page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Activity Summary**.

Charts on the Activity Summary page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Processor	Shows the total CPU utilization across all processors in the machine, categorized into **User Time** and **Privileged Time**. • **User Time** is the total time the processor(s) spends in user mode. This is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. • **Privileged Time** is the time the processor(s) spends in privileged mode. This is designed for operating system components and allows direct access to hardware and all memory. **Note:** The User Time and Privileged Time series are each represented by a single line, and by an item in the legend at the right of the chart.
Paging	Shows the rate at which pages are being swapped in and out of memory. • The **Paging Rate (In)** (or Page Reads) value includes hard pages (paging requests that have to go to the paging file on disk) - not soft pages (requests for memory pages that are not in the program's working set, but still in memory). • The **Paging Rate (Out)** (or Page Writes) value provides the rate of write requests to the paging file on disk. A sustained high rate of paging can cause problems with overall system degradation due to disk thrashing and CPU load. **Note**: Each Paging Rate series is represented by a single line, and by an item in the legend at the right of the chart.
Network Packets	Shows the incoming and outgoing packet rates as an aggregate across all network cards in the system. **Note:** The Packets Sent and Packets Received series are each represented by a single line, and by an item in the legend at the right of the chart.

Chart	Description
Processor Queue Length	Shows the total number of threads (program execution units) that are waiting to be run on ALL processors. A sustained processor queue length greater than three times the number of processors can indicate processor congestion. **Note:** To view Server Work Queues for the individual CPUs in the system, see the CPUs Page in the CPU drilldown.
Disk Queue Length	Shows the number of I/O requests that were queued for each logical disk. Disk Queue Length indicates how heavily loaded a disk subsystem is. High queue lengths mean the disks are struggling to process the I/O load being put on them. **Note:** Each logical disk on the system is represented by a single line on the chart, and by an item in the legend at the right of the chart.
Memory	Shows how much physical memory (RAM) and Virtual Memory Windows is using. Physical Memory usage will normally remain close to the total amount of physical memory installed on the system unless the amount of physical memory you have exceeds the amount of virtual memory that Windows is using. Windows normally keeps some physical memory available (free) for immediate reuse. Virtual Memory usage increases and decreases as Windows processes requests and gives up memory and as processes are started and stopped. A steady increase in Virtual memory usage can indicate that a process on the system has a memory leak. **Note**: The Virtual Used and Physical Used series are each represented by a single line, and by an item in the legend at the right of the chart.

Event Log Page

The Windows Event Log is where application or operating system information is written and can be accessed by system administrators. Spotlight on Windows allows you to analyze event data from your Windows system's various logs, and to display the relevant data in the Event Log drilldown.

Note: If the content of this page is disabled then Event Logs are disabled. You can enable them. Windows Event Log Metrics

To open the Event Log page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Event Log**.

Page Control	Description
Event Log	Select an event log to show in the Event Grid. Depending on the monitored system, the logs to select from may include: • Any – Combine all event logs.

Page Control	Description
	- Application log – General application events and some system service messages. - Directory Service log – Directory service events (only for Windows domain controllers). - DNS Server log – DNS server events (only for Windows DNS servers). - File Replication Service log – File replication service events (only for Windows domain controllers). - Security log – Operating system audit events. Logged according to the audit settings that have been configured on the system. - System log – Operating System and most system service messages.
Event Grid	A display of the selected event log.
Total	The number of entries in the selected log.
Days	The period of time for which the oldest entries in the log have been kept.

About the Event Grid

Notes:

- Not all of these columns will be visible by default. To view hidden columns, right-click the grid heading and choose **Organize Columns**... from the shortcut menu.
- Click an event to display the **Event Log Details** page.
- Right click the body of the grid and select **Event Log Options** to configure the Event Log. You can filter, alarm or expand the list of Event Log items. Windows Event Log Metrics.

Column	Description
Date/Time	The time when the entry was entered into the Windows event log, not when it was brought into Spotlight on Windows.
Event ID	A code that identifies the event.
Message	A description of the event.
Source	The software that logged the event.
Type	The type of event: - Warning – There is something out of the ordinary but not an outright failure. - Error – Something has failed. - Information – General information on what is happening on the system. - Security – There are numerous security based message types. These come from the security log and the system auditing that has been configured on the system.
Category	The category of the event. This is used in auditing, and is primarily used by the security log.

Column	Description
Log	Where the error message was generated: • Application log – General application events and some system service messages. • Directory Service log – Directory service events (only for Windows domain controllers). • DNS Server log – DNS server events (only for Windows DNS servers). • File Replication Service log – File replication service events (only for Windows domain controllers). • Security log – Operating system audit events. Logged according to the audit settings that have been configured on the system. • System log – Operating System and most system service messages. **Note**: The Directory Service and File Replication Service event logs are available only for Windows domain controllers. The DNS Server event log is available only for DNS servers.
Severity	Shows information on alarms associated with the event log item. This is assigned by the Spotlight on Windows event log rules (available through the Windows Event Log Metrics in Spotlight on Windows Options).
Clear	Select to acknowledge the event and return its status to Normal.
Computer	The name of the machine where the event occurred.
User	The user account which instigated the event log item. This is particularly pertinent to security items.

Summary Page

The Summary Page of the Single Application drilldown provides an overview of the Windows performance for the one or more applications chosen in the Windows Applications Metrics options of the Spotlight on Windows Options window.

To open the Summary page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Single Application | Summary**.

Charts on the Summary page

Notes:

• Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

• Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

- To choose applications to measure the performance of, right click the page and select **Select Application**. This opens the Windows Applications Metrics dialog. Windows Files Metrics

Chart	Description
Processors	For the specified applications, this corresponds to the CPU Utilization chart chart in the CPUs page of the CPU drilldown.
Paging	For the specified applications, this corresponds to the Paging chart in the Summary page of the Memory drilldown.
Virtual Memory Size	For the specified applications, this corresponds to the Virtual Memory Size graph on the Memory chart in the Process History page of the Processes drilldown.
Processor Queue Length	For the specified applications, this corresponds to the Processor Queue Length chart in the Activity Summary drilldown.
Disk Queue Length	For the specified applications, this corresponds to the Disk Queue Length chart in the Activity Summary drilldown.
Network Packets	For the specified applications, this corresponds to the Network Packets chart in the Activity Summary drilldown.

CPUs Page

The **CPUs** Page of the Single Application drilldown provides an overview of the CPU performance for the one or more applications chosen in the Windows Applications Metrics options of the Spotlight on Windows Options window.

To open the CPUs page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Single Application | CPUs**.

Charts on the CPUs page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

- To choose applications to measure the performance of, right click the page and select **Select Application**. This opens the Windows Applications Metrics dialog. Windows Files Metrics

Chart	Description
CPUs	For the specified applications, this corresponds to the sum of all graphs on the Processor chart in the Process History page of the Processes drilldown.
CPU per Process	This displays the CPU resources consumed by each application chosen for analysis in the Single Application drilldown. For the specified applications, this corresponds to the %CPU

Chart	Description
	column in the Processes grid in the Processes drilldown.
Queue Length	For the specified applications, this corresponds to the Server Work Queues chart on the CPUs page of the CPU drilldown.
IO Operations	For the specified applications, this corresponds to the IO per second column in the Processes grid on the Processes drilldown.

Memory Page

The Memory Page of the Single Application drilldown provides an overview of the memory performance for the one or more applications chosen in the Windows Applications Metrics options in the Spotlight on Windows Options window.

To open the Memory page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Single Application | Memory**.

Charts on the Memory page
Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

- To choose applications to measure the performance of, right click the page and select **Select Application**. This opens the Windows Applications Metrics dialog. Windows Files Metrics

Chart	Description
Memory Usage: Current OR Peak	If the **Current** option is chosen from the associated drop-down menu, the chart shows the current size of the working set of the selected process. The working set is the set of memory pages touched recently by the threads in the process. If the **Peak** option is chosen from the associated drop-down menu, the chart shows the peak size of the working set of the selected process(es) since they were started.
Virtual Memory Size	For the specified applications, this corresponds to the **Virtual Memory Size** graph on the Memory chart in the Process History page of the Processes drilldown.
Virtual Address Space: Current OR Peak	If the **Current** option is chosen from the associated drop-down menu, the chart shows the current size of the total address space of the process. A process is limited to 2GB of address space no matter how much free RAM may be available. If the **Peak** option is chosen from the drop-down menu, the chart shows the peak total address space of the selected process(es) since they were started.

Chart	Description
Page Faults	For the specified applications, this corresponds to the Page Faults chart in the Process History page of the Processes drilldown.

Disks Page

The Disks Page of the Single Application drilldown provides an overview of the disk performance for the one or more applications chosen in the Windows Applications Metrics options of the Spotlight on Windows Options window.

To open the Disks page

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Single Application | Disk**.

Charts on the Disks page

Notes:

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

- To choose applications to measure the performance of, right click the page and select **Select Application**. This opens the Windows Applications Metrics dialog. Windows Files Metrics

Chart	Description
Disk Reads	For the specified applications, this corresponds to the Disk Reads chart in the Logical Disk Activity page of the Disks drilldown.
Disk Writes	For the specified applications, this corresponds to the Disk Writes chart in the Logical Disk Activity page of the Disks drilldown.
Disk I/O	For the specified applications, this corresponds to the Disk I/O chart in the Logical Disk Activity page of the Disks drilldown.
Disk Queue	For the specified applications, this corresponds to the Disk Queue chart in the Logical Disk Activity page of the Disks drilldown.
Disk Transfer Time	For the specified applications, this corresponds to the Transfer Time chart in the Logical Disk Activity page of the Disks drilldown.
Disk Load	For the specified applications, this corresponds to the Disk Load chart in the Logical Disk Activity page of the Disks drilldown.

Alarm Log

The **Alarm Log** drilldown displays information on the alarms associated with the selected connection.

To open the alarm log

1. Select the **Spotlight** connection in the **Spotlight Browser**.

2. Click **Alarm Log**.

Notes:

- The Alarm Log drilldown is common to all Spotlight applications. See Spotlight Alarms for more information.

- The alarms are specific to the current Spotlight connection. Spotlight Alarms

Spotlight on Windows - View | Options

You can set a number of options that affect the behavior and display of Spotlight on Windows via the Spotlight on Windows Options window.

To open the Spotlight on Windows Options window
Click **View | Options**.

Option	Description
Windows Disks Metrics	Configure how disks are represented in the Disks panel on the Spotlight on Windows home page. These options will NOT be displayed if the system you are connected to contains only one logical drive.
Windows Network Card Display	Configure which network card is represented in the Network panel in the Spotlight on Windows home page. This option will NOT be displayed if the system you are connected to contains only one network card.
Windows Applications Metrics	Specify one or more processes running on the current system for detailed analysis.
Windows Event Log Metrics	Create Event Log rules and configure other Event Log settings.
Windows Files Metrics	Specify one or more files whose file size you want Spotlight to track.

Windows Disks Metrics

Use Windows Disks Metrics to configure how logical disks are represented in the Disks panel on the Spotlight home page.
Note: If the monitored system contains only one logical drive then these options will not be displayed.

To open the Windows Disks Metrics dialog

1. Click **View | Options**
2. Click **Data Collection | Windows Disks Metrics**

The general appearance of the Disks panel for all Spotlight on Windows connections

Option	Description
Allow up to eight disks on home page	By default, the Disks panel on the home page displays the details of four logical disks in the current connection PLUS four "disk container" images that show how full each disk is. Choose this option to display up to eight disks in the Disks panel. If you do, NO container images are shown.
Disk flow units	Choose how you want to display disk-related dataflows: - The number of Read and Write operations performed per second -OR- - The amount of data read and written per second (measured in KB/s).

Connection name - How disk performance is shown for the current Spotlight connection

Option	Description
Automatically configure disk display	When selected, Spotlight determines which disks are displayed in the Disk panel. Deselect this option if you want to configure how disks are displayed in the Disks panel.
By Drive letter/ By Lowest % Free Space	By Drive Letter - Display disk information according to the alphabetical label assigned to each disk. By Lowest % Free Space - Display disk information for disks having the lowest percentage of free space. **Note**: This option is enabled when **Automatically configure disk display** is selected.
Number of disks to show on home page	Type or select the number of disks to be displayed in the Disks panel. **Note**: This option is enabled when **Automatically configure disk display** is deselected.
Disk to edit (numbered top to bottom)	Type or select the number of the disk you want to edit. Disks are numbered from top to bottom starting at one. The number of available disks depends on how many disks you have chosen to display on the home page. **Note**: This option is enabled when **Automatically configure disk display** is deselected.
Display information for	Select the type of information you want to display for the selected disk (Logical Disk, Lowest % Free Space, Smallest Disk Size, Most "Read from" disk, or Most "written to" disk). **Note**: This option is enabled when **Automatically configure disk display** is deselected.
Logical Disk to display	Select the disk whose details you want to display. **Note**: This option is visible when **Display information for** is set to **Logical Disk**.

Windows Network Card Display

The Windows Network Card Display option lists all the NICs (network interface cards) on the Windows system that Spotlight on Windows is connected to.

To select the NIC whose Theoretical Bandwidth Limit you want to show on the Spotlight on Windows home page

1. Click **View | Options**
2. Click **Data Collection | Windows Network Card Display**
3. Select the network card to diagnose.
   Note: Your choice does NOT affect the display of NIC data in any other pages.

Windows Applications Metrics

Spotlight on Windows allows you to single out individual processes (or a group of processes) on a Windows system for particular analysis. The details of that analysis are displayed in the Single Application drilldown.

Use Windows Applications Metrics to choose one or more running processes for analysis.

To configure Windows Applications Metrics

1. Click **View | Options**
2. Click **Data Collection | Windows Applications Metrics**
3. Configure the following options:

Option	Description
Processes	Select the processes on the Windows system that you want to analyze.
CPU	Select the CPUs on the Windows system whose processes you want to analyze.
Disks	Select the disks on the Windows system whose processes you want to analyze.
NICs	Select the NICs on the Windows system whose processes you want to analyze.

Windows Event Log Metrics

The Windows Event Log Metrics dialog allows you to define and edit the logical rules that govern which Windows events are displayed in the Event Log drilldown, and allows you to configure a number of Event Log settings.

To configure Windows Event Log Metrics

1. Click **View | Options**
2. Click **Data Collection | Windows Event Log Metrics**
3. The **Global Rules** tab is to the front. Configure this page.

Global Rules
Enable the keeping of event logs and specify the rules that govern the recording of event log data.

Option	Description
Enable Event Logs	Select to permit the display of event log data in the Event Log drilldown. No event log data is displayed while this is deselected.
List of rules to apply to event log items	The list of event log rules defined for the current connection. • Click on a rule to highlight it. • Select the checkbox associated with the rule to enable it. Disabled rules have no effect on the data displayed in the Event Log drilldown.
Add	Click to create an event log rule. See Event Log Rule Editor for more information.
Edit	Highlight a rule. Click to modify the highlighted rule. See Event Log Rule Editor for more information.
Delete	Highlight a rule. Click to remove the highlighted rule from the list.

Option	Description
Rename	Highlight a rule. Click to change the display name of the highlighted rule.
Copy	Highlight a rule. Click to create a copy of the highlighted rule. Usually this is done as the basis for creating a rule that is similar to the copied rule.
Enable All	Click to enable (select) all the rules in the list.
Disable All	Click to disable (deselect) all the rules in the list.
Move Up	Rules are applied in the order in which they appear in the list. Click to move the highlighted rule higher in the list, and so be applied earlier.
Move Down	Rules are applied in the order in which they appear in the list. Click to move the highlighted rule lower down the list, and so be applied later.

Connection

Specify the details of what events are recorded for the current Spotlight connection and how long they are kept.

Option	Description
Display events from the...	Select the logs to show in the Event Log drilldown. The options may include: • Application log • Directory Service log (Windows domain controllers) • DNS Server log (Windows DNS servers) • File Replication Service log (Windows domain controllers) • Security log • System log The Directory Service and File Replication Service event logs are available only for Windows 2000 domain controllers. The DNS Server event log is available only for DNS servers. Under Windows 2000, the only logs available for view are the Application log, Security log, and System log.
Limit each event log to a maximum of...	Limit how long events are displayed in the Event Log drilldown before being deleted.

Windows Files Metrics

Track the growth of specified files (usually log files) in your Windows system. Spotlight reports an alarm when a page approaches the maximum file size you have set.

To open the Windows Files Metrics dialog

1. Click **View | Options**

2. Click **Data Collection | Windows Files Metrics**

About the dialog

Option	Description
File list	The list of files whose size is monitored for the current connection. There are two columns in this table: the File name and the Maximum size you have specified for the file. **Note**: The maximum file size listed on this page does NOT prevent the specified file from growing beyond that size. However, if the file exceeds that maximum size, Spotlight raises a critical alarm.
Add	Click to add a file to the list of those tracked by Spotlight. See Define Files to Monitor in Windows Files Metrics for more information.
Edit	Highlight a file. Click to change the settings for the highlighted file. See Define Files to Monitor in Windows Files Metrics for more information.
Delete	Highlight a file. Click to remove the highlighted file from the list.

Troubleshooting Spotlight on Windows

This section identifies general problems that you may encounter when using Spotlight on Windows, and details how to address those problems. If you do not find a solution in this section then check the *Spotlight on Windows Release Notes*.

The Registry ... is not accessible.

On connecting to a Windows machine, if Spotlight returns the following error:

> The registry on "<<MachineName>>" is not accessible. You need to connect to this machine with an account with privileges to retrieve server information, query the registry, and access performance monitor objects.

Solution

Login to the Windows machine using an account with administrative rights.

A login with administrative rights satisfies the requirements of Spotlight on Windows, with the privileges to retrieve server information, query the registry, and access WMI and performance monitor objects.

Network Name (Path) Not Found

On connecting to a Windows 2000 machine, if Spotlight returns the following error:

> Network name not found
> or
> Network path not found

Solution

Ensure the Remote Registry service is running on the Windows machine.

Spotlight on Windows will return this error if you attempt to connect to a Windows 2000 machine that has the Remote Registry service stopped.

Connecting As The Logged-in User

To connect to a Windows machine as a user who is already logged in, do NOT enter any user credentials (Domain, User or Password) in the Connection Properties window.

Notes:

- Spotlight on Windows requires Administrator access on all the Windows machines under diagnosis. If the current user is not an administrator on a machine, that Spotlight connection will fail.

- Windows connections can only have one set of credentials in use at a time. If a user connects to a machine by using NO user name and password, and then disconnects and reconnects with an Administrator password, Windows returns an ERROR_SESSION_CREDENTIAL_CONFLICT error.

  For example, if you have mapped a drive to the machine \\serverA\share while logged in as serverA\administrator, you CANNOT map another drive to the same machine while logged in as serverA\user.

  You can, however, connect to the same machine with different credentials if you map to the machine via its IP address.

  In the example above, if serverA has the IP address 192.168.1.100, you CAN employ the net use command to connect to the machine \\192.168.1.100\share as serverA\user.

WMI Connection Problems

If Spotlight on Windows is displaying a WMI access denied error, and the user specified in the connection profile is an administrator on that machine then it may be necessary to adjust your DCOM or Firewall settings.

Complete the following steps to ensure that your DCOM and Firewall settings are configured to allow WMI connections.

To test WMI is working on the Spotlight client:

1. Click **Start | Run**.

2. Enter **WBEMTest** to run the WMI Tester Utility.

3. Click **Connect | Connect**.

4. If an **Access is denied** error is displayed you do not have rights to make WMI connections. In this case, check (and adjust if necessary) your local DCOM setting (see below).

To check the DCOM access permissions for the monitored machine:

1. Login to the monitored machine.

2. Click **Start | Run**.

3. Enter **DCOMCnfg** to run the DCOM configuration utility.

4. Select **Component Services**.

5. Click .

6. Click **COM Security**.

7. Under Access Permissions, click **Edit Permission**.

   Ensure **Remote Access** is set to **Allow** for the user specified in the connection profile (or the group containing this user).

8. Under Launch and Activation Permissions, click **Edit Permission**.

   Ensure **Remote Launch** and **Remote Activation** are set to **Allow** for the user.

   For more information, see Securing a Remote WMI Connection on the MSDN site.

To check if Windows Firewall is active on the Spotlight client:

1. Open the **Control Panel | Windows Firewall**.

2. If Windows Firewall is on, see Connecting Through Windows Firewall on the MSDN site for information on how to configure the firewall to allow WMI connections.

For further information, see How to troubleshoot WMI-related issues in Windows XP SP2 on the Microsoft site.

Display Problems on The Home Page Disks Panel And Disk Drilldown

The Disks panel on the Spotlight on Windows home page and the Disks drilldown are not displaying data. Ensure disk performance counters are enabled on the Windows machine being monitored.

Solution: Enable collection of disk data:

On the Windows system being monitored

1. Start a command line window

2. At the command prompt, type the following

   diskperf -y

3. Restart the system.

Display Problems on The Home Page Network Panel, Flows And Network Drilldown

Spotlight on Windows cannot display data in the Network panel, flows, and drilldowns. The likely cause is that the appropriate performance counters have not been enabled on the Windows machine being diagnosed.

Solution: Enable the collection of network data:

- Check that the **PerfNet** counters are active on the system. To do so, use the **Exctrlst** utility, which you can download from the Microsoft Web site.

- Check that at least one network device is using NBT (NetBIOS over TCP/IP). To do so, check the properties of all network connections (in particular, **Advanced TCP/IP Settings | WINS**) to ensure that the NetBIOS setting is not disabled.

- Check that Service Pack 4 is installed on Windows 2000 systems.

The Processes Page And Process Affinity

In Windows systems where there are multiple processors, an activity (thread) in a process can migrate from processor to processor - and every migration reloads the processor cache.

Process affinity, which you can set on the target Windows machine via the Windows Task Manager, allows you to establish an association that permits specified processes to run ONLY on nominated processors. This limits the number of cache reloads, which may be desirable on systems operating under heavy load.

You CANNOT set process affinity via Spotlight on Windows.

What you CAN do is:

Use the Processes drilldown to indicate the CPUs where the specified process is permitted to run. This enables Spotlight to report the CPU usage for that process with accuracy.

Example: Running a process on an eight-CPU system

Spotlight uses the Windows Performance Monitor (PerfMon) to retrieve important data from target Windows systems, including the total percentage of CPU usage by specified processes.

On a single-CPU system, PerfMon can report a maximum CPU usage of 100%.

On a multiple-CPU system, that maximum value is multiplied by the number of CPUs. For an eight-CPU system, PerfMon can report a maximum CPU usage of 800%. Spotlight on Windows compensates by dividing the reported PerfMon value by the number of CPUs on the system.

However, when the affinity for a process is set so that it can run only on two nominated CPUs of the eight available, the default Spotlight calculation fails. For example, if PerfMon reports a CPU usage of 120% for the process (corresponding to a 60% usage on each permitted CPU), Spotlight initially calculates a value of 15%. Use Spotlight's Configure Affinity function to adjust the Spotlight calculation.

How to set process affinity

Process affinity can be set via the Windows Task Manager on Windows machines that have multiple CPUs.

1. Log in to the monitored Windows machine.

2. Press CTRL+ALT+DEL.

3. Click **Task Manager | Processes**.

4. Right-click the name of the process whose affinity you want to set and select **Set Affinity**.
   Note: This option is unavailable when the system has one CPU.

5. Select which of the available CPUs you want to use for the process.

Note: You can also set process affinity for Microsoft SQL Server through the SQL Server Enterprise Manager.

Installing The SNMP Protocol

Spotlight is unable to display Open Ports information if the SNMP protocol has not been installed on:

- The Windows machine under diagnosis, AND

- The machine running the Spotlight client.

To install the SNMP protocol (Windows XP)

1. Open the **Control Panel | Add or Remove Programs | Add/Remove Windows components**.

2. Select **Management and Monitoring Tools** from the list of components.

3. Click **Details**.

4. Select **Simple Network Management Protocol**.

5. Follow the installation instructions.

Product Authorization Errors

Spotlight on Windows is embedded with a number of other Spotlights, including Spotlight on SQL Server, Spotlight on Active Directory and Spotlight on Exchange.

The embedded version of Spotlight on Windows shares the license of the other Spotlight product. You do not need to license it separately. If Spotlight on Windows reports that its license has expired this means the license of the product Spotlight on Windows shipped with has expired. You can renew the Spotlight on Windows license by renewing the license of the other Spotlight product.

Welcome

Spotlight is powerful diagnostic and problem-resolution tool for Unix and Linux operating systems. Its unique user interface provides you with an intuitive, visual representation of the activity on your host machine.

Spotlight Home Page

The Spotlight home page shows the flow of information and commands between various sub-components and the size and status of internal resources such as processes, disk files and memory structures.

Related operating system statistics are grouped together on panels that are connected by a series of graphical flows and icons. Spotlight updates these flows in real time so you can see how quickly data is moving through the system. The icons change color as Spotlight alarms are raised, upgraded, downgraded and canceled.

The panels on the Spotlight home page are described in more detail next. For a full discussion of the panels, see the Spotlight online help.

Tip: Hover the mouse pointer over a panel component for more information.

What is Spotlight

Spotlight is powerful diagnostic and monitoring tool for Unix operating systems. Its unique user interface provides you with an intuitive, visual representation of the activity on your host machine. Graphical flows illustrate the rate at which data is moving between system components. Icons display the value of key statistics and measurements (metrics).

The power of Spotlight lies in its ability to provide visual and audible warnings if the performance metrics exceed acceptable thresholds. The components and dataflows change color to show you the source of the problem.

A range of visual graphs and tabular grids provide you with detailed information about your Unix hosts. This information can be viewed on the screen or printed.

You can set Spotlight to warn you when a threshold is reached. You may set a number of thresholds so that warning messages are displayed well before the traffic levels into or out of a host become critical. Spotlight uses a number of different techniques to warn you when a Unix host is exceeding a threshold.

When Spotlight detects a condition that it considers is a potential problem, it not only informs you about it, but also advises you what you could look at to diagnose the problem further, and suggests corrective actions.

Features of Spotlight

Some of the main features and benefits of Spotlight on Unix are that it:

- Provides a visual representation of process flows within Unix, allowing you to observe actual host activity in real time.
- Visually identifies bottlenecks and provides extensive drilldown capabilities.
- Displays the details of problem areas, including CPUs, disks, processes, and system calls statistics for rapid problem identification.
- Provides visual and audible warnings to alert you when performance metrics exceed acceptable thresholds.
- Provides detailed information about specific components through the use of drilldowns, therefore allowing you to pin point the source of problems.
- Assesses the normal rate of process flows via a calibration process, and sets the display speed of the visual indicators accordingly.
- Is easy to install.

Configure the Unix System

To monitor a Unix system with Spotlight, ensure the Unix system is configured as follows.

Unix servers and versions.

Spotlight supports the following server operating systems:

- Sun Solaris 9, 10, 11
- HP-UX 11i, 11i v2, 11i v3
- IBM AIX 5.2, 5.3, 6.1
- Red Hat and SUSE operating systems running Linux 2.4 and 2.6 kernels

Unix programs.

Ensure the following Unix programs are accessible to the Unix login (through Spotlight):

- Perl 5.x
- awk
- cat
- date
- df
- grep
- ifconfig
- iostat (not HP-UX)
- netstat

- ps
- sar
- sed
- tr
- uname
- uptime
- vmstat
- wc
- who

For HP-UX, additionally

- bdf
- cstm
- getconf

- swapinfo
- /usr/sbin/ioscan

For AIX, additionally

- lsattr
- lsdev

- lsps
- prtconf

For Solaris, additionally

- /etc/swap
- mpstat

- nawk
- prtconf

Unix User Permissions and Installation Settings

- The Unix user should have no special processing on log-on. In particular there must be no input required from the user, and nor should any special login banners be displayed.
- On AIX, the user must be a member of the "adm" group to be able to run the **sar** command.
- On Linux, the **sysstat** package must be installed to enable the user to get detailed disk I/O information.
- On Linux the **/proc** filesystem must be present.
- For connection using SSH, the **sshd** daemon must be installed and running.

Remote Connectivity: SSH or REXEC

Spotlight on Unix will require you to select the connection type: SSH or REXEC. Information on SSH and REXEC is freely available on the Internet. We recommend SSH as password data is transmitted encrypted. REXEC does not encrypt password data.

Notes:

- Make sure the relevant SSH or REXEC daemon is running on the Unix machine and is configured to receive remote connections.
- Commands to observe system activity (for example, **netstat, vmstat, iostat, sar**) must be accessible to REXEC / SSH sessions for Spotlight to observe Unix activity. Ensure these commands are located in the search path for REXEC / SSH sessions. If not, Spotlight will display an error.

Notes (Specific to SSH):

- Spotlight supports both SSH1 and SSH2 protocols.
- To allow Spotlight to make SSH connections to any Unix or Linux hosts that permit SSH connections, you may need to alter the **PasswordAuthentication** configuration item in the **sshd_config** file. Set the value of **PasswordAuthentication** to **yes**. Once you have modified the **sshd_config** file you must restart the **sshd process** to apply the new setting.
- Public-key encryption is supported under SSH2 only. DSA and RSA are supported.

Note (Specific to REXEC): When Spotlight is monitoring a Unix operating system via **REXEC** with a valid user ID and password, remote commands may not work on the Unix host unless that user ID is added to the **/etc/hosts.equiv** file on the host.

Connect to Unix Systems

1. From the Spotlight Browser select **All Connections | Spotlight on Unix**

2. Select the machine to connect to. If the Unix machine is not listed then follow the instructions to add a new connection.
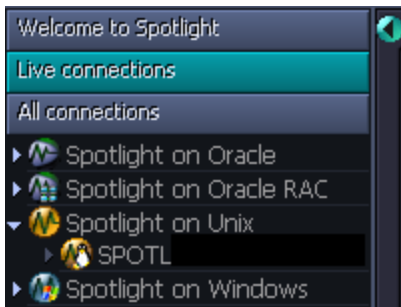
Add a new connection

Note: Ensure the Unix system is configured before you connect to it. Configure the Unix System

1. Click **File | Connect**

2. Select **Spotlight on Unix** on the Connections menu.

3. Double-click **Add new connection**.

4. Fill in the **Details** page of the Properties window:

Field	Description
Connection name	Enter the preferred display name for the Unix machine. **Note**: If left blank, the Connection name field resets to the value of the Address field.

Server Connection Details

Field	Description
Address	Enter the hostname or IP address to the Unix machine.
User	Enter the user name to login to the Unix machine. **Note:** The root user is not allowed. Remote login as "root" has been disallowed for security reasons.
Password	Enter the password to login to the Unix machine. **Note:** Not applicable when **Use SSH Public / Private Keys** is selected. Applicable when **Save password details (for this connection)** is selected.
Connection Type	Select according to the Unix system's configuration. The options are: REXEC and SSH. Configure the Unix System
Port Number	Enter the port number for SSH (secure shell) connection to the Unix machine. The default value is *22*. **Note:** Applicable when the **Connection Type** is **SSH**.

SSH Key Authentication
Note: Applicable when the **Connection Type** is **SSH**.

Field	Description
Use SSH Public / Private Keys	Select according to the Unix machine's configuration.
SSH Key Type	Choose the type of key to use when making the Spotlight connection: RSA or DSA. Configure the Unix System
SSH Private key Filename	Locate the file that contains the private key for the Spotlight connection. Either click [...] or type the filename and location.
SSH Passphrase	Type the passphrase used to decrypt the private key. **Note:** Applicable when **Save password details (for this connection)** is selected.

Select **Save password details (for this connection)** to save all entered password details.

5. Close the dialog.

Click	Description
Connect	Save changes and open the connection in Spotlight. **Note**: If you experience problems connecting there may be an issue with the Unix configuration. Configure the Unix System
OK	Save changes. Do not open the connection in Spotlight.
Cancel	Do not save changes.

Spotlight Home Page

The Spotlight home page shows the flow of information and commands between various sub-components and the size and status of internal resources such as processes, disk files and memory structures.

Related operating system statistics are grouped together on panels that are connected by a series of graphical flows and icons. Spotlight updates these flows in real time so you can see how quickly data is moving through the system. The icons change color as Spotlight alarms are raised, upgraded, downgraded and canceled.

The panels on the Spotlight home page are described in more detail next. For a full discussion of the panels, see the Spotlight online help.

Tip: Hover the mouse pointer over a panel component for more information.

System Panel

The **System** panel indicates the status of the operating system:

- The type and version of operating system.
- The length of time since the machine was last booted.
- The date of the last data collection. (system date)
- The time of the last data collection. (system time)

CPU Panel

The **CPU** panel shows processor and load information for the Unix system:

- The number of CPUs in the machine.
- The CPU load (as a percentage) across all CPUs on the machine.
- The percentage of time the CPU is in the **User** state across all CPUs.
- The percentage of time the CPU is in the **System** state across all CPUs.
- The percentage of time the CPU is in the **Wait** state across all CPUs. **Wait** is where the machine can only wait for I/O, or something similar, to complete.
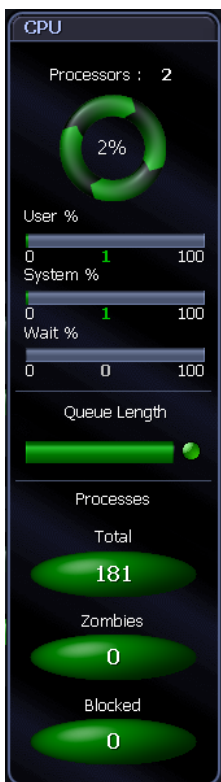- The length of the run queue in which processes are waiting to be executed. Processes in this queue will be run when the CPU becomes available.
- The total number of processes that are running on the machine.
- The number of "zombie" processes: child processes whose termination has not been acknowledged by their parent process.
- The number of processes waiting for some event or condition before they can continue execution.

The flows between the **CPU** panel and the **Memory** panel represent paging information between CPU and memory on the Unix host. The flows indicate the rate data is written out of and read in to memory. The unit of measure is dependent on the Operating System. See "CPU Panel" in the online help for more information.

Memory Panel

The **Memory** panel displays detailed information about the physical and virtual memory on the Unix system. It shows statistics such as:

- The total amount of physical memory in RAM.
- The amount of physical memory that is free.
- The total amount of virtual memory for the system.
- The amount of virtual memory that is free or unused.
- The amount of virtual memory consumed, expressed as a percentage of its maximum size.
- Where applicable, the number of processes whose pages have all been moved into swap space.

The flows between the **Memory** panel and the **Swap Space** panel represent the rate at which processes are being swapped out to disk and swapped in from disk.

Swap Space Panel
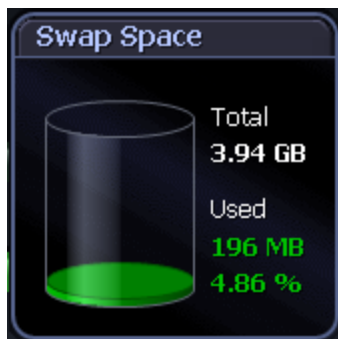
The **Swap Space** panel shows the configured swap space on the Unix system:

- The total amount of configured swap space on the system.
- The amount of swap space currently in use.

Network Panel

The **Network** panel shows the logins, connections and movement of data on the Unix system:

- The number of users currently logged on to the machine.
- The number of external programs with which the machine is currently communicating. The three types of connections shown are **Established**, **Time_Wait**, and **Close_Wait**.
- How fast is data moving through the NICs (network interface cards). The statistics refer to the rate at which network packets and errors are being received and sent through the NICs.

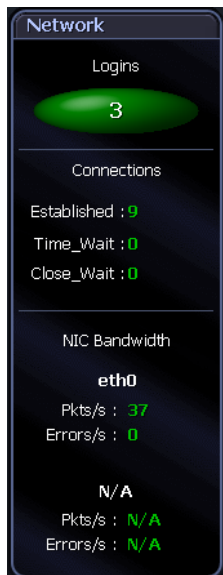The flows between the **Network** panel and the **CPU** panel represent the rate of data transfer between the Unix host and the connected network. The flows include the incoming and outgoing rates of data to and from the Unix host, and the rates at which bad packets are received from, or sent to, the network.

Disk Activity Panel
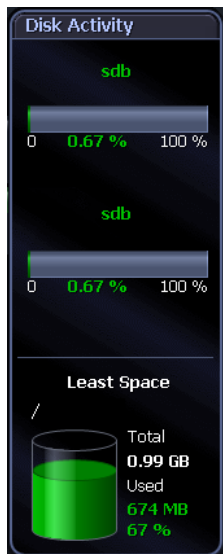
The **Disk Activity** panel shows information on disk usage:

- Activity gauges that display the percentage activity of the disks or partitions with the highest and second-highest read/write rates.
- The name of the mount point of the disk with the least amount of free space that is not already 100% full.
  Note: You can configure this control to display the total and used space on ANY available disk on the system. Click **View | Options**. For more information see the online help.
- The total size of this disk.
- The amount and percentage of disk space that is currently being used by this disk.

The flows between the **Disk Activity** panel and the **Memory** panel represent the number of writes the system makes to the disks per second, and the number of reads the system makes from the disks per second.

Spotlight Alarms

Spotlight alerts you to problems with your system by issuing an alarm. You can configure Spotlight in the level of severity that constitutes an alarm, to disable an alarm, and the actions Spotlight takes on raising the alarm.

When an alarm is raised Spotlight performs one or more of the following actions:

- Changes the color or intensity of relevant components.
- Gives audible warnings.
- Sends email notifications.

Actions you can take on an alarm being raised include:

- View details of the statistics that caused the alarm in a Spotlight drilldown page. Spotlight Drilldowns
- View the Spotlight online help.
- View details of the alarm in the Spotlight **Alarm Log** drilldown.
- Snooze the alarm.
- Save the alarm details.
- Filter the list of alarms.

Following are the alarms specific to Spotlight. For possible solutions to the problems indicated by these alarms or for information on how to diagnose problems further, see Spotlight alarms in the online help.

Alarm	Description
Blocked Processes alarm	Blocked processes are processes that are waiting on information to become available in memory, or that are waiting for a shared resource, or for the completion of disk I/O or network I/O. A high level of blocked processes may indicate an under-configured machine.
Input Error Packets Flow alarm	The Input Error flow represents the number of erroneous network packets coming into the machine per second. The **Input Error Packets Flow** alarm may warrant investigation as even a low error rate can indicate network problems.
Output Error Packets Flow alarm	The Output Error flow represents the number of erroneous network packets sent by the machine per second. The **Output Error Packets Flow** alarm may warrant investigation as even a low error rate can indicate network problems.
Page In alarm	A large number of page-ins may be a symptom of a large number of recent page-outs (see Page Out alarm below). To find out what process is reading these pages in, check the resident size of the processes you are running, and see which ones are increasing in size. Alternatively, a large number of page-ins may occur when a large process has just stopped, and the kernel decides to fill that space with memory pages owned by other processes. Check if an important process has ended unexpectedly.
Page Out alarm	If an active process asks the kernel for more memory than there is immediately available, the kernel will write old memory pages out to swap space. This is known as paging. To stop paging, make sure that there is enough RAM available to support the size of the processes you want to run.
Percentage Busy alarm	The **Most Active Disk** gauge shows the disk with the highest read/write rates. The **Percentage Busy** alarm is triggered when the read/write rate of the most active disk exceeds a specified threshold. The type of alarm that is activated is determined by the percentage of read/write activity experienced by the disk.
Swap Space	If the total amount of swap space allocated to a Unix host becomes full, the machine may halt all processes, or critical actions

Alarm	Description
alarm	
Swap In Flow alarm	The Swap In flow represents the number of processes swapped from disk per second. The **Swap In Flow** alarm may indicate a machine that is under-configured for its workload.
Swap Out Flow alarm	The Swap Out flow represents the number of processes swapped to disk per second. The **Swap Out Flow** alarm may indicate a machine that is under-configured for its workload.
Used Percentage alarm	The **Least Space / Disk Space** container displays the percentage of used disk space for a specified filesystem. The **Used Percentage** alarm indicates that the specified filesystem is full or filling up. If the filesystem fills up completely, no more data can be written to it.
Zombie Processes alarm	The **Zombie** button in the CPU panel shows the number of terminated child processes that have not been acknowledged by their parent process. The **ZombieProcesses** alarm may indicate the presence of an inferior program, or an entry in **/etc/inittab** (or equivalent) that prevents **init** from completing its startup sequence.

Blocked Process Alarm

The Blocked Processes button in the **CPU panel** shows the number of processes waiting on information to become available in memory. Processes can also become blocked when they are waiting for:

- disk I/O completion,
- network I/O completion, or
- a shared resource.

A high level of blocked processes may indicate an under-configured machine.

The type of alarm that is activated is determined by the number of blocked processes waiting on information.

Related Topics
CPU Panel

Input Error Packets Flow Alarm

The Input Error Flow represents the number of erroneous network packets coming in to the machine per second.

The **Input Error Packets Flow** alarm may warrant investigation as even a low error rate can indicate network problems.

The type of alarm that is activated is determined by the number of erroneous network packets coming in to the machine every second.

Related Topics
Network Panel

Output Error Packets Flow Alarm

The Output Error Flow represents the number of erroneous network packets sent by the machine per second.

The **Output Error Packets Flow** alarm may warrant investigation as even a low error rate can indicate network problems.

The type of alarm that is activated is determined by the number of erroneous network packets sent by the machine every second.

Page In Alarm

A large number of page-ins may be a symptom of a large number of recent page-outs. To find out what process is reading these pages in, check the resident size of the processes you are running, and see which ones are increasing in size.

A large number of page-ins may also occur when a large process has just stopped, and the kernel decides to fill that space with memory pages owned by other processes. Make sure that an important process hasn't died unexpectedly.

Page Out Alarm

If an active process asks the kernel for more memory than there is immediately available, the kernel will write old memory pages out to swap space. This is known as **paging**.

To stop paging, make sure that there is enough RAM available to support the size of the processes you want to run.

Percentage Busy Alarm (Most Active Disk)

The **Most Active Disk** gauge in the **Disk Activity** panel shows the disk with the highest read/write rates.

The **Percentage Busy** alarm is triggered when the read/write rate of the most active disk exceeds a specified threshold.

A disk that is more than 20% busy (according to *Sun Performance and Tuning* by A. Cockroft) should be investigated, and possibly have data on it split up and moved to different disks.

The type of alarm that is activated is determined by the percentage of read/write activity experienced by the disk.

Swap Space Alarm

The **Swap Space** panel on the main Spotlight on Unix window shows the total amount of swap space allocated to the machine you are diagnosing, and the amount of swap space currently in use.

If the total amount of swap space allocated to a Unix host becomes full, the machine may halt all processes, or critical actions may be prevented from occurring. Running out of swap space may indicate a runaway process or an under-configured machine.

The type of alarm that is activated is determined by the percentage of swap space currently in use.

Swap In Flow Alarm

The Swap In Flow represents the number of processes swapped from disk per second.

A machine that is swapping processes to or from disk is usually under-configured for its workload.

The type of alarm that is activated is determined by the number of processes swapped from disk every second.

Swap Out Flow Alarm

The Swap Out Flow represents the number of processes swapped to disk per second.

A machine that is swapping processes to or from disk is usually under-configured for its workload.

The type of alarm that is activated is determined by the number of processes swapped to disk every second.

Used Percentage Alarm

The **Least Space / Disk Space** container in the **Disk Activity** panel displays the percentage of used disk space for the specified filesystem.

The **Used Percentage** alarm indicates that the filesystem is full or filling up. If the filesystem fills up completely, no more data can be written to it. For example:

- If the /var filesystem fills, you may have problems with system activities such as e-mail or printing.

- If the /home filesystem fills, user applications may be affected.

- If the /tmp filesystem fills, any applications that use temporary files may be affected.

Notes:

- The data displayed in the **Least Space / Disk Space** container is collected from:

    - The filesystem that has the least amount of free space, OR

    - A user-specified filesystem on the Unix host.

- To choose the filesystem whose details are displayed in the container, right-click the container and select **Disk Options**.

Related Topics
Disk Activity Panel

Zombie Processes Alarm

The Zombie button in the **CPU panel** shows the number of terminated child processes that have not been acknowledged by their parent process.

A high level of zombie processes is indicative of the presence of an inferior program, or an entry in /etc/inittab (or equivalent) preventing init from completing its startup sequence.

The type of alarm that is activated is determined by the number of zombie processes waiting to be exited.

Related Topics

CPU Panel

Spotlight Drilldowns

When you have isolated a problem, you can display a drilldown page, whose charts and tables provide a detailed breakdown of the underlying statistics.

Following are Spotlight drilldowns.

Drilldown	Description
Processes and Services	The pages in the **Processes and Services** drilldown list all the processes running on the Unix machine (including "zombie" processes), and all the services found in the **/etc/services** file (running or not). Useful features in this drilldown include: • **Processes page**: View processes in a standard table or tree structure. In the tree structure secondary processes are shown as "children" of the processes that run them. Click on a process for more detailed information on that process including CPU and Memory utilization. From the Memory Usage chart it is possible to detect memory leaks of the process over the display period. • **Zombies page**: View child processes whose termination has not been acknowledged by their parent process. • **Services page**: Use a shortcut menu option to test whether a specified service is available.
Activity Summary	The pages in the **Activity Summary** drilldown – **Activity, Disk, Memory, CPU, Network**, and **Logins** – show details of the different types of activity on the Unix system.
Filesystems	The **Filesystems** drilldown displays detailed information about the filesystems on the Unix machine. This information can be viewed in chart or table form: right click over the drilldown page.
Alarm Log	The **Alarm Log** drilldown displays information on the alarms associated with the Unix machine.

Processes Page

The **Processes** page lists the processes currently running on the Unix system.

Note: The **Processes** page shows the output from the Unix ps command.

To open the Processes page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Processes | Processes**.

For each process, you can view the information that follows:

Notes:

- Some columns in the table are hidden by default. To view hidden columns, right-click a column heading and choose **Organize Columns…** from the shortcut menu.

- Refer to man ps(1) for further details on the values shown in the **Processes** grid.

- When the table is displayed as a flat grid (not a tree), right click any value in the table and select **Show processes where…** or **Hide processes where…** to filter the table by that value.

Column	Description
Process	The name of the process. **Notes:** • Click on a process to display more information on that process. Process Details Page • Right click on the process and select **Kill Process** (if enabled) to end any processes currently running on the Unix machine. • Some of the processes may be running other, secondary, processes. Right click the grid and select **Show as Tree** to display the secondary processes as branches of the initial process.
Nice	The Nice value describes the relative priority of the specified process. A process with a low Nice value is running at a higher priority than a process with a high Nice value.
PID	The process identifier for the specified process.
PPID	The process identifier for the process that is the parent of the specified process.
% CPU	The percentage of CPU time used by the process in the last sample interval.
State	The state of the process. Process states are platform dependent.

Platform	Value	Description
AIX	A	Active
	W	Swapped
	I	Idle
HP-UX	W	Waiting
	I	Intermediate
	X	Growing
Linux	D	Uninterruptible Sleep
	W	Has no resident pages
	<	High priority process
	N	Low priority tasks
	L	Has pages locked into memory
Solaris	O	Process is running on a processor
	S	Sleeping. That is, the process is waiting for an event to complete.
	R	Runnable. That is, the process is on the run queue.
	Z	Zombie state. That is, the process has been terminated and the parent process is no longer waiting.
	T	The process has been stopped by a job control signal, or because it is being traced.

Column	Description
Terminal	The Unix terminal session where a user started the specified process. If the process was not started by an interactive user, the Terminal value is set to "?".
User	The name of the user to whom the process belongs.
Time	The amount of CPU time the process has consumed.
Virt Mem (MB)	The amount of virtual memory in use by the process, measured in megabytes.

Process Details Page

The **Process Details** page contains a detailed list of properties for the process selected on the **Processes** page.

To open the Process Details page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Processes | Processes**.

3. Select a process to view its details

For the selected process

The **Memory Usage** chart displays data series for the virtual memory and physical memory used by the selected process.

The **CPU Usage** chart displays CPU usage for the selected process.

The information grid is as follows.

Item	Description
Process	The name of the process.
Process ID	The process identifier for the specified process.
User	The name of the user to whom the process belongs.
Parent PID	The process identifier for the process that is the parent of the specified process.
Group ID	The group identifier for the user that owns the specified process.
Processor %	The percentage of CPU time used by the process in the last sample interval.
Elapsed Time	How long the process has been running.
Cumulative Time	The amount of CPU time the process has consumed.
State	The state of the process. Process states are platform dependent.

Item	Description		
	Platform	**Value**	**Description**
	AIX	A	Active
		W	Swapped
		I	Idle
	HP-UX	W	Waiting
		I	Intermediate
		X	Growing
	Linux	D	Uninterruptible Sleep
		W	Has no resident pages
		<	High priority process
		N	Low priority tasks
		L	Has pages locked into memory
	Solaris	O	Process is running on a processor
		S	Sleeping. That is, the process is waiting for an event to complete.
		R	Runnable. That is, the process is on the run queue.
		Z	Zombie state. That is, the process has been terminated and the parent process is no longer waiting.
		T	The process has been stopped by a job control signal, or because it is being traced.
Terminal	The Unix terminal session where a user started the specified process. If the process was not started by an interactive user, the Terminal value is set to "?".		
Priority	The basic priority assigned to the process – the lower the number, the higher the priority. Unix can modify this priority by means of the renice command.		
Nice	The Nice value describes the relative priority of the specified process. A process with a low Nice value is running at a higher priority than a process with a high Nice value.		
Virt Mem (MB)	The amount of virtual memory in use by the process, measured in megabytes.		
Physical Mem (MB)	The amount of physical memory in use by the process, measured in megabytes.		
Command	The command executed by the process.		

Zombies Page

Unix expects a parent process to acknowledge the termination of any child process. If it fails to do so, the terminated child process is classified by the kernel as a **zombie**.

A high number of zombie processes indicates that one or more processes are not handling their child processes properly. You may need to **kill** the parent process to eliminate its zombie child process.

To open the Zombies page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Processes | Zombies**.

For each process, you can view the information that follows:

Column	Description
PID	The process identifier for the specified process.
PPID	The process identifier for the process that is the parent of the specified process.
UID	The user identifier for the user to whom the process belongs.
State	**Z** for Zombie. That is, the process has been terminated and the parent process is no longer waiting.
Priority	The basic priority assigned to the process – the lower the number, the higher the priority.
Nice	The Nice value describes the relative priority of the specified process. A process with a low Nice value is running at a higher priority than a process with a high Nice value.
CPU Utilization	A value representing the amount of CPU time used by the process. The metric used here may differ across Unix implementations.
Terminal	The Unix terminal session where a user started the specified process. If the process was not started by an interactive user, the Terminal value is set to "**?**".
Command	The command executed by the process.

Services Page

The **Services** page lists (by name) the popular services found in the /etc/services file of the Unix machine. These services may or may not be enabled.

To open the Services page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Processes | Services**.

The **Services** table contains the following information.

Column	Description
Service name	The name of the specified service.
Port number	The logical port used by the service to handle data.
Protocol	The transfer protocol used by the service. Possible protocols include TCP (Transmission Control Protocol) and UDP (User

Column	Description
	Datagram Protocol).
Aliases	Alternative names for the service.
Active?	The status of the port used by the service (initially blank). To display the status, right-click on a service in the table and select **Test port**. Possible values are: • Yes - A process on the server is listening on that port. • No - There is no process listening on that port. • Connecting... - Spotlight is waiting for a response from the server.

Activity Page

The **Activity** page contains several charts.

To open the Activity page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Activity Summary | Activity**.

Charts on the Activity page

Notes:

• Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

• Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

• Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
CPU Utilization	The **CPU Utilization** chart shows the CPU activity of the host machine. If the current activity reaches levels of greater than **100** or less than **0**, run **vmstat(1)** to determine the reason for the erroneous values.
Run Queue	The **Run Queue** chart shows the amount of CPU time that will be spent performing tasks. The value is expressed as a fraction. Divide the value by the number of processors running. If the result of the calculation is greater than **1**, you may need more processors. The lines on the chart show the run queue averaged over 1, 5, and 15 minutes. If the value of the **1 minute** line divided by the number of processors exceeds the value **1**, there is merely a spike in activity. However, if the value of the **15 minute** line divided by the number of processors exceeds **1**, you may need more processors.
Total Disk I/O	The **Total Disk I/O** chart shows the total number of I/O operations for each disk and NFS server.
Network Utilization	The **Network Utilization** chart represents the current network activity for the machine you are diagnosing. It shows the rate at which packets are being received by the machine and are being sent by the machine.

Chart	Description
Paging	The **Paging** chart shows the rate of data pages read from, and written out to, disk over time. Typically, high paging rates indicate insufficient system memory, a large number of processes, or a number of very large processes. The unit of measure on the Y-axis is dependent on the Operating System.

Operating System	Unit of Measure
Sun Solaris	Kilobytes per second
HP-UX	Pages per second
IBM AIX	Pages per second
Linux	Kilobytes per second. Old versions of the kernel may measure in blocks per second. Refer to your Linux documentation.

Chart	Description
Free Memory	The **Free Memory** chart shows the amount of available physical memory in the machine. For Solaris, this is typically a low number. If necessary, the page-stealing daemon will seek out memory pages to re-claim them.

Disk Page

The **Disk** page contains several charts.

To open the Disk page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Activity Summary | Disk**.

To view charts on the Disk page

Right-click the **Disk** page and select **View as Chart**.

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Disk Reads	The **Disk Reads** chart shows the rate at which data is being read off each disk.
Service Time	The **Service Time** chart shows the sustained service time for all disks on the machine you are diagnosing. If the sustained service time of any of your disks exceeds **50** milliseconds, move your data onto the under-utilized disks.
Disk Writes	The **Disk Writes** chart shows the rate at which data is being written to each disk.
Queue	The **Queue Length** chart shows the average number of write operations waiting to be completed for each disk.

Chart	Description
Length	
Wait Time	The **Wait Time** chart shows the amount of time that the queue for each disk is not empty. The time is expressed as a percentage of total time.
Busy	The **Busy** chart shows the amount of time each disk spends processing its transactions. The time is expressed as a percentage of total time.

To view the Disk page as a table

Right-click the **Disk** page and select **View as Grid**.

Memory Page

The **Memory** page contains several charts.

To open the network page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Activity Summary | Memory**.

Charts on the Memory page

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Physical	The **Physical Memory** chart shows the total physical memory (RAM) on the Unix host, and the amount of physical memory that is currently free.
Virtual	The **Virtual Memory** chart shows the total amount of memory in use on the Unix host. This includes physical memory plus space in the paging files. A steady increase in virtual memory usage can indicate that a process on the system has a memory leak.
Paging	The Paging chart shows the rate at which pages are being swapped in and out of memory. The chart displays three data series: - **Page In** - The rate at which **read** requests are made to the swap device. - **Page Out** - The rate at which **write** requests are made to the swap device. - **Scan rate** - The rate at which the memory manager scans for a free page in swap. If this rate is rising, it indicates that system memory is under pressure. A sustained high rate of paging can cause problems with overall system degradation due to disk thrashing and CPU load. **Note**: Scan rate is not available for Linux.
Cache	The **Cache** chart shows the percentage of recent file requests (**read** and **write**) that are satisfied by the memory cache, and that do not require a request to the paging areas on disk.

CPU Page

The **CPU** page contains several charts.

Charts on the CPU page

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.

- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.

- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Total CPU Usage	The total CPU utilization for all processors on the Unix host.
CPU Usage by Processor	The CPU utilization for each processor on the Unix host. Each processor is represented by a single line on the chart, and by an item in the legend at the top right of the chart.
Run Queues	The number of threads (program execution units) that are waiting to run on the Unix host.
Context Switching	The rate at which the Unix host has been processing context switches. A context switch occurs when a processor switches from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is pre-empted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to invoke a system service.

Network Page

The **Network** page contains several charts.

Charts on the Network page

Notes:

- Every chart has a **legend** (list of symbols) to its right that describes the various series (line graphs) on the chart.
- Click an item in the legend to highlight its series (line) in the chart. Click a second time to return the series to its normal appearance.
- Move the mouse pointer over an item in the legend to view the current value for that series within the chart.

Chart	Description
Network Utilization	The total rate of data traffic across all network cards (NICs) for the Unix host, measured in packets/second. Separate graphs on the chart indicate the rates at which data packets are being: • Received by the Unix host from the network (**Packets In**). • Transmitted by the Unix host to the network (**Packets Out**).
Connections	The number of external services with which the Unix host is communicating. The graphs on this chart represent three types of connections: • Established - The total number of TCP/IP connections to this machine in the ESTABLISHED state. • TIME_WAIT - The total number of TCP/IP connections to this machine in the TIME_WAIT state, where the local socket has closed, and it is waiting for the remote end to signal that it has done the same. • CLOSE_WAIT - The total number of TCP/IP connections to this machine in the CLOSE_WAIT state, where the remote end of the connection has shut down, and it is waiting for the local end to do the same.
Network Utilization by Kilobytes	The total rate of data traffic across all network cards (NICs) for the Unix host, measured in kilobytes/second. Separate graphs on the chart indicate the rates at which data is being: • Received by the Unix host from the network (**KB In**). • Transmitted by the Unix host to the network (**KB Out**).
Packets by Network Card	This chart shows the rate of data traffic for every network card (NIC) in the Unix host, measured in packets/second. For every NIC, a separate graph represents the data packets: • Received by the NIC from the network (**Packets In**). • Transmitted by the NIC to the network (**Packets Out**).
Error Rates by Network Card	This chart shows error rates for the data transferred by every network card (NIC) in the Unix host, measured in errors/second. For every NIC, a separate graph represents the errors in data: • Received by the NIC from the network (**Errors In**). • Transmitted by the NIC to the network (**Errors Out**).
Collisions by Network Card	A **collision** is the result of two devices trying to transmit data packets across the network at the same time. The network detects the collision of the two transmitted packets as an error, and discards them both. Both devices then need to re-send the data. This chart shows the rate at which collisions occur for every network card (NIC) in the Unix host.

Logins Page

The **Logins** page displays the users currently logged in to the selected connection.

To open the logins page

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Activity Summary | Logins**.

Data in the Logins page

For each login, you can view the information that follows.

Note: Not all of the columns in the Logins table may be visible by default. To view hidden columns, right-click the grid heading and select **Organize columns…**.

Column	Description
UserName	The name of a user logged in to the Unix host.
Line	The terminal device associated with this connection.
LoginTime	The date and time when the user logged in for this session.
Idle	The time elapsed since the user's last activity (if available).
PID	The Process ID of the connection or login.
Comments	Additional information about the connection, including the address from which the connection was made (if available).

Filesystems And Disk Information

The **Filesystems** drilldown displays information on the disk usage for each filesystem on the Unix system, and the proportion of disk space that has been used for each filesystem.

Note: Refer to man df(1) for further details on the values shown in the **Filesystems** page.

To open the Filesystems and Disk Information drilldown

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Filesystems**.

To view disk usage in a chart

Right-click the **Filesystems** page and select **View as Chart**.

The information is displayed on a set of bar graphs that represent used and free disk space. Use the drop-down **By Megabyte / By Percentage** control to specify how to view disk usage.

To view disk usage in a table

Right-click the **Filesystems** page and select **View as Grid**.

Column	Description
Filesystem	The Unix device that holds a filesystem.
Used Space	The amount of used space on the filesystem.
Free Space	The amount of free space on the filesystem.
Mounted On	The mount point for the filesystem.

To view disk usage for NFS mounted disks

NFS (Network File System) was developed by Sun Microsystems to allow a Unix host to use disk space on a remote machine as if it were local storage. Spotlight provides a shortcut option to allow you to show or hide disk usage for NFS mounted disks.

Right-click the **Filesystems** page and select **Show NFS mounted Filesystems**.

Alarm Log

The **Alarm Log** drilldown displays information on the alarms associated with the selected connection.

To open the alarm log

1. Select the **Spotlight on Unix** connection in the **Spotlight Browser**.

2. Click **Alarm Log**.

Notes:

- The Alarm Log drilldown is common to all Spotlight applications. See Spotlight Alarms for more information.
- The alarms are specific to the current Spotlight connection. Spotlight Alarms

Related Topics
Spotlight Alarms
Spotlight Alarms

Spotlight on Unix - View | Options

You can set a number of options that affect the behavior and display of Spotlight on Unix via the Spotlight on Unix Options window.

Each of the pages within the Spotlight on Unix Options window Disks and Other covers one aspect of Spotlight on Unix behavior.

To open the Spotlight on Unix Options window

Click **View | Options**.

Use this option…	To…
Unix Filesystem Metric	Control the display of disk information on the specified Spotlight on Unix drilldown pages.

Unix Filesystem Metric

Use the Unix Filesystem Metric option in the Spotlight on Unix Options window to configure how disks are represented in the Disk Activity panel on the Spotlight home page. The Disk Activity panel represents the logical disks in the system you are analyzing.

To set the Unix Filesystem Metric option

1. Select **View | Options**.
2. Select **Appearance | Unix Filesystem Metric**.
3. Choose one of the following options:

Option	Description
Display Fullest Filesystem	For the current Spotlight connection, choose this option to display the details of the disk with the least available space. When you choose this option, the associated label in the lower half of the Disk Activity panel displays the text Least Space.
Display Following Filesystem	Choose this option to select a specific disk. Use the drop-down list to choose which disk whose details will be displayed. When you choose this option, the associated label in the lower half of the Disk Activity panel changes to Disk Space.

Troubleshooting Spotlight on Unix

This section identifies general problems that you may encounter when using Spotlight on Unix, and details how to address those problems. If you do not find a solution in this section then check the *Spotlight on Unix Release Notes*.

Troubleshooting Spotlight on Unix

This section identifies general problems that you may encounter when using Spotlight on Unix, and details how to address those problems. If you do not find a solution in this section then check the *Spotlight on Unix Release Notes*.

Using Trace To Add Details To Support Bundle

When you contact Questsupport to seek assistance on a Spotlight issue, you can use **Spotlight tracing** to include additional information in the **support bundle** that you send to us.

To turn on Spotlight tracing

1. Close all open Spotlight on Unix connections.

2. Click **View | Options | Tracing**

3. Select **Turn client tracing on** to turn on tracing for the computer Spotlight on Unix is running on.

4. Select **Turn server tracing on** to turn on tracing for the Unix computer system.

5. Click **OK**.

6. Open the connection whose performance you want to trace.

With tracing turned on, Spotlight can collect additional information (such as the output from operating-system commands) for the rest of the current session.

When you create a support bundle with tracing turned on, the support bundle contains the additional information that Questcan use to investigate your support issue.

Note: All tracing information gathered for Spotlight on Unix is from the Spotlight client. There is no server tracing capability.

iostat Not Installed

Spotlight works with full functionality only when the **iostat** command is installed on the Unix system. This is NOT true for (HP-UX).

If **iostat** is NOT installed, Spotlight may display an error message such as **iostat: command not found**. If you encounter this error, you can:

- Choose to **ignore** the error. When you do so, the Spotlight error will no longer appear, but neither will a number of important disk statistics.

- Install the **iostat** command on the server. (On Red Hat Linux, this command is included in the **sysstat** package.)

Issues With Data Availability

Some data is unavailable for certain Unix configurations. See the *Spotlight on Unix Release Notes* for more information.

Authorize Spotlight

Spotlight is supplied with a time-limited trial key so you may test the product. When the trial key expires Spotlight reverts to a preview mode, which limits access to Spotlight. You will be able to see the Spotlight home page and one connection at a time. If you select to open a drilldown in preview mode the following message is displayed:

> The current license does not allow access to this functionality.

Contact your Quest representative to obtain the necessary authorization keys to ensure Spotlight remains fully functional when the trial key expires.

To enter a new authorization key:

While Spotlight is running

1. Click **Help | About Spotlight**.
2. Ensure **Installed Products** is to the front.
3. Locate Spotlight in the list of installed products.
4. Select **View / change product license**.
5. Click **Change this license**.
6. Enter the **Authorization key**.

Notes:

- The Windows user that updates the authorization key must have administrator access to the Windows registry and write access to the **Console\Licenses** sub-folder of the Spotlight installation folder.
- Make sure you enter the authorization key exactly as stated on your Product Authorization sheet or as provided by your Quest distributor. The site message is case sensitive. If you enter the authorization key incorrectly, the following message is displayed: Invalid authorization key
- If Spotlight is in preview mode and the trial period has not expired, check your computer system date is set correctly.

About us

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 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

Third-party contributions

Third-party contributions

This product contains some third-party components (listed below). Copies of their licenses may be found at referencing http://www.quest.com/legal/license-agreements.aspx. Source code for components marked with an asterisk (*) is available at http://opensource.quest.com.

Table 1: List of Third-Party Contributions

Component	License or Acknowledgement
SynEdit 1.0	Portions created by Martin Waldenburg are Copyright (C) 1998 Martin Waldenburg. All Rights Reserved. Mozilla Public License Version 2.0 ================================= 1. Definitions -------------- 1.1. "Contributor" means each individual or legal entity that creates, contributes to the creation of, or owns Covered Software. 1.2. "Contributor Version" means the combination of the Contributions of others (if any) used by a Contributor and that particular Contributor's Contribution. 1.3. "Contribution" means Covered Software of a particular Contributor. 1.4. "Covered Software" means Source Code Form to which the initial Contributor has attached the notice in Exhibit A, the Executable Form of such Source Code Form, and Modifications of such Source Code Form, in each case including portions thereof. 1.5. "Incompatible With Secondary Licenses" means (a) that the initial Contributor has attached the notice described in Exhibit B to the Covered Software; or (b) that the Covered Software was made available under the terms of version 1.1 or earlier of the License, but not also under the terms of a Secondary License. 1.6. "Executable Form" means any form of the work other than Source Code Form. 1.7. "Larger Work" means a work that combines Covered Software with other material, in a separate file or files, that is not Covered Software. 1.8. "License" means this document. 1.9. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently, any and all of the rights conveyed by this License. 1.10. "Modifications" means any of the following: (a) any file in Source Code Form that results from an addition to, deletion from, or modification of the contents of Covered Software; or (b) any new file in Source Code Form that contains any Covered Software. 1.11. "Patent Claims" of a Contributor means any patent claim(s), including without limitation, method, process, and apparatus claims, in any patent Licensable by such Contributor that would be infringed, but for the grant of the License, by the making, using, selling, offering for sale, having made, import, or transfer of either its Contributions or its Contributor Version. 1.12. "Secondary License" means either the GNU General Public License, Version 2.0, the GNU Lesser General Public License, Version 2.1, the

Component	License or Acknowledgement

GNU Affero General Public License, Version 3.0, or any later versions of those licenses. 1.13. "Source Code Form" means the form of the work preferred for making modifications. 1.14. "You" (or "Your") means an individual or a legal entity exercising rights under this License. For legal entities, "You" includes any entity that controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity. 2. License Grants and Conditions -------------------------------- 2.1. Grants Each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license: (a) under intellectual property rights (other than patent or trademark) Licensable by such Contributor to use, reproduce, make available, modify, display, perform, distribute, and otherwise exploit its Contributions, either on an unmodified basis, with Modifications, or as part of a Larger Work; and (b) under Patent Claims of such Contributor to make, use, sell, offer for sale, have made, import, and otherwise transfer either its Contributions or its Contributor Version. 2.2. Effective Date The licenses granted in Section 2.1 with respect to any Contribution become effective for each Contribution on the date the Contributor first distributes such Contribution. 2.3. Limitations on Grant Scope The licenses granted in this Section 2 are the only rights granted under this License. No additional rights or licenses will be implied from the distribution or licensing of Covered Software under this License. Notwithstanding Section 2.1(b) above, no patent license is granted by a Contributor: (a) for any code that a Contributor has removed from Covered Software; or (b) for infringements caused by: (i) Your and any other third party's modifications of Covered Software, or (ii) the combination of its Contributions with other software (except as part of its Contributor Version); or (c) under Patent Claims infringed by Covered Software in the absence of its Contributions. This License does not grant any rights in the trademarks, service marks, or logos of any Contributor (except as may be necessary to comply with the notice requirements in Section 3.4). 2.4. Subsequent Licenses No Contributor makes additional grants as a result of Your choice to distribute the Covered Software under a subsequent version of this License (see Section 10.2) or under the terms of a Secondary License (if permitted under the terms of Section 3.3). 2.5. Representation Each Contributor represents that the Contributor believes its Contributions are its original creation(s) or it has sufficient rights to grant the rights to its Contributions conveyed by this License. 2.6. Fair Use This License is not intended to limit any rights You have under applicable copyright doctrines of fair use, fair dealing, or other equivalents. 2.7. Conditions Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted in Section 2.1. 3. Responsibilities -------------------- 3.1. Distribution of Source Form All distribution of Covered Software in Source Code Form, including any Modifications that You create or to which You contribute, must be under the terms of this License. You must inform recipients that the Source Code Form of the Covered Software is governed by the terms of this License, and how they can obtain a copy of this License. You may not attempt to alter or restrict the recipients' rights in the Source Code Form. 3.2. Distribution of Executable Form If You distribute Covered Software in Executable Form then: (a) such Covered Software must also be made available in Source Code Form, as described in Section 3.1, and You must inform recipients of the Executable Form how they can obtain a copy of such Source Code Form by

mply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be placed in a text file included with all distributions of the Covered Software under this License. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it. 5. Termination -------------- 5.1. The rights granted under this License will terminate automatically if You fail to comply with any of its terms. However, if You become compliant, then the rights granted under this License from a particular Contributor are reinstated (a) provisionally, unless and until such Contributor explicitly and finally terminates Your grants, and (b) on an ongoing basis, if such Contributor fails to notify You of the non-compliance by some reasonable means prior to 60 days after You have come back into compliance. Moreover, Your grants from a particular Contributor are reinstated on an ongoing basis if such Contributor notifies You of the non-compliance by some reasonable means, this is the first time You have received notice of non-compliance with this License from such Contributor, and You become compliant prior to 30 days after Your receipt of the notice. 5.2. If You initiate litigation against any entity by asserting a patent infringement claim (excluding declaratory judgment actions, counter-claims, and cross-claims) alleging that a Contributor Version directly or indirectly infringes any patent, then the rights granted to You by any and all Contributors for the Covered Software under Section 2.1 of this License shall terminate. 5.3. In the event of termination under Sections 5.1 or 5.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or Your distributors under this License prior to termination shall survive termination. ** * * * 6. Disclaimer of Warranty * * ------------------------- * * * * Covered Software is provided under this License on an "as is" * * basis, without warranty of any kind, either expressed, implied, or * * statutory, including, without limitation, warranties that the * * Covered Software is free of defects, merchantable, fit for a * * particular purpose or non-infringing. The entire risk as to the * * quality and performance of the Covered Software is with You. * * Should any Covered Software prove defective in any respect, You * * (not any Contributor) assume the cost of any necessary servicing, * * repair, or correction. This disclaimer of warranty constitutes an * * essential part of this License. No use of any Covered Software is * * authorized under this License except under this disclaimer. * * * ** ** * * * 7. Limitation of Liability * * -------------------------- * * * * Under no circumstances and under no legal theory, whether tort * * (including negligence), contract, or otherwise, shall any * * Contributor, or anyone who distributes Covered Software as * * permitted above, be liable to You for any direct, indirect, * * special, incidental, or consequential damages of any character * * including, without limitation, damages for lost profits, loss of * * goodwill, work stoppage, computer failure or malfunction, or any * * and all other commercial damages or losses, even if such party * * shall have been informed of the possibility of such damages. This * * limitation of liability shall not apply to liability for death or * * personal injury resulting from such party's negligence to the * * extent applicable law prohibits such limitation. Some * * jurisdictions do not allow the exclusion or limitation of * * incidental or consequential damages, so this exclusion and * * limitation may not apply to You. * * * ** 8. Litigation ------------- Any litigation relating to this License may be brought only in the courts of a jurisdiction where the defendant maintains its principal place of business and such litigation shall be governed by laws of that jurisdiction, without reference to its conflict-of-law provisions. Nothing in this Section shall prevent a party's ability to bring cross-claims or counter-claims. 9.

Component   License or Acknowledgement